

The GNU Source-Level Debugger

Ninth Edition, for gdb version 7.0.50.20100218-cvs

(Sourcery G++ Lite 4.4-191)

Richard Stallman, Roland Pesch, Stan Shebs, et al.

(Send bugs and comments on gdb to

https://support.codesourcery.com/GNUToolchain/.)

Debugging with GDB TEXinfo 2008-04-18.10

Published by the Free Software Foundation 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301, USA ISBN 1-882114-77-9

Copyright © 1988, 1989, 1990, 1991, 1992, 1993, 1994, 1995, 1996, 1998, 1999, 2000, 2001, 2002, 2003, 2004, 2005, 2006, 2007, 2008, 2009, 2010 Free Software Foundation, Inc.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.1 or any later version published by the Free Software Foundation; with the Invariant Sections being "Free Software" and "Free Software Needs Free Documentation", with the Front-Cover Texts being "A GNU Manual," and with the Back-Cover Texts as in (a) below.

(a) The FSF's Back-Cover Text is: "You are free to copy and modify this GNU Manual. Buying copies from GNU Press supports the FSF in developing GNU and promoting software freedom."

This edition of the GDB manual is dedicated to the memory of Fred Fish. Fred was a long-standing contributor to GDB and to Free software in general. We will miss him.

Table of Contents

\mathbf{S}_{1}	umn	$\operatorname{nary} \operatorname{of} \operatorname{GDB} \ldots 1$
	Free	Software
	Free	Software Needs Free Documentation
		ributors to GDB
1	A	Sample GDB Session 7
_		Pro GED GEOGRAPIO
2	\mathbf{G}	etting In and Out of GDB 11
	2.1	Invoking GDB
	2.	1.1 Choosing Files
	2.	1.2 Choosing Modes
	2.	1.3 What gdb Does During Startup
	2.2	Quitting gdb
	2.3	Shell Commands
	2.4	Logging Output
3	GI	OB Commands
	3.1	Command Syntax
	3.2	Command Completion
	3.3	Getting Help
4	R	unning Programs Under GDB 25
•	4.1	Compiling for Debugging
	4.1	Starting your Program
	4.3	Your Program's Arguments
	4.4	Your Program's Environment. 29
	4.5	Your Program's Working Directory
	4.6	Your Program's Input and Output
	4.7	Debugging an Already-running Process
	4.8	Killing the Child Process
	4.9	Debugging Multiple Inferiors and Programs
	4.10	Debugging Programs with Multiple Threads
	4.11	Debugging Forks
	4.12	Setting a Bookmark to Return to Later
	4.	12.1 A Non-obvious Benefit of Using Checkpoints

5	Stop	ping and Continuing	. 43
	5.1 Bre	eakpoints, Watchpoints, and Catchpoints	43
	5.1.1	Setting Breakpoints	
	5.1.2	Setting Watchpoints	49
	5.1.3	Setting Catchpoints	51
	5.1.4	Deleting Breakpoints	55
	5.1.5	Disabling Breakpoints	55
	5.1.6	Break Conditions	56
	5.1.7	Breakpoint Command Lists	58
	5.1.8	"Cannot insert breakpoints"	
	5.1.9	"Breakpoint address adjusted"	59
	5.2 Con	ntinuing and Stepping	
	5.3 Sign	nals	63
	5.4 Sto	pping and Starting Multi-thread Programs	65
	5.4.1	All-Stop Mode	65
	5.4.2	Non-Stop Mode	67
	5.4.3	Background Execution	
	5.4.4	Thread-Specific Breakpoints	69
	5.4.5	Interrupted System Calls	69
7		ording Inferior's Execution and Replayi	_
8	Exar	mining the Stack	. 77
		ck Frames	
		cktraces	
		ecting a Frame	
	8.4 Info	ecting a Frame	80
9	8.4 Info	ecting a Frame prmation About a Frame	80
		9	80
	Exar	ormation About a Frame	80 81
	Example 9.1 Prin	mining Source Files	80 81 . 83 83
	Exar 9.1 Pri 9.2 Spe	mining Source Files	80 81 . 83 83 84
	Exar 9.1 Pri 9.2 Spe	mining Source Files nting Source Lines ecifying a Location	80 81 . 83 83 84 85
	Exar 9.1 Pri 9.2 Spe 9.3 Edi 9.3.1	mining Source Files nting Source Lines ecifying a Location ting Source Files	80 81 83 83 84 85 85
	Exar 9.1 Pri 9.2 Spe 9.3 Edi 9.3.1 9.4 Sea	mining Source Files nting Source Lines ecifying a Location ting Source Files Choosing your Editor	80 81 83 83 84 85 85

10 E	Examining Data	93
10.1	Expressions	. 93
10.2	Ambiguous Expressions	
10.3	Program Variables	. 95
10.4	Artificial Arrays	. 97
10.5	Output Formats	. 97
10.6	Examining Memory	. 99
10.7	Automatic Display	101
10.8	Print Settings	102
10.9	Value History	108
10.10	Convenience Variables	109
10.11	Registers	110
10.12	Floating Point Hardware	. 112
10.13	Vector Unit	. 112
10.14	Operating System Auxiliary Information	112
10.15	Memory Region Attributes	. 113
10	.15.1 Attributes	114
	10.15.1.1 Memory Access Mode	114
	10.15.1.2 Memory Access Size	. 114
	10.15.1.3 Data Cache	114
10	.15.2 Memory Access Checking	115
10.16	Copy Between Memory and a File	115
10.17	How to Produce a Core File from Your Program	116
10.18	Character Sets	116
10.19	Caching Data of Remote Targets	119
10.20	Search Memory	119
T		100
11 I	Debugging Optimized Code	123
11.1	Inline Functions	123
12	C Preprocessor Macros	125
13 T	racepoints	129
13.1	Commands to Set Tracepoints	129
13.1	.1.1 Create and Delete Tracepoints	
	.1.2 Enable and Disable Tracepoints	
	.1.3 Tracepoint Passcounts	
	.1.4 Tracepoint Conditions	
	.1.5 Trace State Variables	
	.1.6 Tracepoint Action Lists	
	.1.7 Listing Tracepoints	
	.1.8 Starting and Stopping Trace Experiments	
13.2		
	.2.1 tfind n	
	.2.2 tdump	
	.2.3 save-tracepoints filename	
13.3	Convenience Variables for Tracepoints	
13.4	Using Trace Files	

14 Debugg	ing Programs That Use Overlays	
		141
14.1 How Ove	erlays Work	141
14.2 Overlay	Commands	142
14.3 Automa	tic Overlay Debugging	144
14.4 Overlay	Sample Program	145
15 Using G	FDB with Different Languages	147
	g Between Source Languages	
15.1.1 List	t of Filename Extensions and Languages	147
	ting the Working Language	
	ving GDB Infer the Source Language	
	ng the Language	
	d Range Checking	
	Overview of Type Checking	
	Overview of Range Checking	
	ed Languages	
15.4.1 C a 15.4.1.1	nd C++ C and C++ Operators	
15.4.1.1	C and C++ Constants	
15.4.1.3	C++ Expressions	
15.4.1.4	C and C++ Defaults	
15.4.1.5	C and C++ Type and Range Checks	
15.4.1.6	GDB and C	
15.4.1.7	GDB Features for C++	
15.4.1.8	Decimal Floating Point format	
	jective-C	
15.4.2.1	Method Names in Commands	157
15.4.2.2	The Print Command With Objective-C	158
15.4.3 For	tran	
15.4.3.1	Fortran Operators and Expressions	158
15.4.3.2	Fortran Defaults	
15.4.3.3	Special Fortran Commands	
	cal	
	dula-2	
15.4.5.1	Operators	
15.4.5.2	Built-in Functions and Procedures	
15.4.5.3	Constants	
15.4.5.4	Modula-2 Types	
15.4.5.5 15.4.5.6	Modula-2 Defaults	
15.4.5.7	Modula-2 Type and Range Checks	
15.4.5.8	The Scope Operators :: and	
15.4.5.9	GDB and Modula-2	
	GDB and Woddia-2	
15.4.6.1	Introduction	
15.4.6.2	Omissions from Ada	
15.4.6.3	Additions to Ada	

		15.4.6.5 Extensions for Ada Tasks	169 169
		15.4.6.6 Tasking Support when Debugging Core Files	172
	15.5	15.4.6.7 Known Peculiarities of Ada Mode	172173
	10.0	Unsupported Languages	110
16	6 E	Examining the Symbol Table 1	75
17	7 A	Altering Execution 1	81
	17.1	Assignment to Variables	
	17.2	Continuing at a Different Address	
	17.3	Giving your Program a Signal	
	17.4	Returning from a Function	
	17.5 17.6	Calling Program Functions Patching Programs	
	11.0	1 atoming 1 rograms	100
18	3 G	DB Files	87
	18.1	Commands to Specify Files	187
	18.2	Debugging Information in Separate Files	
	18.3	Errors Reading Symbol Files	
	18.4	GDB Data Files	198
19	S	specifying a Debugging Target 1	99
	19.1	Active Targets	199
	19.2	Commands for Managing Targets	200
	19.3	Choosing Target Byte Order	202
~ ~			
20) L	Debugging Remote Programs 2	
	20.1	Connecting to a Remote Target	
	20.2	Sending files to a remote system	
		S S	
	20	0.3.1 Running gdbserver 20.3.1.1 Attaching to a Running Program	205
		20.3.1.1 Attaching to a Ruming Frogram	206
		<u> </u>	207
	20	0.3.2 Connecting to gdbserver	207
		0.3.3 Monitor Commands for gdbserver	207
	20.4	Remote Configuration	208
	20.5	Implementing a Remote Stub	212
		0.5.1 What the Stub Can Do for You	213
		0.5.2 What You Must Do for the Stub	214
		0.5.3 Putting it All Together	015

21	Conf	figuration-Specific Information	217
	21.1 Nati	ive	. 217
	21.1.1	HP-UX	. 217
	21.1.2	BSD libkvm Interface	. 217
	21.1.3	SVR4 Process Information	. 217
	21.1.4	Features for Debugging DJGPP Programs	. 219
	21.1.5	Features for Debugging MS Windows PE Executables .	. 221
	21.1	.5.1 Support for DLLs without Debugging Symbols	. 222
	21.1	.5.2 DLL Name Prefixes	. 223
	21.1	0	
	21.1.6	Commands Specific to GNU Hurd Systems	. 224
	21.1.7	QNX Neutrino	. 226
	21.1.8	Darwin	. 226
	21.2 Emb	pedded Operating Systems	. 227
	21.2.1	Using gdb with VxWorks	
	21.2	.1.1 Connecting to VxWorks	
	21.2	.1.2 VxWorks Download	
	21.2		
	21.3 Emb	pedded Processors	
	21.3.1	ARM	
	21.3.2	Renesas M32R/D and M32R/SDI	
	21.3.3	M68k	
	21.3.4	MicroBlaze	
	21.3.5	MIPS Embedded	
	21.3.6	OpenRISC 1000.	
	21.3.7	PowerPC Embedded	
	21.3.8	HP PA Embedded	
	21.3.9	Tsqware Sparclet	
	21.3	0	
	21.3	0 1	
	21.3	1	
	21.3	0 00 0	
	21.3.10	-5 ··· - 1	
		Zilog Z8000	
	21.3.12	Atmel AVR	
	21.3.13		
	21.3.14	1	
		hitectures	
	21.4.1	x86 Architecture-specific Issues	
	21.4.2	A29K	
	21.4.3 $21.4.4$	MIPS	
	21.4.4 $21.4.5$	HPPA	
	21.4.5 $21.4.6$	Cell Broadband Engine SPU architecture	
	21.4.0 $21.4.7$	PowerPC	0.4.4
	41.4.1	1 (7)() (1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1	. 444

22	C	Controlling GDB	245
2	2.1	Prompt	. 245
2	2.2	Command Editing	
2	2.3	Command History	
	2.4	Screen Size	
2	2.5	Numbers	
	2.6	Configuring the Current ABI	
	2.7	Optional Warnings and Messages	
	2.8	Optional Messages about Internal Happenings	
	2.9	Other Miscellaneous Settings	
23	E	vtonding CDD	255
		xtending GDB	
2	3.1	Canned Sequences of Commands	
		.1.1 User-defined Commands	
		.1.2 User-defined Command Hooks	
	_	.1.3 Command Files	
		.1.4 Commands for Controlled Output	
2	3.2		
		.2.1 Python Commands	
	23.	.2.2 Python API	
		23.2.2.1 Basic Python	
		23.2.2.2 Exception Handling	
		23.2.2.3 Auto-loading	
		23.2.2.4 Values From Inferior	
		23.2.2.5 Types In Python	
		23.2.2.6 Pretty Printing	
		23.2.2.7 Selecting Pretty-Printers	
		23.2.2.8 Commands In Python	
		23.2.2.9 Writing new convenience functions	
		23.2.2.10 Objfiles In Python	
		23.2.2.11 Acessing inferior stack frames from Python	
		23.2.2.12 Python representation of lazy strings	. 277
24	C	Command Interpreters	27 9
~ =			201
25	G	DB Text User Interface	281
2	5.1	TUI Overview	. 281
2	5.2	TUI Key Bindings	. 282
2	5.3	TUI Single Key Mode	. 283
2	5.4	TUI-specific Commands	. 283
2	5.5	TUI Configuration Variables	
26	U	Jsing GDB under GNU Emacs	287

27	7 T	The GDB/MI Interface	289
	Funct	ion and Purpose	289
	Notat	ion and Terminology	. 289
	27.3	GDB/MI General Design	. 289
	27	.3.1 Context management	. 290
	27	.3.2 Asynchronous command execution and non-stop mode	291
	27	.3.3 Thread groups	. 291
	27.4	GDB/MI Command Syntax	. 292
	27	.4.1 Gdb/mi Input Syntax	
	27	.4.2 Gdb/mi Output Syntax	. 293
	27.5	, 1	
	27.6	GDB/MI Development and Front Ends	
	27.7	/ 1	
	27	.7.1 GDB/MI Result Records	
	27	.7.2 GDB/MI Stream Records	
	27	.7.3 GDB/MI Async Records	
		.7.4 GDB/MI Frame Information	
	27	.7.5 GDB/MI Thread Information	
	27.8	Simple Examples of GDB/MI Interaction	
	27.9	GDB/MI Command Description Format	
	27.10	/ 1	
	27.11	, 8	
	27.12	GDB/MI Thread Commands	
	27.13	GDB/MI Program Execution	
	27.14	GDB/MI Stack Manipulation Commands	
	27.15	GDB/MI Variable Objects	
	27.16	GDB/MI Data Manipulation	
	27.17	GDB/MI Tracepoint Commands	
	27.18	GDB/MI Symbol Query Commands	
	27.19	GDB/MI File Commands	
	27.20	GDB/MI Target Manipulation Commands	
	27.21	GDB/MI File Transfer Commands	
	27.22	Miscellaneous gdb/mi Commands	. 347
28	3 G	DB Annotations	355
	28.1	What is an Annotation?	. 355
	28.2	The Server Prefix	
	28.3	Annotation for gdb Input	
	28.4	Errors	
	28.5	Invalidation Notices	
	28.6	Running the Program	
	28.7	Displaying Source	
29	J .J	IT Compilation Interface	359
	29.1	JIT Declarations	
	29.1	Registering Code	
		Unregistering Code	360

30	Reporting Bugs in GDB	361
	30.1 Have You Found a Bug?	. 361
	30.2 How to Report Bugs	. 361
31	Command Line Editing	365
	31.1 Introduction to Line Editing	. 365
	31.2 Readline Interaction	
	31.2.1 Readline Bare Essentials	. 365
	31.2.2 Readline Movement Commands	
	31.2.3 Readline Killing Commands	
	31.2.4 Readline Arguments	
	31.2.5 Searching for Commands in the History	
	31.3 Readline Init File	
	31.3.1 Readline Init File Syntax	
	31.3.2 Conditional Init Constructs	
	31.3.3 Sample Init File	
	31.4 Bindable Readline Commands	
	31.4.1 Commands For Moving	
	31.4.2 Commands For Manipulating The History	
	31.4.3 Commands For Changing Text	
	31.4.5 Specifying Numeric Arguments	
	31.4.6 Letting Readline Type For You	
	31.4.7 Keyboard Macros	
	31.4.8 Some Miscellaneous Commands	
	31.5 Readline vi Mode	
		. 000
32	2 Using History Interactively	385
	32.1 History Expansion	
	32.1.1 Event Designators	
	32.1.2 Word Designators	
	32.1.3 Modifiers	
A	ppendix A Formatting Documentation	389
Α	ppendix B Installing GDB	391
	B.1 Requirements for Building GDB	
	B.2 Invoking the GDB 'configure' Script	
	B.3 Compiling GDB in Another Directory	
	B.4 Specifying Names for Hosts and Targets	
	B.5 'configure' Options	
	B.6 System-wide configuration and settings	
	0	
A	ppendix C Maintenance Commands	397

Appendix D GDB Remote Serial Protocol	403
D.1 Overview	403
D.2 Standard Replies	404
D.3 Packets	
D.4 Stop Reply Packets	413
D.5 General Query Packets	
D.6 Architecture-Specific Protocol Details	
D.6.1 ARM	
D.6.1.1 Breakpoint Kinds	
D.6.2 MIPS	
D.6.2.1 Register Packet Format	
D.6.2.2 Breakpoint Kinds	
D.7 Tracepoint Packets	
D.8 Host I/O Packets	
D.9 Interrupts	
D.10 Notification Packets	
D.11 Remote Protocol Support for Non-Stop Mode	
D.12 Packet Acknowledgment	
D.13 Examples	
D.14 File-I/O Remote Protocol Extension	
D.14.1 File-I/O Overview	
D.14.2 Protocol Basics	
D.14.3 The F Request Packet	
D.14.4 The F Reply Packet	
D.14.5 The 'Ctrl-C' Message	
D.14.6 Console I/O	
D.14.7 List of Supported Calls	
close	
read	
write	
lseek	
rename	
unlink	
stat/fstat	
gettimeofday	
isatty	
system	
D.14.8 Protocol-specific Representation of Datatypes	
Integral Datatypes	
Pointer Values	
Memory Transfer	
struct stat	
struct timeval	
D.14.9 Constants	
Open Flags	
mode_t Values	
Errno Values	

Lseek Flags	
Limits	448
D.14.10 File-I/O Examples	448
D.15 Library List Format	449
D.16 Memory Map Format	450
D.17 Thread List Format	451
Appendix E The GDB Agent Expression	
Mechanism	
E.1 General Bytecode Design	
· · · · · · · · · · · · · · · · · · ·	
J · · · · · · · · · · · · · · · · · · ·	
E.4 Varying Target Capabilities E.5 Rationale	
E.5 Rationale	400
Annualis E There Elle Especial	405
Appendix F Trace File Format	405
Appendix G Target Descriptions	467
G.1 Retrieving Descriptions	467
G.2 Target Description Format	
G.2.1 Inclusion	468
G.2.2 Architecture	468
G.2.3 OS ABI	469
G.2.4 Compatible Architecture	469
G.2.5 Features	469
G.2.6 Types	469
G.2.7 Registers	
G.3 Predefined Target Types	
G.4 Standard Target Features	471
G.4.1 ARM Features	
G.4.2 i386 Features	472
G.4.3 MIPS Features	473
G.4.4 M68K Features	473
G.4.5 PowerPC Features	473
Appendix H Operating System Information	on
	475
H.1 Process list	475
11.1 1100055 1150	
Appendix I GNU GENERAL PUBLIC	
LICENSE	177
LICENSE	411
A	•
Appendix J GNU Free Documentation L	
	489
J.1 ADDENDUM: How to use this License for your docume	ents 495

xii	Debugging with GDB
GDB Index	497

Summary of GDB

Summary of GDB

The purpose of a debugger such as GDB is to allow you to see what is going on "inside" another program while it executes—or what another program was doing at the moment it crashed.

GDB can do four main kinds of things (plus other things in support of these) to help you catch bugs in the act:

- Start your program, specifying anything that might affect its behavior.
- Make your program stop on specified conditions.
- Examine what has happened, when your program has stopped.
- Change things in your program, so you can experiment with correcting the effects of one bug and go on to learn about another.

You can use GDB to debug programs written in C and C++. For more information, see Section 15.4 [Supported Languages], page 151. For more information, see Section 15.4.1 [C and C++], page 151.

Support for Modula-2 is partial. For information on Modula-2, see Section 15.4.5 [Modula-2], page 159.

Debugging Pascal programs which use sets, subranges, file variables, or nested functions does not currently work. GDB does not support entering expressions, printing values, or similar features using Pascal syntax.

GDB can be used to debug programs written in Fortran, although it may be necessary to refer to some variables with a trailing underscore.

GDB can be used to debug programs written in Objective-C, using either the Apple/NeXT or the GNU Objective-C runtime.

Free Software

GDB is free software, protected by the GNU General Public License (GPL). The GPL gives you the freedom to copy or adapt a licensed program—but every person getting a copy also gets with it the freedom to modify that copy (which means that they must get access to the source code), and the freedom to distribute further copies. Typical software companies use copyrights to limit your freedoms; the Free Software Foundation uses the GPL to preserve these freedoms.

Fundamentally, the General Public License is a license which says that you have these freedoms and that you cannot take these freedoms away from anyone else.

Free Software Needs Free Documentation

The biggest deficiency in the free software community today is not in the software—it is the lack of good free documentation that we can include with the free software. Many of our most important programs do not come with free reference manuals and free introductory texts. Documentation is an essential part of any software package; when an important free software package does not come with a free manual and a free tutorial, that is a major gap. We have many such gaps today.

Consider Perl, for instance. The tutorial manuals that people normally use are non-free. How did this come about? Because the authors of those manuals published them with restrictive terms—no copying, no modification, source files not available—which exclude them from the free software world.

That wasn't the first time this sort of thing happened, and it was far from the last. Many times we have heard a GNU user eagerly describe a manual that he is writing, his intended contribution to the community, only to learn that he had ruined everything by signing a publication contract to make it non-free.

Free documentation, like free software, is a matter of freedom, not price. The problem with the non-free manual is not that publishers charge a price for printed copies—that in itself is fine. (The Free Software Foundation sells printed copies of manuals, too.) The problem is the restrictions on the use of the manual. Free manuals are available in source code form, and give you permission to copy and modify. Non-free manuals do not allow this.

The criteria of freedom for a free manual are roughly the same as for free software. Redistribution (including the normal kinds of commercial redistribution) must be permitted, so that the manual can accompany every copy of the program, both on-line and on paper.

Permission for modification of the technical content is crucial too. When people modify the software, adding or changing features, if they are conscientious they will change the manual too—so they can provide accurate and clear documentation for the modified program. A manual that leaves you no choice but to write a new manual to document a changed version of the program is not really available to our community.

Some kinds of limits on the way modification is handled are acceptable. For example, requirements to preserve the original author's copyright notice, the distribution terms, or the list of authors, are ok. It is also no problem to require modified versions to include notice that they were modified. Even entire sections that may not be deleted or changed are acceptable, as long as they deal with nontechnical topics (like this one). These kinds of restrictions are acceptable because they don't obstruct the community's normal use of the manual.

However, it must be possible to modify all the *technical* content of the manual, and then distribute the result in all the usual media, through all the usual channels. Otherwise, the restrictions obstruct the use of the manual, it is not free, and we need another manual to replace it.

Please spread the word about this issue. Our community continues to lose manuals to proprietary publishing. If we spread the word that free software needs free reference manuals and free tutorials, perhaps the next person who wants to contribute by writing documentation will realize, before it is too late, that only free manuals contribute to the free software community.

If you are writing documentation, please insist on publishing it under the GNU Free Documentation License or another free documentation license. Remember that this decision requires your approval—you don't have to let the publisher decide. Some commercial publishers will use a free license if you insist, but they will not propose the option; it is up to you to raise the issue and say firmly that this is what you want. If the publisher you are dealing with refuses, please try other publishers. If you're not sure whether a proposed license is free, write to licensing@gnu.org.

Summary of gdb 3

You can encourage commercial publishers to sell more free, copylefted manuals and tutorials by buying them, and particularly by buying copies from the publishers that paid for their writing or for major improvements. Meanwhile, try to avoid buying non-free documentation at all. Check the distribution terms of a manual before you buy it, and insist that whoever seeks your business must respect your freedom. Check the history of the book, and try to reward the publishers that have paid or pay the authors to work on it.

The Free Software Foundation maintains a list of free documentation published by other publishers, at http://www.fsf.org/doc/other-free-books.html.

Contributors to GDB

Richard Stallman was the original author of GDB, and of many other GNU programs. Many others have contributed to its development. This section attempts to credit major contributors. One of the virtues of free software is that everyone is free to contribute to it; with regret, we cannot actually acknowledge everyone here. The file 'ChangeLog' in the GDB distribution approximates a blow-by-blow account.

Changes much prior to version 2.0 are lost in the mists of time.

Plea: Additions to this section are particularly welcome. If you or your friends (or enemies, to be evenhanded) have been unfairly omitted from this list, we would like to add your names!

So that they may not regard their many labors as thankless, we particularly thank those who shepherded gdb through major releases: Andrew Cagney (releases 6.3, 6.2, 6.1, 6.0, 5.3, 5.2, 5.1 and 5.0); Jim Blandy (release 4.18); Jason Molenda (release 4.17); Stan Shebs (release 4.14); Fred Fish (releases 4.16, 4.15, 4.13, 4.12, 4.11, 4.10, and 4.9); Stu Grossman and John Gilmore (releases 4.8, 4.7, 4.6, 4.5, and 4.4); John Gilmore (releases 4.3, 4.2, 4.1, 4.0, and 3.9); Jim Kingdon (releases 3.5, 3.4, and 3.3); and Randy Smith (releases 3.2, 3.1, and 3.0).

Richard Stallman, assisted at various times by Peter TerMaat, Chris Hanson, and Richard Mlynarik, handled releases through 2.8.

Michael Tiemann is the author of most of the GNU C++ support in GDB, with significant additional contributions from Per Bothner and Daniel Berlin. James Clark wrote the GNU C++ demangler. Early work on C++ was by Peter TerMaat (who also did much general update work leading to release 3.0).

GDB uses the BFD subroutine library to examine multiple object-file formats; BFD was a joint project of David V. Henkel-Wallace, Rich Pixley, Steve Chamberlain, and John Gilmore.

David Johnson wrote the original COFF support; Pace Willison did the original support for encapsulated COFF.

Brent Benson of Harris Computer Systems contributed DWARF 2 support.

Adam de Boor and Bradley Davis contributed the ISI Optimum V support. Per Bothner, Noboyuki Hikichi, and Alessandro Forin contributed MIPS support. Jean-Daniel Fekete contributed Sun 386i support. Chris Hanson improved the HP9000 support. Noboyuki Hikichi and Tomoyuki Hasei contributed Sony/News OS 3 support. David Johnson contributed Encore Umax support. Jyrki Kuoppala contributed Altos 3068 support. Jeff Law contributed HP PA and SOM support. Keith Packard contributed NS32K support.

Doug Rabson contributed Acorn Risc Machine support. Bob Rusk contributed Harris Nighthawk CX-UX support. Chris Smith contributed Convex support (and Fortran debugging). Jonathan Stone contributed Pyramid support. Michael Tiemann contributed SPARC support. Tim Tucker contributed support for the Gould NP1 and Gould Powernode. Pace Willison contributed Intel 386 support. Jay Vosburgh contributed Symmetry support. Marko Mlinar contributed OpenRISC 1000 support.

Andreas Schwab contributed M68K GNU/Linux support.

Rich Schaefer and Peter Schauer helped with support of SunOS shared libraries.

Jay Fenlason and Roland McGrath ensured that GDB and GAS agree about several machine instruction sets.

Patrick Duval, Ted Goldstein, Vikram Koka and Glenn Engel helped develop remote debugging. Intel Corporation, Wind River Systems, AMD, and ARM contributed remote debugging modules for the i960, VxWorks, A29K UDI, and RDI targets, respectively.

Brian Fox is the author of the readline libraries providing command-line editing and command history.

Andrew Beers of SUNY Buffalo wrote the language-switching code, the Modula-2 support, and contributed the Languages chapter of this manual.

Fred Fish wrote most of the support for Unix System Vr4. He also enhanced the command-completion support to cover C++ overloaded symbols.

Hitachi America (now Renesas America), Ltd. sponsored the support for H8/300, H8/500, and Super-H processors.

NEC sponsored the support for the v850, Vr4xxx, and Vr5xxx processors.

Mitsubishi (now Renesas) sponsored the support for D10V, D30V, and M32R/D processors.

Toshiba sponsored the support for the TX39 Mips processor.

Matsushita sponsored the support for the MN10200 and MN10300 processors.

Fujitsu sponsored the support for SPARClite and FR30 processors.

Kung Hsu, Jeff Law, and Rick Sladkey added support for hardware watchpoints.

Michael Snyder added support for tracepoints.

Stu Grossman wrote gdbserver.

Jim Kingdon, Peter Schauer, Ian Taylor, and Stu Grossman made nearly innumerable bug fixes and cleanups throughout GDB.

The following people at the Hewlett-Packard Company contributed support for the PARISC 2.0 architecture, HP-UX 10.20, 10.30, and 11.0 (narrow mode), HP's implementation of kernel threads, HP's aC++ compiler, and the Text User Interface (nee Terminal User Interface): Ben Krepp, Richard Title, John Bishop, Susan Macchia, Kathy Mann, Satish Pai, India Paul, Steve Rehrauer, and Elena Zannoni. Kim Haase provided HP-specific information in this manual.

DJ Delorie ported gdb to MS-DOS, for the DJGPP project. Robert Hoehne made significant contributions to the DJGPP port.

Cygnus Solutions has sponsored gdb maintenance and much of its development since 1991. Cygnus engineers who have worked on gdb fulltime include Mark Alexander, Jim Summary of gdb 5

Blandy, Per Bothner, Kevin Buettner, Edith Epstein, Chris Faylor, Fred Fish, Martin Hunt, Jim Ingham, John Gilmore, Stu Grossman, Kung Hsu, Jim Kingdon, John Metzler, Fernando Nasser, Geoffrey Noer, Dawn Perchik, Rich Pixley, Zdenek Radouch, Keith Seitz, Stan Shebs, David Taylor, and Elena Zannoni. In addition, Dave Brolley, Ian Carmichael, Steve Chamberlain, Nick Clifton, JT Conklin, Stan Cox, DJ Delorie, Ulrich Drepper, Frank Eigler, Doug Evans, Sean Fagan, David Henkel-Wallace, Richard Henderson, Jeff Holcomb, Jeff Law, Jim Lemke, Tom Lord, Bob Manson, Michael Meissner, Jason Merrill, Catherine Moore, Drew Moseley, Ken Raeburn, Gavin Romig-Koch, Rob Savoye, Jamie Smith, Mike Stump, Ian Taylor, Angela Thomas, Michael Tiemann, Tom Tromey, Ron Unrau, Jim Wilson, and David Zuhn have made contributions both large and small.

Andrew Cagney, Fernando Nasser, and Elena Zannoni, while working for Cygnus Solutions, implemented the original GDB/MI interface.

Jim Blandy added support for preprocessor macros, while working for Red Hat.

Andrew Cagney designed gdb's architecture vector. Many people including Andrew Cagney, Stephane Carrez, Randolph Chung, Nick Duffek, Richard Henderson, Mark Kettenis, Grace Sainsbury, Kei Sakamoto, Yoshinori Sato, Michael Snyder, Andreas Schwab, Jason Thorpe, Corinna Vinschen, Ulrich Weigand, and Elena Zannoni, helped with the migration of old architectures to this new framework.

Andrew Cagney completely re-designed and re-implemented GDB's unwinder framework, this consisting of a fresh new design featuring frame IDs, independent frame sniffers, and the sentinel frame. Mark Kettenis implemented the DWARF 2 unwinder, Jeff Johnston the libunwind unwinder, and Andrew Cagney the dummy, sentinel, tramp, and trad unwinders. The architecture-specific changes, each involving a complete rewrite of the architecture's frame code, were carried out by Jim Blandy, Joel Brobecker, Kevin Buettner, Andrew Cagney, Stephane Carrez, Randolph Chung, Orjan Friberg, Richard Henderson, Daniel Jacobowitz, Jeff Johnston, Mark Kettenis, Theodore A. Roth, Kei Sakamoto, Yoshinori Sato, Michael Snyder, Corinna Vinschen, and Ulrich Weigand.

Christian Zankel, Ross Morley, Bob Wilson, and Maxim Grigoriev from Tensilica, Inc. contributed support for Xtensa processors. Others who have worked on the Xtensa port of GDB in the past include Steve Tjiang, John Newlin, and Scott Foehner.

Michael Eager and staff of Xilinx, Inc., contributed support for the Xilinx MicroBlaze architecture.

1 A Sample gdb Session

You can use this manual at your leisure to read all about GDB. However, a handful of commands are enough to get started using the debugger. This chapter illustrates those commands.

In this sample session, we emphasize user input like this: **input**, to make it easier to pick out from the surrounding output.

One of the preliminary versions of GNU m4 (a generic macro processor) exhibits the following bug: sometimes, when we change its quote strings from the default, the commands used to capture one macro definition within another stop working. In the following short m4 session, we define a macro foo which expands to 0000; we then use the m4 built-in defn to define bar as the same thing. However, when we change the open quote string to <QUOTE> and the close quote string to <UNQUOTE>, the same procedure fails to define a new synonym baz:

```
$ cd gnu/m4
      $./m4
      define(foo,0000)
      foo
      0000
      define(bar,defn('foo'))
      bar
      0000
      changequote(<QUOTE>,<UNQUOTE>)
      define(baz,defn(<QUOTE>foo<UNQUOTE>))
      baz
      Ctrl-d
      m4: End of input: 0: fatal error: EOF in string
Let us use GDB to try to see what is going on.
      GDB is free software and you are welcome to distribute copies
       of it under certain conditions; type "show copying" to see
       the conditions.
      There is absolutely no warranty for GDB; type "show warranty"
       for details.
      GDB 7.0.50.20100218-cvs, Copyright 1999 Free Software Foundation, Inc...
      (gdb)
```

GDB reads only enough symbol data to know where to find the rest when needed; as a result, the first prompt comes up very quickly. We now tell GDB to use a narrower display width than usual, so that examples fit in this manual.

```
(gdb) set width 70
```

We need to see how the m4 built-in changequote works. Having looked at the source, we know the relevant subroutine is m4_changequote, so we set a breakpoint there with the GDB break command.

```
(gdb) break m4_changequote
Breakpoint 1 at 0x62f4: file builtin.c, line 879.
```

Using the run command, we start m4 running under GDB control; as long as control does not reach the m4_changequote subroutine, the program runs as usual:

```
(gdb) run
Starting program: /work/Editorial/gdb/gnu/m4/m4
define(foo,0000)

foo
0000
```

To trigger the breakpoint, we call changequote. GDB suspends execution of m4, displaying information about the context where it stops.

```
changequote(<QUOTE>,<UNQUOTE>)
```

```
Breakpoint 1, m4_changequote (argc=3, argv=0x33c70)
    at builtin.c:879
879     if (bad_argc(TOKEN_DATA_TEXT(argv[0]),argc,1,3))
```

Now we use the command n (next) to advance execution to the next line of the current function.

set_quotes looks like a promising subroutine. We can go into it by using the command s (step) instead of next. step goes to the next line to be executed in *any* subroutine, so it steps into set_quotes.

```
(gdb) s
set_quotes (lq=0x34c78 "<QUOTE>", rq=0x34c88 "<UNQUOTE>")
    at input.c:530
530     if (lquote != def_lquote)
```

The display that shows the subroutine where m4 is now suspended (and its arguments) is called a stack frame display. It shows a summary of the stack. We can use the backtrace command (which can also be spelled bt), to see where we are in the stack as a whole: the backtrace command displays a stack frame for each active subroutine.

```
(gdb) bt
#0 set_quotes (1q=0x34c78 "<QUOTE>", rq=0x34c88 "<UNQUOTE>")
    at input.c:530
#1 0x6344 in m4_changequote (argc=3, argv=0x33c70)
    at builtin.c:882
#2 0x8174 in expand_macro (sym=0x33320) at macro.c:242
#3 0x7a88 in expand_token (obs=0x0, t=209696, td=0xf7fffa30)
    at macro.c:71
#4 0x79dc in expand_input () at macro.c:40
#5 0x2930 in main (argc=0, argv=0xf7fffb20) at m4.c:195
```

We step through a few more lines to see what happens. The first two times, we can use 's'; the next two times we use n to avoid falling into the xstrdup subroutine.

The last line displayed looks a little odd; we can examine the variables lquote and rquote to see if they are in fact the new left and right quotes we specified. We use the command p (print) to see their values.

```
(gdb) p lquote

$1 = 0x35d40 "<QUOTE>"

(gdb) p rquote

$2 = 0x35d50 "<UNQUOTE>"
```

lquote and rquote are indeed the new left and right quotes. To look at some context, we can display ten lines of source surrounding the current line with the 1 (list) command.

```
(gdb) l
533
                xfree(rquote);
534
            lquote = (lq == nil || *lq == '\0') ? def_lquote\
535
: xstrdup (lq);
            rquote = (rq == nil || *rq == '\0') ? def_rquote\
536
: xstrdup (rq);
537
538
            len_lquote = strlen(rquote);
539
            len_rquote = strlen(lquote);
540
        }
541
542
```

Let us step past the two lines that set len_lquote and len_rquote, and then examine the values of those variables.

```
(gdb) n
539         len_rquote = strlen(lquote);
(gdb) n
540    }
(gdb) p len_lquote
$3 = 9
(gdb) p len_rquote
$4 = 7
```

That certainly looks wrong, assuming len_lquote and len_rquote are meant to be the lengths of lquote and rquote respectively. We can set them to better values using the p command, since it can print the value of any expression—and that expression can include subroutine calls and assignments.

```
(gdb) p len_lquote=strlen(lquote)
$5 = 7
(gdb) p len_rquote=strlen(rquote)
$6 = 9
```

Is that enough to fix the problem of using the new quotes with the m4 built-in defn? We can allow m4 to continue executing with the c (continue) command, and then try the example that caused trouble initially:

```
(gdb) c
Continuing.
define(baz,defn(<QUOTE>foo<UNQUOTE>))
baz
0000
```

Debugging with GDB

Success! The new quotes now work just as well as the default ones. The problem seems to have been just the two typos defining the wrong lengths. We allow m4 exit by giving it an EOF as input:

Ctrl-d

Program exited normally.

The message 'Program exited normally.' is from GDB; it indicates m4 has finished executing. We can end our GDB session with the GDB quit command.

(gdb) quit

2 Getting In and Out of GDB

This chapter discusses how to start GDB, and how to get out of it. The essentials are:

- type 'gdb' to start GDB.
- type quit or Ctrl-d to exit.

2.1 Invoking GDB

Invoke GDB by running the program gdb. Once started, GDB reads commands from the terminal until you tell it to exit.

You can also run gdb with a variety of arguments and options, to specify more of your debugging environment at the outset.

The command-line options described here are designed to cover a variety of situations; in some environments, some of these options may effectively be unavailable.

The most usual way to start GDB is with one argument, specifying an executable program: gdb program

You can also start with both an executable program and a core file specified:

```
gdb program core
```

You can, instead, specify a process ID as a second argument, if you want to debug a running process:

```
gdb program 1234
```

would attach GDB to process 1234 (unless you also have a file named '1234'; GDB does check for a core file first).

Taking advantage of the second command-line argument requires a fairly complete operating system; when you use GDB as a remote debugger attached to a bare board, there may not be any notion of "process", and there is often no way to get a core dump. GDB will warn you if it is unable to attach or to read core dumps.

You can optionally have gdb pass any arguments after the executable file to the inferior using --args. This option stops option processing.

```
gdb --args gcc -02 -c foo.c
```

This will cause gdb to debug gcc, and to set gcc's command-line arguments (see Section 4.3 [Arguments], page 28) to '-02 -c foo.c'.

You can run gdb without printing the front material, which describes gdb's non-warranty, by specifying -silent:

```
gdb -silent
```

You can further control how GDB starts up by using command-line options. GDB itself can remind you of the options available.

Type

```
gdb -help
```

to display all available options and briefly describe their use ('gdb -h' is a shorter equivalent).

All options and command line arguments you give are processed in sequential order. The order makes a difference when the '-x' option is used.

Debugging with GDB

2.1.1 Choosing Files

When GDB starts, it reads any arguments other than options as specifying an executable file and core file (or process ID). This is the same as if the arguments were specified by the '-se' and '-c' (or '-p') options respectively. (GDB reads the first argument that does not have an associated option flag as equivalent to the '-se' option followed by that argument; and the second argument that does not have an associated option flag, if any, as equivalent to the '-c'/'-p' option followed by that argument.) If the second argument begins with a decimal digit, GDB will first attempt to attach to it as a process, and if that fails, attempt to open it as a corefile. If you have a corefile whose name begins with a digit, you can prevent GDB from treating it as a pid by prefixing it with './', e.g. './12345'.

If gdb has not been configured to included core file support, such as for most embedded targets, then it will complain about a second argument and ignore it.

Many options have both long and short forms; both are shown in the following list. GDB also recognizes the long forms if you truncate them, so long as enough of the option is present to be unambiguous. (If you prefer, you can flag option arguments with '--' rather than '-', though we illustrate the more usual convention.)

```
-symbols file
```

-s file Read symbol table from file file.

-exec file

-e file Use file as the executable file to execute when appropriate, and for examining pure data in conjunction with a core dump.

-se file Read symbol table from file file and use it as the executable file.

-core file

-c file Use file as a core dump to examine.

-pid number

-p number

Connect to process ID number, as with the attach command.

-command file

-x file Execute commands from file file. The contents of this file is evaluated exactly as the source command would. See Section 23.1.3 [Command files], page 258.

-eval-command command

-ex command

Execute a single gdb command.

This option may be used multiple times to call multiple commands. It may also be interleaved with '-command' as required.

```
gdb -ex 'target sim' -ex 'load' \
  -x setbreakpoints -ex 'run' a.out
```

-directory directory

-d directory

Add directory to the path to search for source and script files.

-r

-readnow Read each symbol file's entire symbol table immediately, rather than the default, which is to read it incrementally as it is needed. This makes startup slower, but makes future operations faster.

2.1.2 Choosing Modes

You can run gdb in various alternative modes—for example, in batch mode or quiet mode.

-nx

-n Do not execute commands found in any initialization files. Normally, GDB executes the commands in these files after all the command options and arguments have been processed. See Section 23.1.3 [Command Files], page 258.

-quiet
-silent

-q "Quiet". Do not print the introductory and copyright messages. These messages are also suppressed in batch mode.

-batch Run in batch mode. Exit with status 0 after processing all the command files specified with '-x' (and all commands from initialization files, if not inhibited with '-n'). Exit with nonzero status if an error occurs in executing the GDB commands in the command files.

Batch mode may be useful for running GDB as a filter, for example to download and run a program on another computer; in order to make this more useful, the message

Program exited normally.

(which is ordinarily issued whenever a program running under GDB control terminates) is not issued when running in batch mode.

-batch-silent

Run in batch mode exactly like '-batch', but totally silently. All GDB output to stdout is prevented (stderr is unaffected). This is much quieter than '-silent' and would be useless for an interactive session.

This is particularly useful when using targets that give 'Loading section' messages, for example.

Note that targets that give their output via GDB, as opposed to writing directly to stdout, will also be made silent.

-return-child-result

The return code from GDB will be the return code from the child process (the process being debugged), with the following exceptions:

- GDB exits abnormally. E.g., due to an incorrect argument or an internal error. In this case the exit code is the same as it would have been without '-return-child-result'.
- The user quits with an explicit value. E.g., 'quit 1'.
- The child process never runs, or is not allowed to terminate, in which case the exit code will be -1.

This option is useful in conjunction with '-batch' or '-batch-silent', when GDB is being used as a remote program loader or simulator interface.

-nowindows

-nw

"No windows". If GDB comes with a graphical user interface (GUI) built in, then this option tells GDB to only use the command-line interface. If no GUI is available, this option has no effect.

-windows

−w

If gdb includes a GUI, then this option requires it to be used if possible.

-cd directory

Run gdb using directory as its working directory, instead of the current directory.

-fullname

-f

GNU Emacs sets this option when it runs GDB as a subprocess. It tells GDB to output the full file name and line number in a standard, recognizable fashion each time a stack frame is displayed (which includes each time your program stops). This recognizable format looks like two '\032' characters, followed by the file name, line number and character position separated by colons, and a newline. The Emacs-to-GDB interface program uses the two '\032' characters as a signal to display the source code for the frame.

-epoch

The Epoch Emacs-GDB interface sets this option when it runs GDB as a subprocess. It tells GDB to modify its print routines so as to allow Epoch to display values of expressions in a separate window.

-annotate level

This option sets the annotation level inside GDB. Its effect is identical to using 'set annotate level' (see Chapter 28 [Annotations], page 355). The annotation level controls how much information GDB prints together with its prompt, values of expressions, source lines, and other types of output. Level 0 is the normal, level 1 is for use when GDB is run as a subprocess of GNU Emacs, level 3 is the maximum annotation suitable for programs that control GDB, and level 2 has been deprecated.

The annotation mechanism has largely been superseded by GDB/MI (see Chapter 27 [GDB/MI], page 289).

--args

Change interpretation of command line so that arguments following the executable file are passed as command line arguments to the inferior. This option stops option processing.

-baud bps

-b bps

Set the line speed (baud rate or bits per second) of any serial interface used by GDB for remote debugging.

-1 timeout

Set the timeout (in seconds) of any communication used by GDB for remote debugging.

-tty device

-t device

Run using device for your program's standard input and output.

Activate the Text User Interface when starting. The Text User Interface manages several text windows on the terminal, showing source, assembly, registers and GDB command outputs (see Chapter 25 [GDB Text User Interface], page 281). Alternatively, the Text User Interface can be enabled by invoking the program 'gdbtui'. Do not use this option if you run GDB from Emacs (see Chapter 26 [Using GDB under GNU Emacs], page 287).

-interpreter interp

Use the interpreter *interp* for interface with the controlling program or device. This option is meant to be set by programs which communicate with GDB using it as a back end. See Chapter 24 [Command Interpreters], page 279.

'--interpreter=mi' (or '--interpreter=mi2') causes GDB to use the GDB/MI interface (see Chapter 27 [The GDB/MI Interface], page 289) included since GDB version 6.0. The previous GDB/MI interface, included in GDB version 5.3 and selected with '--interpreter=mi1', is deprecated. Earlier GDB/MI interfaces are no longer supported.

-write Open the executable and core files for both reading and writing. This is equivalent to the 'set write on' command inside GDB (see Section 17.6 [Patching], page 185).

-statistics

This option causes gdb to print statistics about time and memory usage after it completes each command and returns to the prompt.

-version This option causes GDB to print its version number and no-warranty blurb, and exit.

2.1.3 What gdb Does During Startup

Here's the description of what GDB does during session startup:

- 1. Sets up the command interpreter as specified by the command line (see Section 2.1.2 [Mode Options], page 13).
- 2. Reads the system-wide *init file* (if '--with-system-gdbinit' was used when building GDB; see Section B.6 [System-wide configuration and settings], page 395) and executes all the commands in that file.
- 3. Reads the init file (if any) in your home directory¹ and executes all the commands in that file.
- 4. Processes command line options and operands.
- 5. Reads and executes the commands from init file (if any) in the current working directory. This is only done if the current directory is different from your home directory. Thus, you can have more than one init file, one generic in your home directory, and another, specific to the program you are debugging, in the directory where you invoke GDB.

 $^{^{1}}$ On DOS/Windows systems, the home directory is the one pointed to by the HOME environment variable.

6. Reads command files specified by the '-x' option. See Section 23.1.3 [Command Files], page 258, for more details about GDB command files.

7. Reads the command history recorded in the *history file*. See Section 22.3 [Command History], page 245, for more details about the command history and the files where GDB records it.

Init files use the same syntax as command files (see Section 23.1.3 [Command Files], page 258) and are processed by GDB in the same way. The init file in your home directory can set options (such as 'set complaints') that affect subsequent processing of command line options and operands. Init files are not executed if you use the '-nx' option (see Section 2.1.2 [Choosing Modes], page 13).

To display the list of init files loaded by gdb at startup, you can use gdb --help.

The GDB init files are normally called '.gdbinit'. The DJGPP port of GDB uses the name 'gdb.ini', due to the limitations of file names imposed by DOS filesystems. The Windows ports of GDB use the standard name, but if they find a 'gdb.ini' file, they warn you about that and suggest to rename the file to the standard name.

2.2 Quitting GDB

quit [expression]

To exit gdb, use the quit command (abbreviated q), or type an end-of-file character (usually Ctrl-d). If you do not supply expression, gdb will terminate normally; otherwise it will terminate using the result of expression as the error code.

An interrupt (often Ctr1-c) does not exit from GDB, but rather terminates the action of any GDB command that is in progress and returns to GDB command level. It is safe to type the interrupt character at any time because GDB does not allow it to take effect until a time when it is safe.

If you have been using GDB to control an attached process or device, you can release it with the detach command (see Section 4.7 [Debugging an Already-running Process], page 31).

2.3 Shell Commands

If you need to execute occasional shell commands during your debugging session, there is no need to leave or suspend GDB; you can just use the shell command.

shell command string

Invoke a standard shell to execute *command string*. If it exists, the environment variable SHELL determines which shell to run. Otherwise GDB uses the default shell ('/bin/sh' on Unix systems, 'COMMAND.COM' on MS-DOS, etc.).

The utility make is often needed in development environments. You do not have to use the shell command for this purpose in GDB:

make make-args

Execute the make program with the specified arguments. This is equivalent to 'shell make make-args'.

2.4 Logging Output

You may want to save the output of GDB commands to a file. There are several commands to control GDB's logging.

set logging on

Enable logging.

set logging off

Disable logging.

set logging file file

Change the name of the current logfile. The default logfile is 'gdb.txt'.

set logging overwrite [on|off]

By default, GDB will append to the logfile. Set overwrite if you want set logging on to overwrite the logfile instead.

set logging redirect [on|off]

By default, GDB output will go to both the terminal and the logfile. Set redirect if you want output to go only to the log file.

show logging

Show the current values of the logging settings.

3 GDB Commands

You can abbreviate a GDB command to the first few letters of the command name, if that abbreviation is unambiguous; and you can repeat certain GDB commands by typing just RET. You can also use the TAB key to get GDB to fill out the rest of a word in a command (or to show you the alternatives available, if there is more than one possibility).

3.1 Command Syntax

A GDB command is a single line of input. There is no limit on how long it can be. It starts with a command name, which is followed by arguments whose meaning depends on the command name. For example, the command step accepts an argument which is the number of times to step, as in 'step 5'. You can also use the step command with no arguments. Some commands do not allow any arguments.

GDB command names may always be truncated if that abbreviation is unambiguous. Other possible command abbreviations are listed in the documentation for individual commands. In some cases, even ambiguous abbreviations are allowed; for example, **s** is specially defined as equivalent to **step** even though there are other commands whose names start with **s**. You can test abbreviations by using them as arguments to the **help** command.

A blank line as input to GDB (typing just RET) means to repeat the previous command. Certain commands (for example, run) will not repeat this way; these are commands whose unintentional repetition might cause trouble and which you are unlikely to want to repeat. User-defined commands can disable this feature; see Section 23.1.1 [Define], page 255.

The list and x commands, when you repeat them with RET, construct new arguments rather than repeating exactly as typed. This permits easy scanning of source or memory.

GDB can also use RET in another way: to partition lengthy output, in a way similar to the common utility more (see Section 22.4 [Screen Size], page 247). Since it is easy to press one RET too many in this situation, GDB disables command repetition after any command that generates this sort of display.

Any text from a # to the end of the line is a comment; it does nothing. This is useful mainly in command files (see Section 23.1.3 [Command Files], page 258).

The Ctrl-o binding is useful for repeating a complex sequence of commands. This command accepts the current line, like RET, and then fetches the next line relative to the current line from the history for editing.

3.2 Command Completion

GDB can fill in the rest of a word in a command for you, if there is only one possibility; it can also show you what the valid possibilities are for the next word in a command, at any time. This works for GDB commands, GDB subcommands, and the names of symbols in your program.

Press the TAB key whenever you want GDB to fill out the rest of a word. If there is only one possibility, GDB fills in the word, and waits for you to finish the command (or press RET to enter it). For example, if you type

(gdb) info bre TAB

GDB fills in the rest of the word 'breakpoints', since that is the only info subcommand beginning with 'bre':

```
(gdb) info breakpoints
```

You can either press RET at this point, to run the info breakpoints command, or backspace and enter something else, if 'breakpoints' does not look like the command you expected. (If you were sure you wanted info breakpoints in the first place, you might as well just type RET immediately after 'info bre', to exploit command abbreviations rather than command completion).

If there is more than one possibility for the next word when you press TAB, GDB sounds a bell. You can either supply more characters and try again, or just press TAB a second time; GDB displays all the possible completions for that word. For example, you might want to set a breakpoint on a subroutine whose name begins with 'make_', but when you type b make_TAB GDB just sounds the bell. Typing TAB again displays all the function names in your program that begin with those characters, for example:

After displaying the available possibilities, GDB copies your partial input ('b make_' in the example) so you can finish the command.

If you just want to see the list of alternatives in the first place, you can press M-? rather than pressing TAB twice. M-? means META?. You can type this either by holding down a key designated as the META shift on your keyboard (if there is one) while typing ?, or as ESC followed by ?.

Sometimes the string you need, while logically a "word", may contain parentheses or other characters that GDB normally excludes from its notion of a word. To permit word completion to work in this situation, you may enclose words in ' (single quote marks) in GDB commands.

The most likely situation where you might need this is in typing the name of a C++ function. This is because C++ allows function overloading (multiple definitions of the same function, distinguished by argument type). For example, when you want to set a breakpoint you may need to distinguish whether you mean the version of name that takes an int parameter, name(int), or the version that takes a float parameter, name(float). To use the word-completion facilities in this situation, type a single quote ' at the beginning of the function name. This alerts GDB that it may need to consider more information than usual when you press TAB or M-? to request word completion:

```
(gdb) b 'bubble( M-?
bubble(double,double) bubble(int,int)
(gdb) b 'bubble(
```

In some cases, GDB can tell that completing a name requires using quotes. When this happens, GDB inserts the quote for you (while completing as much as it can) if you do not type the quote in the first place:

```
(gdb) b bub TAB
GDB alters your input line to the following, and rings a bell:
    (gdb) b 'bubble(
```

In general, GDB can tell that a quote is needed (and inserts it) if you have not yet started typing the argument list when you ask for completion on an overloaded symbol.

For more information about overloaded functions, see Section 15.4.1.3 [C++ Expressions], page 154. You can use the command set overload-resolution off to disable overload resolution; see Section 15.4.1.7 [GDB Features for C++], page 156.

When completing in an expression which looks up a field in a structure, GDB also tries¹ to limit completions to the field names available in the type of the left-hand-side:

```
(gdb) p gdb_stdout.M-?
magic to_delete to_fputs to_put to_rewind
to_data to_flush to_isatty to_read to_write
```

This is because the gdb_stdout is a variable of the type struct ui_file that is defined in GDB sources as follows:

```
struct ui_file
{
   int *magic;
   ui_file_flush_ftype *to_flush;
   ui_file_write_ftype *to_write;
   ui_file_fputs_ftype *to_fputs;
   ui_file_read_ftype *to_read;
   ui_file_delete_ftype *to_delete;
   ui_file_isatty_ftype *to_isatty;
   ui_file_rewind_ftype *to_rewind;
   ui_file_put_ftype *to_put;
   void *to_data;
}
```

3.3 Getting Help

You can always ask GDB itself for information on its commands, using the command help.

help

h

You can use help (abbreviated h) with no arguments to display a short list of named classes of commands:

```
(gdb) help
List of classes of commands:
aliases -- Aliases of other commands
breakpoints -- Making program stop at certain points
data -- Examining data
files -- Specifying and examining files
internals -- Maintenance commands
obscure -- Obscure features
running -- Running the program
stack -- Examining the stack
status -- Status inquiries
support -- Support facilities
tracepoints -- Tracing of program execution without
              stopping the program
user-defined -- User-defined commands
Type "help" followed by a class name for a list of
```

¹ The completer can be confused by certain kinds of invalid expressions. Also, it only examines the static type of the expression, not the dynamic type.

```
commands in that class.

Type "help" followed by command name for full documentation.

Command name abbreviations are allowed if unambiguous.

(gdb)
```

help class

Using one of the general help classes as an argument, you can get a list of the individual commands in that class. For example, here is the help display for the class status:

help command

With a command name as help argument, GDB displays a short paragraph on how to use that command.

apropos args

The apropos command searches through all of the gdb commands, and their documentation, for the regular expression specified in args. It prints out all matches found. For example:

```
apropos reload
results in:
set symbol-reloading -- Set dynamic symbol table reloading
multiple times in one run
show symbol-reloading -- Show dynamic symbol table reloading
multiple times in one run
```

complete args

The complete args command lists all the possible completions for the beginning of a command. Use args to specify the beginning of the command you want completed. For example:

```
complete i
results in:
    if
    ignore
    info
    inspect
```

This is intended for use by GNU Emacs.

In addition to help, you can use the GDB commands info and show to inquire about the state of your program, or the state of GDB itself. Each command supports many topics of inquiry; this manual introduces each of them in the appropriate context. The listings under info and under show in the Index point to all the sub-commands. See [GDB Index], page 497.

info

This command (abbreviated i) is for describing the state of your program. For example, you can show the arguments passed to a function with info args, list the registers currently in use with info registers, or list the breakpoints you have set with info breakpoints. You can get a complete list of the info sub-commands with help info.

set

You can assign the result of an expression to an environment variable with set. For example, you can set the GDB prompt to a \$-sign with set prompt \$.

show

In contrast to info, show is for describing the state of GDB itself. You can change most of the things you can show, by using the related command set; for example, you can control what number system is used for displays with set radix, or simply inquire which is currently in use with show radix.

To display all the settable parameters and their current values, you can use show with no arguments; you may also use info set. Both commands produce the same display.

Here are three miscellaneous show subcommands, all of which are exceptional in lacking corresponding set commands:

show version

Show what version of GDB is running. You should include this information in GDB bug-reports. If multiple versions of GDB are in use at your site, you may need to determine which version of GDB you are running; as GDB evolves, new commands are introduced, and old ones may wither away. Also, many system vendors ship variant versions of GDB, and there are variant versions of GDB in GNU/Linux distributions as well. The version number is the same as the one announced when you start GDB.

show copying info copying

Display information about permission for copying GDB.

show warranty info warranty

Display the GNU "NO WARRANTY" statement, or a warranty, if your version of GDB comes with one.

4 Running Programs Under GDB

When you run a program under GDB, you must first generate debugging information when you compile it.

You may start GDB with its arguments, if any, in an environment of your choice. If you are doing native debugging, you may redirect your program's input and output, debug an already running process, or kill a child process.

4.1 Compiling for Debugging

In order to debug a program effectively, you need to generate debugging information when you compile it. This debugging information is stored in the object file; it describes the data type of each variable or function and the correspondence between source line numbers and addresses in the executable code.

To request debugging information, specify the '-g' option when you run the compiler.

Programs that are to be shipped to your customers are compiled with optimizations, using the '-0' compiler option. However, some compilers are unable to handle the '-g' and '-0' options together. Using those compilers, you cannot generate optimized executables containing debugging information.

GCC, the GNU C/C++ compiler, supports '-g' with or without '-0', making it possible to debug optimized code. We recommend that you *always* use '-g' whenever you compile a program. You may think your program is correct, but there is no sense in pushing your luck. For more information, see Chapter 11 [Optimized Code], page 123.

Older versions of the GNU C compiler permitted a variant option '-gg' for debugging information. GDB no longer supports this format; if your GNU C compiler has this option, do not use it.

[Macros], page 125). Most compilers do not include information about preprocessor macros in the debugging information if you specify the '-g' flag alone, because this information is rather large. Version 3.1 and later of GCC, the GNU C compiler, provides macro information if you specify the options '-gdwarf-2' and '-g3'; the former option requests debugging information in the Dwarf 2 format, and the latter requests "extra information". In the future, we hope to find more compact ways to represent macro information, so that it can be included with '-g' alone.

4.2 Starting your Program

run

Use the run command to start your program under GDB. You must first specify the program name (except on VxWorks) with an argument to GDB (see Chapter 2 [Getting In and Out of GDB], page 11), or by using the file or exec-file command (see Section 18.1 [Commands to Specify Files], page 187).

If you are running your program in an execution environment that supports processes, run creates an inferior process and makes that process run your program. In some environments without processes, run jumps to the start of your program. Other targets, like 'remote', are always running. If you get an error message like this one:

```
The "remote" target does not support "run". Try "help target" or "continue".
```

then use continue to run your program. You may need load first (see [load], page 201).

The execution of a program is affected by certain information it receives from its superior. GDB provides ways to specify this information, which you must do *before* starting your program. (You can change it after starting your program, but such changes only affect your program the next time you start it.) This information may be divided into four categories:

The arguments.

Specify the arguments to give your program as the arguments of the run command. If a shell is available on your target, the shell is used to pass the arguments, so that you may use normal conventions (such as wildcard expansion or variable substitution) in describing the arguments. In Unix systems, you can control which shell is used with the SHELL environment variable. See Section 4.3 [Your Program's Arguments], page 28.

The environment.

Your program normally inherits its environment from GDB, but you can use the GDB commands set environment and unset environment to change parts of the environment that affect your program. See Section 4.4 [Your Program's Environment], page 29.

The working directory.

Your program inherits its working directory from GDB. You can set the GDB working directory with the cd command in GDB. See Section 4.5 [Your Program's Working Directory], page 30.

The standard input and output.

Your program normally uses the same device for standard input and standard output as GDB is using. You can redirect input and output in the run command line, or you can use the tty command to set a different device for your program. See Section 4.6 [Your Program's Input and Output], page 30.

Warning: While input and output redirection work, you cannot use pipes to pass the output of the program you are debugging to another program; if you attempt this, GDB is likely to wind up debugging the wrong program.

When you issue the run command, your program begins to execute immediately. See Chapter 5 [Stopping and Continuing], page 43, for discussion of how to arrange for your program to stop. Once your program has stopped, you may call functions in your program, using the print or call commands. See Chapter 10 [Examining Data], page 93.

If the modification time of your symbol file has changed since the last time GDB read its symbols, GDB discards its symbol table, and reads it again. When it does this, GDB tries to retain your current breakpoints.

The name of the main procedure can vary from language to language. With C or C++, the main procedure name is always main, but other languages such as Ada do not require a specific name for their main procedure. The debugger provides a convenient way to start the execution of the program and to stop at the beginning of the main procedure, depending on the language used.

The 'start' command does the equivalent of setting a temporary breakpoint at the beginning of the main procedure and then invoking the 'run' command.

Some programs contain an *elaboration* phase where some startup code is executed before the main procedure is called. This depends on the languages used to write your program. In C++, for instance, constructors for static and global objects are executed before main is called. It is therefore possible that the debugger stops before reaching the main procedure. However, the temporary breakpoint will remain to halt execution.

Specify the arguments to give to your program as arguments to the 'start' command. These arguments will be given verbatim to the underlying 'run' command. Note that the same arguments will be reused if no argument is provided during subsequent calls to 'start' or 'run'.

It is sometimes necessary to debug the program during elaboration. In these cases, using the start command would stop the execution of your program too late, as the program would have already completed the elaboration phase. Under these circumstances, insert breakpoints in your elaboration code before running your program.

```
set exec-wrapper wrapper
show exec-wrapper
unset exec-wrapper
```

When 'exec-wrapper' is set, the specified wrapper is used to launch programs for debugging. GDB starts your program with a shell command of the form exec wrapper program. Quoting is added to program and its arguments, but not to wrapper, so you should add quotes if appropriate for your shell. The wrapper runs until it executes your program, and then GDB takes control.

You can use any program that eventually calls execve with its arguments as a wrapper. Several standard Unix utilities do this, e.g. env and nohup. Any Unix shell script ending with exec "\$0" will also work.

For example, you can use **env** to pass an environment variable to the debugged program, without setting the variable in your shell's environment:

```
(gdb) set exec-wrapper env 'LD_PRELOAD=libtest.so' (gdb) run
```

This command is available when debugging locally on most targets, excluding DJGPP, Cygwin, MS Windows, and QNX Neutrino.

```
set disable-randomization set disable-randomization on
```

This option (enabled by default in GDB) will turn off the native randomization of the virtual address space of the started program. This option is useful for multiple debugging sessions to make the execution better reproducible and memory addresses reusable across debugging sessions.

This feature is implemented only on GNU/Linux. You can get the same behavior using

```
(gdb) set exec-wrapper setarch 'uname -m' -R
```

set disable-randomization off

Leave the behavior of the started executable unchanged. Some bugs rear their ugly heads only when the program is loaded at certain addresses. If your bug disappears when you run the program under GDB, that might be because GDB by default disables the address randomization on platforms, such as GNU/Linux, which do that for stand-alone programs. Use set disable-randomization off to try to reproduce such elusive bugs.

The virtual address space randomization is implemented only on GNU/Linux. It protects the programs against some kinds of security attacks. In these cases the attacker needs to know the exact location of a concrete executable code. Randomizing its location makes it impossible to inject jumps misusing a code at its expected addresses.

Prelinking shared libraries provides a startup performance advantage but it makes addresses in these libraries predictable for privileged processes by having just unprivileged access at the target system. Reading the shared library binary gives enough information for assembling the malicious code misusing it. Still even a prelinked shared library can get loaded at a new random address just requiring the regular relocation process during the startup. Shared libraries not already prelinked are always loaded at a randomly chosen address.

Position independent executables (PIE) contain position independent code similar to the shared libraries and therefore such executables get loaded at a randomly chosen address upon startup. PIE executables always load even already prelinked shared libraries at a random address. You can build such executable using gcc -fPIE -pie.

Heap (malloc storage), stack and custom mmap areas are always placed randomly (as long as the randomization is enabled).

show disable-randomization

Show the current setting of the explicit disable of the native randomization of the virtual address space of the started program.

4.3 Your Program's Arguments

The arguments to your program can be specified by the arguments of the run command. They are passed to a shell, which expands wildcard characters and performs redirection of I/O, and thence to your program. Your SHELL environment variable (if it exists) specifies what shell GDB uses. If you do not define SHELL, GDB uses the default shell ('/bin/sh' on Unix).

On non-Unix systems, the program is usually invoked directly by GDB, which emulates I/O redirection via the appropriate system calls, and the wildcard characters are expanded by the startup code of the program, not by the shell.

run with no arguments uses the same arguments used by the previous run, or those set by the set args command.

set args Specify the arguments to be used the next time your program is run. If set args has no arguments, run executes your program with no arguments. Once you have run your program with arguments, using set args before the next run is the only way to run it again without arguments.

show args Show the arguments to give your program when it is started.

4.4 Your Program's Environment

The environment consists of a set of environment variables and their values. Environment variables conventionally record such things as your user name, your home directory, your terminal type, and your search path for programs to run. Usually you set up environment variables with the shell and they are inherited by all the other programs you run. When debugging, it can be useful to try running your program with a modified environment without having to start GDB over again.

path directory

Add directory to the front of the PATH environment variable (the search path for executables) that will be passed to your program. The value of PATH used by GDB does not change. You may specify several directory names, separated by whitespace or by a system-dependent separator character (':' on Unix, ';' on MS-DOS and MS-Windows). If directory is already in the path, it is moved to the front, so it is searched sooner.

You can use the string '\$cwd' to refer to whatever is the current working directory at the time GDB searches the path. If you use '.' instead, it refers to the directory where you executed the path command. GDB replaces '.' in the directory argument (with the current path) before adding directory to the search path.

show paths

Display the list of search paths for executables (the PATH environment variable).

show environment [varname]

Print the value of environment variable *varname* to be given to your program when it starts. If you do not supply *varname*, print the names and values of all environment variables to be given to your program. You can abbreviate environment as env.

set environment varname [=value]

Set environment variable varname to value. The value changes for your program only, not for GDB itself. value may be any string; the values of environment variables are just strings, and any interpretation is supplied by your program itself. The value parameter is optional; if it is eliminated, the variable is set to a null value.

For example, this command:

```
set env USER = foo
```

tells the debugged program, when subsequently run, that its user is named 'foo'. (The spaces around '=' are used for clarity here; they are not actually required.)

unset environment varname

Remove variable varname from the environment to be passed to your program. This is different from 'set env varname ='; unset environment removes the variable from the environment, rather than assigning it an empty value.

Warning: On Unix systems, GDB runs your program using the shell indicated by your SHELL environment variable if it exists (or /bin/sh if not). If your SHELL variable names a shell that runs an initialization file—such as '.cshrc' for C-shell, or '.bashrc' for BASH—any variables you set in that file affect your program. You may wish to move setting of environment variables to files that are only run when you sign on, such as '.login' or '.profile'.

4.5 Your Program's Working Directory

Each time you start your program with run, it inherits its working directory from the current working directory of GDB. The GDB working directory is initially whatever it inherited from its parent process (typically the shell), but you can specify a new working directory in GDB with the cd command.

The GDB working directory also serves as a default for the commands that specify files for GDB to operate on. See Section 18.1 [Commands to Specify Files], page 187.

cd directory

Set the gdb working directory to directory.

pwd Print the GDB working directory.

It is generally impossible to find the current working directory of the process being debugged (since a program can change its directory during its run). If you work on a system where GDB is configured with the '/proc' support, you can use the info proc command (see Section 21.1.3 [SVR4 Process Information], page 217) to find out the current working directory of the debuggee.

4.6 Your Program's Input and Output

By default, the program you run under GDB does input and output to the same terminal that GDB uses. GDB switches the terminal to its own terminal modes to interact with you, but it records the terminal modes your program was using and switches back to them when you continue running your program.

info terminal

Displays information recorded by GDB about the terminal modes your program is using.

You can redirect your program's input and/or output using shell redirection with the run command. For example,

```
run > outfile
```

starts your program, diverting its output to the file 'outfile'.

Another way to specify where your program should do input and output is with the tty command. This command accepts a file name as argument, and causes this file to be the default for future run commands. It also resets the controlling terminal for the child process, for future run commands. For example,

```
tty /dev/ttyb
```

directs that processes started with subsequent run commands default to do input and output on the terminal '/dev/ttyb' and have that as their controlling terminal.

An explicit redirection in run overrides the tty command's effect on the input/output device, but not its effect on the controlling terminal.

When you use the tty command or redirect input in the run command, only the input for your program is affected. The input for GDB still comes from your terminal. tty is an alias for set inferior-tty.

You can use the **show inferior-tty** command to tell GDB to display the name of the terminal that will be used for future runs of your program.

set inferior-tty /dev/ttyb

Set the tty for the program being debugged to /dev/ttyb.

show inferior-tty

Show the current tty for the program being debugged.

4.7 Debugging an Already-running Process

attach process-id

This command attaches to a running process—one that was started outside GDB. (info files shows your active targets.) The command takes as argument a process ID. The usual way to find out the *process-id* of a Unix process is with the ps utility, or with the 'jobs -1' shell command.

attach does not repeat if you press RET a second time after executing the command.

To use attach, your program must be running in an environment which supports processes; for example, attach does not work for programs on bare-board targets that lack an operating system. You must also have permission to send the process a signal.

When you use attach, the debugger finds the program running in the process first by looking in the current working directory, then (if the program is not found) by using the source file search path (see Section 9.5 [Specifying Source Directories], page 86). You can also use the file command to load the program. See Section 18.1 [Commands to Specify Files], page 187.

The first thing GDB does after arranging to debug the specified process is to stop it. You can examine and modify an attached process with all the GDB commands that are ordinarily available when you start processes with run. You can insert breakpoints; you can step and continue; you can modify storage. If you would rather the process continue running, you may use the continue command after attaching GDB to the process.

When you have finished debugging the attached process, you can use the detach command to release it from GDB control. Detaching the process continues its execution. After the detach command, that process and GDB become completely independent once more, and you are ready to attach another process or start one with run. detach does not repeat if you press RET again after executing the command.

If you exit GDB while you have an attached process, you detach that process. If you use the run command, you kill that process. By default, GDB asks for confirmation if you try to do either of these things; you can control whether or not you need to confirm by using the set confirm command (see Section 22.7 [Optional Warnings and Messages], page 249).

4.8 Killing the Child Process

kill Kill the child process in which your program is running under GDB.

This command is useful if you wish to debug a core dump instead of a running process. GDB ignores any core dump file while your program is running.

On some operating systems, a program cannot be executed outside GDB while you have breakpoints set on it inside GDB. You can use the kill command in this situation to permit running your program outside the debugger.

The kill command is also useful if you wish to recompile and relink your program, since on many systems it is impossible to modify an executable file while it is running in a process. In this case, when you next type run, GDB notices that the file has changed, and reads the symbol table again (while trying to preserve your current breakpoint settings).

4.9 Debugging Multiple Inferiors and Programs

GDB lets you run and debug multiple programs in a single session. In addition, GDB on some systems may let you run several programs simultaneously (otherwise you have to exit from one before starting another). In the most general case, you can have multiple threads of execution in each of multiple processes, launched from multiple executables.

GDB represents the state of each program execution with an object called an *inferior*. An inferior typically corresponds to a process, but is more general and applies also to targets that do not have processes. Inferiors may be created before a process runs, and may be retained after a process exits. Inferiors have unique identifiers that are different from process ids. Usually each inferior will also have its own distinct address space, although some embedded targets may have several inferiors running in different parts of a single address space. Each inferior may in turn have multiple threads running in it.

To find out what inferiors exist at any moment, use info inferiors:

info inferiors

Print a list of all inferiors currently being managed by GDB.

GDB displays for each inferior (in this order):

- 1. the inferior number assigned by GDB
- 2. the target system's inferior identifier
- 3. the name of the executable the inferior is running.

An asterisk '*' preceding the gdb inferior number indicates the current inferior.

For example,

(gdb) info inferiors
Num Description Executable
2 process 2307 hello
* 1 process 3401 goodbye

To switch focus between inferiors, use the inferior command:

inferior infno

Make inferior number *infno* the current inferior. The argument *infno* is the inferior number assigned by GDB, as shown in the first field of the 'info inferiors' display.

You can get multiple executables into a debugging session via the add-inferior and clone-inferior commands. On some systems GDB can add inferiors to the debug session automatically by following calls to fork and exec. To remove inferiors from the debugging session use the remove-inferior command.

add-inferior [-copies n] [-exec executable]

Adds n inferiors to be run using executable as the executable. n defaults to 1. If no executable is specified, the inferiors begins empty, with no program. You can still assign or change the program assigned to the inferior at any time by using the file command with the executable name as its argument.

clone-inferior [-copies n] [infno]

Adds n inferiors ready to execute the same program as inferior infno. n defaults to 1. infno defaults to the number of the current inferior. This is a convenient command when you want to run another instance of the inferior you are debugging.

```
(gdb) info inferiors
 Num Description
                        Executable
     process 29964
                        helloworld
* 1
(gdb) clone-inferior
Added inferior 2.
1 inferiors added.
(gdb) info inferiors
  Num Description
                        Executable
  2
       <null>
                        helloworld
      process 29964
* 1
                        helloworld
```

You can now simply switch focus to inferior 2 and run it.

remove-inferior infno

Removes the inferior *infno*. It is not possible to remove an inferior that is running with this command. For those, use the kill or detach command first.

To quit debugging one of the running inferiors that is not the current inferior, you can either detach from it by using the detach inferior command (allowing it to run independently), or kill it using the kill inferior command:

detach inferior infno

Detach from the inferior identified by GDB inferior number *infno*, and remove it from the inferior list.

kill inferior infno

Kill the inferior identified by GDB inferior number infno, and remove it from the inferior list.

After the successful completion of a command such as detach, detach inferior, kill or kill inferior, or after a normal process exit, the inferior is still valid and listed with info inferiors, ready to be restarted.

To be notified when inferiors are started or exit under GDB's control use set print inferior-events:

```
set print inferior-events
set print inferior-events on
set print inferior-events off
```

The set print inferior-events command allows you to enable or disable printing of messages when GDB notices that new inferiors have started or that inferiors have exited or have been detached. By default, these messages will not be printed.

show print inferior-events

Show whether messages will be printed when GDB detects that inferiors have started, exited or have been detached.

Many commands will work the same with multiple programs as with a single program: e.g., print myglobal will simply display the value of myglobal in the current inferior.

Occasionaly, when debugging GDB itself, it may be useful to get more info about the relationship of inferiors, programs, address spaces in a debug session. You can do that with the maint info program-spaces command.

maint info program-spaces

Print a list of all program spaces currently being managed by GDB.

GDB displays for each program space (in this order):

- 1. the program space number assigned by GDB
- 2. the name of the executable loaded into the program space, with e.g., the file command.

An asterisk '*' preceding the GDB program space number indicates the current program space.

In addition, below each program space line, GDB prints extra information that isn't suitable to display in tabular form. For example, the list of inferiors bound to the program space.

```
(gdb) maint info program-spaces
  Id Executable
  2   goodbye
        Bound inferiors: ID 1 (process 21561)
* 1   hello
```

Here we can see that no inferior is running the program hello, while process 21561 is running the program goodbye. On some targets, it is possible that multiple inferiors are bound to the same program space. The most common example is that of debugging both the parent and child processes of a vfork call. For example,

```
(gdb) maint info program-spaces
  Id     Executable
* 1     vfork-test
          Bound inferiors: ID 2 (process 18050), ID 1 (process 18045)
```

Here, both inferior 2 and inferior 1 are running in the same program space as a result of inferior 1 having executed a vfork call.

4.10 Debugging Programs with Multiple Threads

In some operating systems, such as HP-UX and Solaris, a single program may have more than one *thread* of execution. The precise semantics of threads differ from one operating system to another, but in general the threads of a single program are akin to multiple processes—except that they share one address space (that is, they can all examine and modify the same variables). On the other hand, each thread has its own registers and execution stack, and perhaps private memory.

GDB provides these facilities for debugging multi-thread programs:

- automatic notification of new threads
- 'thread threadno', a command to switch among threads
- 'info threads', a command to inquire about existing threads
- 'thread apply [threadno] [all] args', a command to apply a command to a list of threads
- thread-specific breakpoints
- 'set print thread-events', which controls printing of messages on thread start and exit.
- 'set libthread-db-search-path path', which lets the user specify which libthread_db to use if the default choice isn't compatible with the program.

Warning: These facilities are not yet available on every GDB configuration where the operating system supports threads. If your GDB does not support threads, these commands have no effect. For example, a system without thread support shows no output from 'info threads', and always rejects the thread command, like this:

```
(gdb) info threads
(gdb) thread 1
Thread ID 1 not known. Use the "info threads" command to
see the IDs of currently known threads.
```

The GDB thread debugging facility allows you to observe all threads while your program runs—but whenever GDB takes control, one thread in particular is always the focus of debugging. This thread is called the *current thread*. Debugging commands show program information from the perspective of the current thread.

Whenever GDB detects a new thread in your program, it displays the target system's identification for the thread with a message in the form '[New systag]'. systag is a thread identifier whose form varies depending on the particular system. For example, on GNU/Linux, you might see

```
[New Thread 46912507313328 (LWP 25582)]
```

when GDB notices a new thread. In contrast, on an SGI system, the systag is simply something like 'process 368', with no further qualifier.

For debugging purposes, GDB associates its own thread number—always a single integer—with each thread in your program.

info threads

Display a summary of all threads currently in your program. GDB displays for each thread (in this order):

- 1. the thread number assigned by GDB
- 2. the target system's thread identifier (systag)
- 3. the current stack frame summary for that thread

An asterisk '*' to the left of the gdb thread number indicates the current thread.

For example,

On HP-UX systems:

For debugging purposes, GDB associates its own thread number—a small integer assigned in thread-creation order—with each thread in your program.

Whenever GDB detects a new thread in your program, it displays both GDB's thread number and the target system's identification for the thread with a message in the form '[New systag]'. systag is a thread identifier whose form varies depending on the particular system. For example, on HP-UX, you see

[New thread 2 (system thread 26594)] when GDB notices a new thread.

info threads

Display a summary of all threads currently in your program. GDB displays for each thread (in this order):

- 1. the thread number assigned by GDB
- 2. the target system's thread identifier (systag)
- 3. the current stack frame summary for that thread

An asterisk '*' to the left of the GDB thread number indicates the current thread.

For example,

On Solaris, you can display more information about user threads with a Solaris-specific command:

maint info sol-threads

Display info on Solaris user threads.

thread threadno

Make thread number threadno the current thread. The command argument threadno is the internal GDB thread number, as shown in the first field of the 'info threads' display. GDB responds by displaying the system identifier of the thread you selected, and its current stack frame summary:

```
(gdb) thread 2
[Switching to process 35 thread 23]
0x34e5 in sigpause ()
```

As with the '[New ...]' message, the form of the text after 'Switching to' depends on your system's conventions for identifying threads.

thread apply [threadno] [all] command

The thread apply command allows you to apply the named *command* to one or more threads. Specify the numbers of the threads that you want affected with the command argument *threadno*. It can be a single thread number, one of the numbers shown in the first field of the 'info threads' display; or it could be a range of thread numbers, as in 2-4. To apply a command to all threads, type *thread apply all command*.

```
set print thread-events
set print thread-events on
set print thread-events off
```

The set print thread-events command allows you to enable or disable printing of messages when GDB notices that new threads have started or that threads have exited. By default, these messages will be printed if detection of these events is supported by the target. Note that these messages cannot be disabled on all targets.

show print thread-events

Show whether messages will be printed when GDB detects that threads have started and exited.

See Section 5.4 [Stopping and Starting Multi-thread Programs], page 65, for more information about how GDB behaves when you stop and start programs with multiple threads.

See Section 5.1.2 [Setting Watchpoints], page 49, for information about watchpoints in programs with multiple threads.

set libthread-db-search-path [path]

If this variable is set, path is a colon-separated list of directories GDB will use to search for libthread_db. If you omit path, 'libthread-db-search-path' will be reset to an empty list.

On GNU/Linux and Solaris systems, GDB uses a "helper" libthread_db library to obtain information about threads in the inferior process. GDB will use 'libthread-db-search-path' to find libthread_db. If that fails, GDB will continue with default system shared library directories, and finally the directory from which libpthread was loaded in the inferior process.

For any libthread_db library GDB finds in above directories, GDB attempts to initialize it with the current inferior process. If this initialization fails (which could happen because of a version mismatch between libthread_db and libpthread), GDB will unload libthread_db, and continue with the next directory. If none of libthread_db libraries initialize successfully, GDB will issue a warning and thread debugging will be disabled.

Setting libthread-db-search-path is currently implemented only on some platforms.

show libthread-db-search-path

Display current libthread_db search path.

4.11 Debugging Forks

On most systems, GDB has no special support for debugging programs which create additional processes using the fork function. When a program forks, GDB will continue to debug the parent process and the child process will run unimpeded. If you have set a breakpoint in any code which the child then executes, the child will get a SIGTRAP signal which (unless it catches the signal) will cause it to terminate.

However, if you want to debug the child process there is a workaround which isn't too painful. Put a call to sleep in the code which the child process executes after the fork. It may be useful to sleep only if a certain environment variable is set, or a certain file exists, so that the delay need not occur when you don't want to run GDB on the child. While the child is sleeping, use the ps program to get its process ID. Then tell GDB (a new invocation of GDB if you are also debugging the parent process) to attach to the child process (see Section 4.7 [Attach], page 31). From that point on you can debug the child process just like any other process which you attached to.

On some systems, GDB provides support for debugging programs that create additional processes using the fork or vfork functions. Currently, the only platforms with this feature are HP-UX (11.x and later only?) and GNU/Linux (kernel version 2.5.60 and later).

By default, when a program forks, GDB will continue to debug the parent process and the child process will run unimpeded.

If you want to follow the child process instead of the parent process, use the command set follow-fork-mode.

set follow-fork-mode mode

Set the debugger response to a program call of fork or vfork. A call to fork or vfork creates a new process. The *mode* argument can be:

parent The original process is debugged after a fork. The child process runs unimpeded. This is the default.

child The new process is debugged after a fork. The parent process runs unimpeded.

show follow-fork-mode

Display the current debugger response to a fork or vfork call.

On Linux, if you want to debug both the parent and child processes, use the command set detach-on-fork.

set detach-on-fork mode

Tells gdb whether to detach one of the processes after a fork, or retain debugger control over them both.

on The child process (or parent process, depending on the value of follow-fork-mode) will be detached and allowed to run independently. This is the default.

off

Both processes will be held under the control of GDB. One process (child or parent, depending on the value of follow-fork-mode) is debugged as usual, while the other is held suspended.

show detach-on-fork

Show whether detach-on-fork mode is on/off.

If you choose to set 'detach-on-fork' mode off, then GDB will retain control of all forked processes (including nested forks). You can list the forked processes under the control of GDB by using the info inferiors command, and switch from one fork to another by using the inferior command (see Section 4.9 [Debugging Multiple Inferiors and Programs], page 32).

To quit debugging one of the forked processes, you can either detach from it by using the detach inferior command (allowing it to run independently), or kill it using the kill inferior command. See Section 4.9 [Debugging Multiple Inferiors and Programs], page 32.

If you ask to debug a child process and a vfork is followed by an exec, GDB executes the new target up to the first breakpoint in the new target. If you have a breakpoint set on main in your original program, the breakpoint will also be set on the child process's main.

On some systems, when a child process is spawned by vfork, you cannot debug the child or parent until an exec call completes.

If you issue a run command to GDB after an exec call executes, the new target restarts. To restart the parent process, use the file command with the parent executable name as its argument. By default, after an exec call executes, GDB discards the symbols of the previous executable image. You can change this behaviour with the set follow-exec-mode command.

set follow-exec-mode mode

Set debugger response to a program call of exec. An exec call replaces the program image of a process.

follow-exec-mode can be:

new

GDB creates a new inferior and rebinds the process to this new inferior. The program the process was running before the exec call can be restarted afterwards by restarting the original inferior.

For example:

```
(gdb) info inferiors
(gdb) info inferior
      Description
                     Executable
* 1
       <null>
                     prog1
(gdb) run
process 12020 is executing new program: prog2
Program exited normally.
(gdb) info inferiors
  Ιd
       Description
                     Executable
* 2
       <null>
                      prog2
  1
       <n1111>
                     prog1
```

same

GDB keeps the process bound to the same inferior. The new executable image replaces the previous executable loaded in the inferior. Restarting the inferior after the exec call, with e.g., the run

command, restarts the executable the process was running after the exec call. This is the default mode.

For example:

You can use the catch command to make GDB stop whenever a fork, vfork, or exec call is made. See Section 5.1.3 [Setting Catchpoints], page 51.

4.12 Setting a *Bookmark* to Return to Later

On certain operating systems¹, GDB is able to save a *snapshot* of a program's state, called a *checkpoint*, and come back to it later.

Returning to a checkpoint effectively undoes everything that has happened in the program since the checkpoint was saved. This includes changes in memory, registers, and even (within some limits) system state. Effectively, it is like going back in time to the moment when the checkpoint was saved.

Thus, if you're stepping thru a program and you think you're getting close to the point where things go wrong, you can save a checkpoint. Then, if you accidentally go too far and miss the critical statement, instead of having to restart your program from the beginning, you can just go back to the checkpoint and start again from there.

This can be especially useful if it takes a lot of time or steps to reach the point where you think the bug occurs.

To use the checkpoint/restart method of debugging:

checkpoint

Save a snapshot of the debugged program's current execution state. The checkpoint command takes no arguments, but each checkpoint is assigned a small integer id, similar to a breakpoint id.

info checkpoints

List the checkpoints that have been saved in the current debugging session. For each checkpoint, the following information will be listed:

Checkpoint ID
Process ID
Code Address
Source line, or label

restart checkpoint-id

Restore the program state that was saved as checkpoint number *checkpoint-id*. All program variables, registers, stack frames etc. will be returned to the values

¹ Currently, only GNU/Linux.

that they had when the checkpoint was saved. In essence, gdb will "wind back the clock" to the point in time when the checkpoint was saved.

Note that breakpoints, GDB variables, command history etc. are not affected by restoring a checkpoint. In general, a checkpoint only restores things that reside in the program being debugged, not in the debugger.

delete checkpoint checkpoint-id

Delete the previously-saved checkpoint identified by checkpoint-id.

Returning to a previously saved checkpoint will restore the user state of the program being debugged, plus a significant subset of the system (OS) state, including file pointers. It won't "un-write" data from a file, but it will rewind the file pointer to the previous location, so that the previously written data can be overwritten. For files opened in read mode, the pointer will also be restored so that the previously read data can be read again.

Of course, characters that have been sent to a printer (or other external device) cannot be "snatched back", and characters received from eg. a serial device can be removed from internal program buffers, but they cannot be "pushed back" into the serial pipeline, ready to be received again. Similarly, the actual contents of files that have been changed cannot be restored (at this time).

However, within those constraints, you actually can "rewind" your program to a previously saved point in time, and begin debugging it again — and you can change the course of events so as to debug a different execution path this time.

Finally, there is one bit of internal program state that will be different when you return to a checkpoint — the program's process id. Each checkpoint will have a unique process id (or *pid*), and each will be different from the program's original *pid*. If your program has saved a local copy of its process id, this could potentially pose a problem.

4.12.1 A Non-obvious Benefit of Using Checkpoints

On some systems such as GNU/Linux, address space randomization is performed on new processes for security reasons. This makes it difficult or impossible to set a breakpoint, or watchpoint, on an absolute address if you have to restart the program, since the absolute location of a symbol will change from one execution to the next.

A checkpoint, however, is an *identical* copy of a process. Therefore if you create a checkpoint at (eg.) the start of main, and simply return to that checkpoint instead of restarting the process, you can avoid the effects of address randomization and your symbols will all stay in the same place.

5 Stopping and Continuing

The principal purposes of using a debugger are so that you can stop your program before it terminates; or so that, if your program runs into trouble, you can investigate and find out why.

Inside GDB, your program may stop for any of several reasons, such as a signal, a breakpoint, or reaching a new line after a GDB command such as step. You may then examine and change variables, set new breakpoints or remove old ones, and then continue execution. Usually, the messages shown by GDB provide ample explanation of the status of your program—but you can also explicitly request this information at any time.

info program

Display information about the status of your program: whether it is running or not, what process it is, and why it stopped.

5.1 Breakpoints, Watchpoints, and Catchpoints

A breakpoint makes your program stop whenever a certain point in the program is reached. For each breakpoint, you can add conditions to control in finer detail whether your program stops. You can set breakpoints with the break command and its variants (see Section 5.1.1 [Setting Breakpoints], page 44), to specify the place where your program should stop by line number, function name or exact address in the program.

On some systems, you can set breakpoints in shared libraries before the executable is run. There is a minor limitation on HP-UX systems: you must wait until the executable is run in order to set breakpoints in shared library routines that are not called directly by the program (for example, routines that are arguments in a pthread_create call).

A watchpoint is a special breakpoint that stops your program when the value of an expression changes. The expression may be a value of a variable, or it could involve values of one or more variables combined by operators, such as 'a + b'. This is sometimes called data breakpoints. You must use a different command to set watchpoints (see Section 5.1.2 [Setting Watchpoints], page 49), but aside from that, you can manage a watchpoint like any other breakpoint: you enable, disable, and delete both breakpoints and watchpoints using the same commands.

You can arrange to have values from your program displayed automatically whenever GDB stops at a breakpoint. See Section 10.7 [Automatic Display], page 101.

A catchpoint is another special breakpoint that stops your program when a certain kind of event occurs, such as the throwing of a C++ exception or the loading of a library. As with watchpoints, you use a different command to set a catchpoint (see Section 5.1.3 [Setting Catchpoints], page 51), but aside from that, you can manage a catchpoint like any other breakpoint. (To stop when your program receives a signal, use the handle command; see Section 5.3 [Signals], page 63.)

GDB assigns a number to each breakpoint, watchpoint, or catchpoint when you create it; these numbers are successive integers starting with one. In many of the commands for controlling various features of breakpoints you use the breakpoint number to say which breakpoint you want to change. Each breakpoint may be *enabled* or *disabled*; if disabled, it has no effect on your program until you enable it again.

Some GDB commands accept a range of breakpoints on which to operate. A breakpoint range is either a single breakpoint number, like '5', or two such numbers, in increasing order, separated by a hyphen, like '5-7'. When a breakpoint range is given to a command, all breakpoints in that range are operated on.

5.1.1 Setting Breakpoints

Breakpoints are set with the break command (abbreviated b). The debugger convenience variable '\$bpnum' records the number of the breakpoint you've set most recently; see Section 10.10 [Convenience Variables], page 109, for a discussion of what you can do with convenience variables.

break location

Set a breakpoint at the given *location*, which can specify a function name, a line number, or an address of an instruction. (See Section 9.2 [Specify Location], page 84, for a list of all the possible ways to specify a *location*.) The breakpoint will stop your program just before it executes any of the code in the specified *location*.

When using source languages that permit overloading of symbols, such as C++, a function name may refer to more than one possible place to break. See Section 10.2 [Ambiguous Expressions], page 94, for a discussion of that situation.

It is also possible to insert a breakpoint that will stop the program only if a specific thread (see Section 5.4.4 [Thread-Specific Breakpoints], page 69) or a specific task (see Section 15.4.6.5 [Ada Tasks], page 169) hits that breakpoint.

break

When called without any arguments, break sets a breakpoint at the next instruction to be executed in the selected stack frame (see Chapter 8 [Examining the Stack], page 77). In any selected frame but the innermost, this makes your program stop as soon as control returns to that frame. This is similar to the effect of a finish command in the frame inside the selected frame—except that finish does not leave an active breakpoint. If you use break without an argument in the innermost frame, GDB stops the next time it reaches the current location; this may be useful inside loops.

GDB normally ignores breakpoints when it resumes execution, until at least one instruction has been executed. If it did not do this, you would be unable to proceed past a breakpoint without first disabling the breakpoint. This rule applies whether or not the breakpoint already existed when your program stopped.

break ... if cond

Set a breakpoint with condition *cond*; evaluate the expression *cond* each time the breakpoint is reached, and stop only if the value is nonzero—that is, if *cond* evaluates as true. '...' stands for one of the possible arguments described above (or no argument) specifying where to break. See Section 5.1.6 [Break Conditions], page 56, for more information on breakpoint conditions.

tbreak args

Set a breakpoint enabled only for one stop. args are the same as for the break command, and the breakpoint is set in the same way, but the breakpoint is auto-

matically deleted after the first time your program stops there. See Section 5.1.5 [Disabling Breakpoints], page 55.

hbreak args

Set a hardware-assisted breakpoint. args are the same as for the break command and the breakpoint is set in the same way, but the breakpoint requires hardware support and some target hardware may not have this support. The main purpose of this is EPROM/ROM code debugging, so you can set a breakpoint at an instruction without changing the instruction. This can be used with the new trap-generation provided by SPARClite DSU and most x86-based targets. These targets will generate traps when a program accesses some data or instruction address that is assigned to the debug registers. However the hardware breakpoint registers can take a limited number of breakpoints. For example, on the DSU, only two data breakpoints can be set at a time, and GDB will reject this command if more than two are used. Delete or disable unused hardware breakpoints before setting new ones (see Section 5.1.5 [Disabling Breakpoints], page 55). See Section 5.1.6 [Break Conditions], page 56. For remote targets, you can restrict the number of hardware breakpoints GDB will use, see [set remote hardware-breakpoint-limit], page 209.

thbreak args

Set a hardware-assisted breakpoint enabled only for one stop. args are the same as for the hbreak command and the breakpoint is set in the same way. However, like the tbreak command, the breakpoint is automatically deleted after the first time your program stops there. Also, like the hbreak command, the breakpoint requires hardware support and some target hardware may not have this support. See Section 5.1.5 [Disabling Breakpoints], page 55. See also Section 5.1.6 [Break Conditions], page 56.

rbreak regex

Set breakpoints on all functions matching the regular expression regex. This command sets an unconditional breakpoint on all matches, printing a list of all breakpoints it set. Once these breakpoints are set, they are treated just like the breakpoints set with the break command. You can delete them, disable them, or make them conditional the same way as any other breakpoint.

The syntax of the regular expression is the standard one used with tools like 'grep'. Note that this is different from the syntax used by shells, so for instance foo* matches all functions that include an fo followed by zero or more os. There is an implicit .* leading and trailing the regular expression you supply, so to match only functions that begin with foo, use ^foo.

When debugging C++ programs, rbreak is useful for setting breakpoints on overloaded functions that are not members of any special classes.

The rbreak command can be used to set breakpoints in all the functions in a program, like this:

(gdb) rbreak .

info breakpoints [n]
info break [n]
info watchpoints [n]

Print a table of all breakpoints, watchpoints, and catchpoints set and not deleted. Optional argument n means print information only about the specified breakpoint (or watchpoint or catchpoint). For each breakpoint, following columns are printed:

 $Breakpoint\ Numbers$

Type Breakpoint, watchpoint, or catchpoint.

Disposition

Whether the breakpoint is marked to be disabled or deleted when hit.

Enabled or Disabled

Enabled breakpoints are marked with 'y'. 'n' marks breakpoints that are not enabled.

Address Where the breakpoint is in your program, as a memory address.

For a pending breakpoint whose address is not yet known, this field will contain '<PENDING>'. Such breakpoint won't fire until a shared library that has the symbol or line referred by breakpoint is loaded. See below for details. A breakpoint with several locations

will have '<MULTIPLE>' in this field—see below for details.

What Where the breakpoint is in the source for your program, as a file and line number. For a pending breakpoint, the original string passed to the breakpoint command will be listed as it cannot be resolved until the appropriate shared library is loaded in the future.

If a breakpoint is conditional, info break shows the condition on the line following the affected breakpoint; breakpoint commands, if any, are listed after that. A pending breakpoint is allowed to have a condition specified for it. The condition is not parsed for validity until a shared library is loaded that allows the pending breakpoint to resolve to a valid location.

info break with a breakpoint number n as argument lists only that breakpoint. The convenience variable $\$ _ and the default examining-address for the x command are set to the address of the last breakpoint listed (see Section 10.6 [Examining Memory], page 99).

info break displays a count of the number of times the breakpoint has been hit. This is especially useful in conjunction with the ignore command. You can ignore a large number of breakpoint hits, look at the breakpoint info to see how many times the breakpoint was hit, and then run again, ignoring one less than that number. This will get you quickly to the last hit of that breakpoint.

GDB allows you to set any number of breakpoints at the same place in your program. There is nothing silly or meaningless about this. When the breakpoints are conditional, this is even useful (see Section 5.1.6 [Break Conditions], page 56).

It is possible that a breakpoint corresponds to several locations in your program. Examples of this situation are:

- For a C++ constructor, the GCC compiler generates several instances of the function body, used in different cases.
- For a C++ template function, a given line in the function can correspond to any number of instantiations.
- For an inlined function, a given source line can correspond to several places where that function is inlined.

In all those cases, GDB will insert a breakpoint at all the relevant locations¹.

A breakpoint with multiple locations is displayed in the breakpoint table using several rows—one header row, followed by one row for each breakpoint location. The header row has '<MULTIPLE>' in the address column. The rows for individual locations contain the actual addresses for locations, and show the functions to which those locations belong. The number column for a location is of the form breakpoint-number.location-number.

For example:

```
Num Type Disp Enb Address What
1 breakpoint keep y <MULTIPLE>
stop only if i==1
breakpoint already hit 1 time
1.1 y 0x080486a2 in void foo<int>() at t.cc:8
1.2 y 0x080486ca in void foo<double>() at t.cc:8
```

Each location can be individually enabled or disabled by passing breakpoint-number.location-number as argument to the enable and disable commands. Note that you cannot delete the individual locations from the list, you can only delete the entire list of locations that belong to their parent breakpoint (with the delete num command, where num is the number of the parent breakpoint, 1 in the above example). Disabling or enabling the parent breakpoint (see Section 5.1.5 [Disabling], page 55) affects all of the locations that belong to that breakpoint.

It's quite common to have a breakpoint inside a shared library. Shared libraries can be loaded and unloaded explicitly, and possibly repeatedly, as the program is executed. To support this use case, GDB updates breakpoint locations whenever any shared library is loaded or unloaded. Typically, you would set a breakpoint in a shared library at the beginning of your debugging session, when the library is not loaded, and when the symbols from the library are not available. When you try to set breakpoint, GDB will ask you if you want to set a so called *pending breakpoint*—breakpoint whose address is not yet resolved.

After the program is run, whenever a new shared library is loaded, GDB reevaluates all the breakpoints. When a newly loaded shared library contains the symbol or line referred to by some pending breakpoint, that breakpoint is resolved and becomes an ordinary breakpoint. When a library is unloaded, all breakpoints that refer to its symbols or source lines become pending again.

This logic works for breakpoints with multiple locations, too. For example, if you have a breakpoint in a C++ template function, and a newly loaded shared library has an instantiation of that template, a new location is added to the list of locations for the breakpoint.

¹ As of this writing, multiple-location breakpoints work only if there's line number information for all the locations. This means that they will generally not work in system libraries, unless you have debug info with line numbers for them.

Except for having unresolved address, pending breakpoints do not differ from regular breakpoints. You can set conditions or commands, enable and disable them and perform other breakpoint operations.

GDB provides some additional commands for controlling what happens when the 'break' command cannot resolve breakpoint address specification to an address:

set breakpoint pending auto

This is the default behavior. When GDB cannot find the breakpoint location, it queries you whether a pending breakpoint should be created.

set breakpoint pending on

This indicates that an unrecognized breakpoint location should automatically result in a pending breakpoint being created.

set breakpoint pending off

This indicates that pending breakpoints are not to be created. Any unrecognized breakpoint location results in an error. This setting does not affect any pending breakpoints previously created.

show breakpoint pending

Show the current behavior setting for creating pending breakpoints.

The settings above only affect the **break** command and its variants. Once breakpoint is set, it will be automatically updated as shared libraries are loaded and unloaded.

For some targets, GDB can automatically decide if hardware or software breakpoints should be used, depending on whether the breakpoint address is read-only or read-write. This applies to breakpoints set with the break command as well as to internal breakpoints set by commands like next and finish. For breakpoints set with hbreak, GDB will always use hardware breakpoints.

You can control this automatic behaviour with the following commands::

set breakpoint auto-hw on

This is the default behavior. When GDB sets a breakpoint, it will try to use the target memory map to decide if software or hardware breakpoint must be used.

set breakpoint auto-hw off

This indicates GDB should not automatically select breakpoint type. If the target provides a memory map, GDB will warn when trying to set software breakpoint at a read-only address.

GDB normally implements breakpoints by replacing the program code at the breakpoint address with a special instruction, which, when executed, given control to the debugger. By default, the program code is so modified only when the program is resumed. As soon as the program stops, GDB restores the original instructions. This behaviour guards against leaving breakpoints inserted in the target should gdb abrubptly disconnect. However, with slow remote targets, inserting and removing breakpoint can reduce the performance. This behavior can be controlled with the following commands::

set breakpoint always-inserted off

All breakpoints, including newly added by the user, are inserted in the target only when the target is resumed. All breakpoints are removed from the target when it stops.

set breakpoint always-inserted on

Causes all breakpoints to be inserted in the target at all times. If the user adds a new breakpoint, or changes an existing breakpoint, the breakpoints in the target are updated immediately. A breakpoint is removed from the target only when breakpoint itself is removed.

set breakpoint always-inserted auto

This is the default mode. If GDB is controlling the inferior in non-stop mode (see Section 5.4.2 [Non-Stop Mode], page 67), gdb behaves as if breakpoint always-inserted mode is on. If GDB is controlling the inferior in all-stop mode, GDB behaves as if breakpoint always-inserted mode is off.

GDB itself sometimes sets breakpoints in your program for special purposes, such as proper handling of longjmp (in C programs). These internal breakpoints are assigned negative numbers, starting with -1; 'info breakpoints' does not display them. You can see these breakpoints with the GDB maintenance command 'maint info breakpoints' (see [maint info breakpoints], page 397).

5.1.2 Setting Watchpoints

You can use a watchpoint to stop execution whenever the value of an expression changes, without having to predict a particular place where this may happen. (This is sometimes called a *data breakpoint*.) The expression may be as simple as the value of a single variable, or as complex as many variables combined by operators. Examples include:

- A reference to the value of a single variable.
- An address cast to an appropriate data type. For example, '*(int *)0x12345678' will watch a 4-byte region at the specified address (assuming an int occupies 4 bytes).
- An arbitrarily complex expression, such as 'a*b + c/d'. The expression can use any operators valid in the program's native language (see Chapter 15 [Languages], page 147).

You can set a watchpoint on an expression even if the expression can not be evaluated yet. For instance, you can set a watchpoint on '*global_ptr' before 'global_ptr' is initialized. GDB will stop when your program sets 'global_ptr' and the expression produces a valid value. If the expression becomes valid in some other way than changing a variable (e.g. if the memory pointed to by '*global_ptr' becomes readable as the result of a malloc call), GDB may not stop until the next time the expression changes.

Depending on your system, watchpoints may be implemented in software or hardware. GDB does software watchpointing by single-stepping your program and testing the variable's value each time, which is hundreds of times slower than normal execution. (But this may still be worth it, to catch errors where you have no clue what part of your program is the culprit.)

On some systems, such as HP-UX, PowerPC, GNU/Linux and most other x86-based targets, GDB includes support for hardware watchpoints, which do not slow down the running of your program.

watch expr [thread threadnum]

Set a watchpoint for an expression. GDB will break when the expression expr is written into by the program and its value changes. The simplest (and the most popular) use of this command is to watch the value of a single variable:

(gdb) watch foo

If the command includes a [thread threadnum] clause, GDB breaks only when the thread identified by threadnum changes the value of expr. If any other threads change the value of expr, GDB will not break. Note that watchpoints restricted to a single thread in this way only work with Hardware Watchpoints.

rwatch expr [thread threadnum]

Set a watchpoint that will break when the value of expr is read by the program.

awatch expr [thread threadnum]

Set a watchpoint that will break when *expr* is either read from or written into by the program.

info watchpoints

This command prints a list of watchpoints, breakpoints, and catchpoints; it is the same as info break (see Section 5.1.1 [Set Breaks], page 44).

GDB sets a hardware watchpoint if possible. Hardware watchpoints execute very quickly, and the debugger reports a change in value at the exact instruction where the change occurs. If GDB cannot set a hardware watchpoint, it sets a software watchpoint, which executes more slowly and reports the change in value at the next *statement*, not the instruction, after the change occurs.

You can force GDB to use only software watchpoints with the set can-use-hw-watchpoints 0 command. With this variable set to zero, GDB will never try to use hardware watchpoints, even if the underlying system supports them. (Note that hardware-assisted watchpoints that were set before setting can-use-hw-watchpoints to zero will still use the hardware mechanism of watching expression values.)

set can-use-hw-watchpoints

Set whether or not to use hardware watchpoints.

show can-use-hw-watchpoints

Show the current mode of using hardware watchpoints.

For remote targets, you can restrict the number of hardware watchpoints GDB will use, see [set remote hardware-breakpoint-limit], page 209.

When you issue the watch command, GDB reports

Hardware watchpoint num: expr

if it was able to set a hardware watchpoint.

Currently, the awatch and rwatch commands can only set hardware watchpoints, because accesses to data that don't change the value of the watched expression cannot be detected without examining every instruction as it is being executed, and GDB does not do that currently. If GDB finds that it is unable to set a hardware breakpoint with the awatch or rwatch command, it will print a message like this:

Expression cannot be implemented with read/access watchpoint.

Sometimes, GDB cannot set a hardware watchpoint because the data type of the watched expression is wider than what a hardware watchpoint on the target machine can handle. For example, some systems can only watch regions that are up to 4 bytes wide; on such systems you cannot set hardware watchpoints for an expression that yields a double-precision

floating-point number (which is typically 8 bytes wide). As a work-around, it might be possible to break the large region into a series of smaller ones and watch them with separate watchpoints.

If you set too many hardware watchpoints, GDB might be unable to insert all of them when you resume the execution of your program. Since the precise number of active watchpoints is unknown until such time as the program is about to be resumed, GDB might not be able to warn you about this when you set the watchpoints, and the warning will be printed only when the program is resumed:

Hardware watchpoint num: Could not insert watchpoint

If this happens, delete or disable some of the watchpoints.

Watching complex expressions that reference many variables can also exhaust the resources available for hardware-assisted watchpoints. That's because gdb needs to watch every variable in the expression with separately allocated resources.

If you call a function interactively using print or call, any watchpoints you have set will be inactive until GDB reaches another kind of breakpoint or the call completes.

GDB automatically deletes watchpoints that watch local (automatic) variables, or expressions that involve such variables, when they go out of scope, that is, when the execution leaves the block in which these variables were defined. In particular, when the program being debugged terminates, all local variables go out of scope, and so only watchpoints that watch global variables remain set. If you rerun the program, you will need to set all such watchpoints again. One way of doing that would be to set a code breakpoint at the entry to the main function and when it breaks, set all the watchpoints.

In multi-threaded programs, watchpoints will detect changes to the watched expression from every thread.

Warning: In multi-threaded programs, software watchpoints have only limited usefulness. If GDB creates a software watchpoint, it can only watch the value of an expression in a single thread. If you are confident that the expression can only change due to the current thread's activity (and if you are also confident that no other thread can become current), then you can use software watchpoints as usual. However, GDB may not notice when a non-current thread's activity changes the expression. (Hardware watchpoints, in contrast, watch an expression in all threads.)

See [set remote hardware-watchpoint-limit], page 209.

5.1.3 Setting Catchpoints

You can use *catchpoints* to cause the debugger to stop for certain kinds of program events, such as C++ exceptions or the loading of a shared library. Use the *catch* command to set a catchpoint.

catch event

Stop when event occurs. event can be any of the following:

throw The throwing of a C++ exception.

catch The catching of a C++ exception.

exception

An Ada exception being raised. If an exception name is specified at the end of the command (eg catch exception Program_Error), the debugger will stop only when this specific exception is raised. Otherwise, the debugger stops execution when any Ada exception is raised.

When inserting an exception catchpoint on a user-defined exception whose name is identical to one of the exceptions defined by the language, the fully qualified name must be used as the exception name. Otherwise, GDB will assume that it should stop on the pre-defined exception rather than the user-defined one. For instance, assuming an exception called Constraint_Error is defined in package Pck, then the command to use to catch such exceptions is catch exception Pck. Constraint_Error.

exception unhandled

An exception that was raised but is not handled by the program.

assert A failed Ada assertion.

exec A call to exec. This is currently only available for HP-UX and GNU/Linux.

syscall
syscall [name | number] ...

A call to or return from a system call, a.k.a. syscall. A syscall is a mechanism for application programs to request a service from the operating system (OS) or one of the OS system services. GDB can catch some or all of the syscalls issued by the debuggee, and show the related information for each syscall. If no argument is specified, calls to and returns from all system calls will be caught.

name can be any system call name that is valid for the underlying OS. Just what syscalls are valid depends on the OS. On GNU and Unix systems, you can find the full list of valid syscall names on 'usr/include/asm/unistd.h'.

Normally, GDB knows in advance which syscalls are valid for each OS, so you can use the GDB command-line completion facilities (see Section 3.2 [command completion], page 19) to list the available choices.

You may also specify the system call numerically. A syscall's number is the value passed to the OS's syscall dispatcher to identify the requested service. When you specify the syscall by its name, GDB uses its database of syscalls to convert the name into the corresponding numeric code, but using the number directly may be useful if GDB's database does not have the complete list of syscalls on your system (e.g., because GDB lags behind the OS upgrades).

The example below illustrates how this command works if you don't provide arguments to it:

```
(gdb) catch syscall
      Catchpoint 1 (syscall)
      (gdb) r
      Starting program: /tmp/catch-syscall
      Catchpoint 1 (call to syscall 'close'), \
         Oxffffe424 in __kernel_vsyscall ()
      (gdb) c
      Continuing.
      Catchpoint 1 (returned from syscall 'close'), \
      0xffffe424 in __kernel_vsyscall ()
      (gdb)
Here is an example of catching a system call by name:
      (gdb) catch syscall chroot
      Catchpoint 1 (syscall 'chroot' [61])
      (gdb) r
      Starting program: /tmp/catch-syscall
      Catchpoint 1 (call to syscall 'chroot'), \
         0xffffe424 in __kernel_vsyscall ()
      (gdb) c
      Continuing.
      Catchpoint 1 (returned from syscall 'chroot'), \
      0xffffe424 in __kernel_vsyscall ()
      (gdb)
```

An example of specifying a system call numerically. In the case below, the syscall number has a corresponding entry in the XML file, so gdb finds its name and prints it:

However, there can be situations when there is no corresponding name in XML file for that syscall number. In this case, GDB prints a warning message saying that it was not able to find the syscall name, but the catchpoint will be set anyway. See the example below:

```
(gdb) catch syscall 764 warning: The number '764' does not represent a known syscall. Catchpoint 2 (syscall 764) (gdb)
```

If you configure GDB using the '--without-expat' option, it will not be able to display syscall names. Also, if your architecture does not have an XML file describing its system calls, you will not be

able to see the syscall names. It is important to notice that these two features are used for accessing the syscall name database. In either case, you will see a warning like this:

```
(gdb) catch syscall
warning: Could not open "syscalls/i386-linux.xml"
warning: Could not load the syscall XML file 'syscalls/i386-linux.xml'.
GDB will not be able to display syscall names.
Catchpoint 1 (syscall)
(gdb)
```

Of course, the file name will change depending on your architecture and system.

Still using the example above, you can also try to catch a syscall by its number. In this case, you would see something like:

```
(gdb) catch syscall 252
Catchpoint 1 (syscall(s) 252)
```

Again, in this case GDB would not be able to display syscall's names.

fork A call to fork. This is currently only available for HP-UX and GNU/Linux.

vfork A call to vfork. This is currently only available for HP-UX and GNU/Linux.

tcatch event

Set a catchpoint that is enabled only for one stop. The catchpoint is automatically deleted after the first time the event is caught.

Use the info break command to list the current catchpoints.

There are currently some limitations to C++ exception handling (catch throw and catch catch) in GDB:

- If you call a function interactively, GDB normally returns control to you when the function has finished executing. If the call raises an exception, however, the call may bypass the mechanism that returns control to you and cause your program either to abort or to simply continue running until it hits a breakpoint, catches a signal that GDB is listening for, or exits. This is the case even if you set a catchpoint for the exception; catchpoints on exceptions are disabled within interactive calls.
- You cannot raise an exception interactively.
- You cannot install an exception handler interactively.

Sometimes catch is not the best way to debug exception handling: if you need to know exactly where an exception is raised, it is better to stop *before* the exception handler is called, since that way you can see the stack before any unwinding takes place. If you set a breakpoint in an exception handler instead, it may not be easy to find out where the exception was raised.

To stop just before an exception handler is called, you need some knowledge of the implementation. In the case of GNU C++, exceptions are raised by calling a library function named __raise_exception which has the following ANSI C interface:

```
/* addr is where the exception identifier is stored.
  id is the exception identifier. */
```

```
void __raise_exception (void **addr, void *id);
```

To make the debugger catch all exceptions before any stack unwinding takes place, set a breakpoint on __raise_exception (see Section 5.1 [Breakpoints; Watchpoints; and Exceptions], page 43).

With a conditional breakpoint (see Section 5.1.6 [Break Conditions], page 56) that depends on the value of *id*, you can stop your program when a specific exception is raised. You can use multiple conditional breakpoints to stop your program when any of a number of exceptions are raised.

5.1.4 Deleting Breakpoints

It is often necessary to eliminate a breakpoint, watchpoint, or catchpoint once it has done its job and you no longer want your program to stop there. This is called *deleting* the breakpoint. A breakpoint that has been deleted no longer exists; it is forgotten.

With the clear command you can delete breakpoints according to where they are in your program. With the delete command you can delete individual breakpoints, watchpoints, or catchpoints by specifying their breakpoint numbers.

It is not necessary to delete a breakpoint to proceed past it. GDB automatically ignores breakpoints on the first instruction to be executed when you continue execution without changing the execution address.

clear

Delete any breakpoints at the next instruction to be executed in the selected stack frame (see Section 8.3 [Selecting a Frame], page 80). When the innermost frame is selected, this is a good way to delete a breakpoint where your program just stopped.

clear location

Delete any breakpoints set at the specified *location*. See Section 9.2 [Specify Location], page 84, for the various forms of *location*; the most useful ones are listed below:

clear function

clear filename: function

Delete any breakpoints set at entry to the named function.

clear linenum

clear filename: linenum

Delete any breakpoints set at or within the code of the specified *linenum* of the specified *filename*.

delete [breakpoints] [range...]

Delete the breakpoints, watchpoints, or catchpoints of the breakpoint ranges specified as arguments. If no argument is specified, delete all breakpoints (GDB asks confirmation, unless you have set confirm off). You can abbreviate this command as d.

5.1.5 Disabling Breakpoints

Rather than deleting a breakpoint, watchpoint, or catchpoint, you might prefer to disable it. This makes the breakpoint inoperative as if it had been deleted, but remembers the information on the breakpoint so that you can enable it again later.

You disable and enable breakpoints, watchpoints, and catchpoints with the enable and disable commands, optionally specifying one or more breakpoint numbers as arguments. Use info break or info watch to print a list of breakpoints, watchpoints, and catchpoints if you do not know which numbers to use.

Disabling and enabling a breakpoint that has multiple locations affects all of its locations.

A breakpoint, watchpoint, or catchpoint can have any of four different states of enablement:

- Enabled. The breakpoint stops your program. A breakpoint set with the break command starts out in this state.
- Disabled. The breakpoint has no effect on your program.
- Enabled once. The breakpoint stops your program, but then becomes disabled.
- Enabled for deletion. The breakpoint stops your program, but immediately after it does so it is deleted permanently. A breakpoint set with the tbreak command starts out in this state.

You can use the following commands to enable or disable breakpoints, watchpoints, and catchpoints:

disable [breakpoints] [range...]

Disable the specified breakpoints—or all breakpoints, if none are listed. A disabled breakpoint has no effect but is not forgotten. All options such as ignore-counts, conditions and commands are remembered in case the breakpoint is enabled again later. You may abbreviate disable as dis.

enable [breakpoints] [range...]

Enable the specified breakpoints (or all defined breakpoints). They become effective once again in stopping your program.

enable [breakpoints] once range...

Enable the specified breakpoints temporarily. GDB disables any of these breakpoints immediately after stopping your program.

enable [breakpoints] delete range...

Enable the specified breakpoints to work once, then die. GDB deletes any of these breakpoints as soon as your program stops there. Breakpoints set by the tbreak command start out in this state.

Except for a breakpoint set with tbreak (see Section 5.1.1 [Setting Breakpoints], page 44), breakpoints that you set are initially enabled; subsequently, they become disabled or enabled only when you use one of the commands above. (The command until can set and delete a breakpoint of its own, but it does not change the state of your other breakpoints; see Section 5.2 [Continuing and Stepping], page 60.)

5.1.6 Break Conditions

The simplest sort of breakpoint breaks every time your program reaches a specified place. You can also specify a *condition* for a breakpoint. A condition is just a Boolean expression in your programming language (see Section 10.1 [Expressions], page 93). A breakpoint with a condition evaluates the expression each time your program reaches it, and your program stops only if the condition is *true*.

This is the converse of using assertions for program validation; in that situation, you want to stop when the assertion is violated—that is, when the condition is false. In C, if you want to test an assertion expressed by the condition assert, you should set the condition '! assert' on the appropriate breakpoint.

Conditions are also accepted for watchpoints; you may not need them, since a watchpoint is inspecting the value of an expression anyhow—but it might be simpler, say, to just set a watchpoint on a variable name, and specify a condition that tests whether the new value is an interesting one.

Break conditions can have side effects, and may even call functions in your program. This can be useful, for example, to activate functions that log program progress, or to use your own print functions to format special data structures. The effects are completely predictable unless there is another enabled breakpoint at the same address. (In that case, GDB might see the other breakpoint first and stop your program without checking the condition of this one.) Note that breakpoint commands are usually more convenient and flexible than break conditions for the purpose of performing side effects when a breakpoint is reached (see Section 5.1.7 [Breakpoint Command Lists], page 58).

Break conditions can be specified when a breakpoint is set, by using 'if' in the arguments to the break command. See Section 5.1.1 [Setting Breakpoints], page 44. They can also be changed at any time with the condition command.

You can also use the if keyword with the watch command. The catch command does not recognize the if keyword; condition is the only way to impose a further condition on a catchpoint.

condition bnum expression

Specify expression as the break condition for breakpoint, watchpoint, or catchpoint number bnum. After you set a condition, breakpoint bnum stops your program only if the value of expression is true (nonzero, in C). When you use condition, GDB checks expression immediately for syntactic correctness, and to determine whether symbols in it have referents in the context of your breakpoint. If expression uses symbols not referenced in the context of the breakpoint, GDB prints an error message:

No symbol "foo" in current context.

GDB does not actually evaluate expression at the time the condition command (or a command that sets a breakpoint with a condition, like break if ...) is given, however. See Section 10.1 [Expressions], page 93.

condition bnum

Remove the condition from breakpoint number *bnum*. It becomes an ordinary unconditional breakpoint.

A special case of a breakpoint condition is to stop only when the breakpoint has been reached a certain number of times. This is so useful that there is a special way to do it, using the *ignore count* of the breakpoint. Every breakpoint has an ignore count, which is an integer. Most of the time, the ignore count is zero, and therefore has no effect. But if your program reaches a breakpoint whose ignore count is positive, then instead of stopping, it just decrements the ignore count by one and continues. As a result, if the ignore count value is n, the breakpoint does not stop the next n times your program reaches it.

ignore bnum count

Set the ignore count of breakpoint number bnum to count. The next count times the breakpoint is reached, your program's execution does not stop; other than to decrement the ignore count, GDB takes no action.

To make the breakpoint stop the next time it is reached, specify a count of zero.

When you use continue to resume execution of your program from a breakpoint, you can specify an ignore count directly as an argument to continue, rather than using ignore. See Section 5.2 [Continuing and Stepping], page 60.

If a breakpoint has a positive ignore count and a condition, the condition is not checked. Once the ignore count reaches zero, GDB resumes checking the condition.

You could achieve the effect of the ignore count with a condition such as '\$foo-- <= 0' using a debugger convenience variable that is decremented each time. See Section 10.10 [Convenience Variables], page 109.

Ignore counts apply to breakpoints, watchpoints, and catchpoints.

5.1.7 Breakpoint Command Lists

You can give any breakpoint (or watchpoint or catchpoint) a series of commands to execute when your program stops due to that breakpoint. For example, you might want to print the values of certain expressions, or enable other breakpoints.

```
{\tt commands} \; [{\tt bnum}]
```

... command-list ...

end Specify a l

Specify a list of commands for breakpoint number *bnum*. The commands themselves appear on the following lines. Type a line containing just **end** to terminate the commands.

To remove all commands from a breakpoint, type commands and follow it immediately with end; that is, give no commands.

With no *bnum* argument, **commands** refers to the last breakpoint, watchpoint, or catchpoint set (not to the breakpoint most recently encountered).

Pressing RET as a means of repeating the last gdb command is disabled within a command-list.

You can use breakpoint commands to start your program up again. Simply use the continue command, or step, or any other command that resumes execution.

Any other commands in the command list, after a command that resumes execution, are ignored. This is because any time you resume execution (even with a simple next or step), you may encounter another breakpoint—which could have its own command list, leading to ambiguities about which list to execute.

If the first command you specify in a command list is silent, the usual message about stopping at a breakpoint is not printed. This may be desirable for breakpoints that are to print a specific message and then continue. If none of the remaining commands print anything, you see no sign that the breakpoint was reached. silent is meaningful only at the beginning of a breakpoint command list.

The commands echo, output, and printf allow you to print precisely controlled output, and are often useful in silent breakpoints. See Section 23.1.4 [Commands for Controlled Output], page 259.

For example, here is how you could use breakpoint commands to print the value of x at entry to foo whenever x is positive.

```
break foo if x>0
commands
silent
printf "x is %d\n",x
cont
end
```

One application for breakpoint commands is to compensate for one bug so you can test for another. Put a breakpoint just after the erroneous line of code, give it a condition to detect the case in which something erroneous has been done, and give it commands to assign correct values to any variables that need them. End with the continue command so that your program does not stop, and start with the silent command so that no output is produced. Here is an example:

```
break 403
commands
silent
set x = y + 4
cont
end
```

5.1.8 "Cannot insert breakpoints"

If you request too many active hardware-assisted breakpoints and watchpoints, you will see this error message:

```
Stopped; cannot insert breakpoints.
You may have requested too many hardware breakpoints and watchpoints.
```

This message is printed when you attempt to resume the program, since only then GDB knows exactly how many hardware breakpoints and watchpoints it needs to insert.

When this message is printed, you need to disable or remove some of the hardware-assisted breakpoints and watchpoints, and then continue.

5.1.9 "Breakpoint address adjusted..."

Some processor architectures place constraints on the addresses at which breakpoints may be placed. For architectures thus constrained, GDB will attempt to adjust the breakpoint's address to comply with the constraints dictated by the architecture.

One example of such an architecture is the Fujitsu FR-V. The FR-V is a VLIW architecture in which a number of RISC-like instructions may be bundled together for parallel execution. The FR-V architecture constrains the location of a breakpoint instruction within such a bundle to the instruction with the lowest address. GDB honors this constraint by adjusting a breakpoint's address to the first in the bundle.

It is not uncommon for optimized code to have bundles which contain instructions from different source statements, thus it may happen that a breakpoint's address will be adjusted from one source statement to another. Since this adjustment may significantly alter GDB's breakpoint related behavior from what the user expects, a warning is printed when the breakpoint is first set and also when the breakpoint is hit.

A warning like the one below is printed when setting a breakpoint that's been subject to address adjustment:

```
warning: Breakpoint address adjusted from 0x00010414 to 0x00010410.
```

Such warnings are printed both for user settable and GDB's internal breakpoints. If you see one of these warnings, you should verify that a breakpoint set at the adjusted address will have the desired affect. If not, the breakpoint in question may be removed and other breakpoints may be set which will have the desired behavior. E.g., it may be sufficient to place the breakpoint at a later instruction. A conditional breakpoint may also be useful in some cases to prevent the breakpoint from triggering too often.

GDB will also issue a warning when stopping at one of these adjusted breakpoints:

```
warning: Breakpoint 1 address previously adjusted from 0x00010414 to 0x00010410.
```

When this warning is encountered, it may be too late to take remedial action except in cases where the breakpoint is hit earlier or more frequently than expected.

5.2 Continuing and Stepping

Continuing means resuming program execution until your program completes normally. In contrast, stepping means executing just one more "step" of your program, where "step" may mean either one line of source code, or one machine instruction (depending on what particular command you use). Either when continuing or when stepping, your program may stop even sooner, due to a breakpoint or a signal. (If it stops due to a signal, you may want to use handle, or use 'signal 0' to resume execution. See Section 5.3 [Signals], page 63.)

```
continue [ignore-count]
c [ignore-count]
fg [ignore-count]
```

Resume program execution, at the address where your program last stopped; any breakpoints set at that address are bypassed. The optional argument ignore-count allows you to specify a further number of times to ignore a breakpoint at this location; its effect is like that of ignore (see Section 5.1.6 [Break Conditions], page 56).

The argument *ignore-count* is meaningful only when your program stopped due to a breakpoint. At other times, the argument to **continue** is ignored.

The synonyms c and fg (for foreground, as the debugged program is deemed to be the foreground program) are provided purely for convenience, and have exactly the same behavior as continue.

To resume execution at a different place, you can use return (see Section 17.4 [Returning from a Function], page 183) to go back to the calling function; or jump (see Section 17.2 [Continuing at a Different Address], page 182) to go to an arbitrary location in your program.

A typical technique for using stepping is to set a breakpoint (see Section 5.1 [Breakpoints; Watchpoints; and Catchpoints], page 43) at the beginning of the function or the section of your program where a problem is believed to lie, run your program until it stops at that breakpoint, and then step through the suspect area, examining the variables that are interesting, until you see the problem happen.

Step Continue running your program until control reaches a different source line, then stop it and return control to GDB. This command is abbreviated s.

Warning: If you use the step command while control is within a function that was compiled without debugging information, execution proceeds until control reaches a function that does have debugging information. Likewise, it will not step into a function which is compiled without debugging information. To step through functions without debugging information, use the stepi command, described below.

The step command only stops at the first instruction of a source line. This prevents the multiple stops that could otherwise occur in switch statements, for loops, etc. step continues to stop if a function that has debugging information is called within the line. In other words, step steps inside any functions called within the line.

Also, the step command only enters a function if there is line number information for the function. Otherwise it acts like the next command. This avoids problems when using cc -gl on MIPS machines. Previously, step entered subroutines if there was any debugging information about the routine.

step count

Continue running as in step, but do so *count* times. If a breakpoint is reached, or a signal not related to stepping occurs before *count* steps, stepping stops right away.

next [count]

Continue to the next source line in the current (innermost) stack frame. This is similar to step, but function calls that appear within the line of code are executed without stopping. Execution stops when control reaches a different line of code at the original stack level that was executing when you gave the next command. This command is abbreviated n.

An argument *count* is a repeat count, as for step.

The next command only stops at the first instruction of a source line. This prevents multiple stops that could otherwise occur in switch statements, for loops, etc.

set step-mode set step-mode on

The set step-mode on command causes the step command to stop at the first instruction of a function which contains no debug line information rather than stepping over it.

This is useful in cases where you may be interested in inspecting the machine instructions of a function which has no symbolic info and do not want GDB to automatically skip over this function.

set step-mode off

Causes the step command to step over any functions which contains no debug information. This is the default.

62 Debugging with GDB

show step-mode

Show whether GDB will stop in or step over functions without source line debug information.

finish Continue running until just after function in the selected stack frame returns. Print the returned value (if any). This command can be abbreviated as fin. Contrast this with the return command (see Section 17.4 [Returning from a Function, page 183).

until

Continue running until a source line past the current line, in the current stack frame, is reached. This command is used to avoid single stepping through a loop more than once. It is like the next command, except that when until encounters a jump, it automatically continues execution until the program counter is greater than the address of the jump.

This means that when you reach the end of a loop after single stepping though it, until makes your program continue execution until it exits the loop. In contrast, a next command at the end of a loop simply steps back to the beginning of the loop, which forces you to step through the next iteration.

until always stops your program if it attempts to exit the current stack frame. until may produce somewhat counterintuitive results if the order of machine code does not match the order of the source lines. For example, in the following excerpt from a debugging session, the f (frame) command shows that execution is stopped at line 206; yet when we use until, we get to line 195:

```
(gdb) f
#0 main (argc=4, argv=0xf7fffae8) at m4.c:206
                    expand_input();
(gdb) until
                for ( ; argc > 0; NEXTARG) {
195
```

This happened because, for execution efficiency, the compiler had generated code for the loop closure test at the end, rather than the start, of the loop even though the test in a C for-loop is written before the body of the loop. The until command appeared to step back to the beginning of the loop when it advanced to this expression; however, it has not really gone to an earlier statement—not in terms of the actual machine code.

until with no argument works by means of single instruction stepping, and hence is slower than until with an argument.

until location u location

Continue running your program until either the specified location is reached, or the current stack frame returns. location is any of the forms described in Section 9.2 [Specify Location], page 84. This form of the command uses temporary breakpoints, and hence is quicker than until without an argument. The specified location is actually reached only if it is in the current frame. This implies that until can be used to skip over recursive function invocations. For instance in the code below, if the current location is line 96, issuing until 99 will execute the program up to line 99 in the same invocation of factorial, i.e., after the inner invocations have returned.

11

```
94 int factorial (int value)
95 {
96    if (value > 1) {
97         value *= factorial (value - 1);
98    }
99    return (value);
100    }
```

advance location

Continue running the program up to the given *location*. An argument is required, which should be of one of the forms described in Section 9.2 [Specify Location], page 84. Execution will also stop upon exit from the current stack frame. This command is similar to until, but advance will not skip over recursive function calls, and the target location doesn't have to be in the same frame as the current one.

stepi stepi *arg*

si

Execute one machine instruction, then stop and return to the debugger.

It is often useful to do 'display/i \$pc' when stepping by machine instructions. This makes GDB automatically display the next instruction to be executed, each time your program stops. See Section 10.7 [Automatic Display], page 101.

An argument is a repeat count, as in step.

nexti nexti arg

ni nextrarg

Execute one machine instruction, but if it is a function call, proceed until the function returns.

An argument is a repeat count, as in next.

5.3 Signals

A signal is an asynchronous event that can happen in a program. The operating system defines the possible kinds of signals, and gives each kind a name and a number. For example, in Unix SIGINT is the signal a program gets when you type an interrupt character (often Ctr1-c); SIGSEGV is the signal a program gets from referencing a place in memory far away from all the areas in use; SIGALRM occurs when the alarm clock timer goes off (which happens only if your program has requested an alarm).

Some signals, including SIGALRM, are a normal part of the functioning of your program. Others, such as SIGSEGV, indicate errors; these signals are *fatal* (they kill your program immediately) if the program has not specified in advance some other way to handle the signal. SIGINT does not indicate an error in your program, but it is normally fatal so it can carry out the purpose of the interrupt: to kill the program.

GDB has the ability to detect any occurrence of a signal in your program. You can tell GDB in advance what to do for each kind of signal.

Normally, GDB is set up to let the non-erroneous signals like SIGALRM be silently passed to your program (so as not to interfere with their role in the program's functioning) but to stop your program immediately whenever an error signal happens. You can change these settings with the handle command.

info signals info handle

Print a table of all the kinds of signals and how GDB has been told to handle each one. You can use this to see the signal numbers of all the defined types of signals.

info signals sig

Similar, but print information only about the specified signal number.

info handle is an alias for info signals.

handle signal [keywords...]

Change the way GDB handles signal signal. signal can be the number of a signal or its name (with or without the 'SIG' at the beginning); a list of signal numbers of the form 'low-high'; or the word 'all', meaning all the known signals. Optional arguments keywords, described below, say what change to make.

The keywords allowed by the handle command can be abbreviated. Their full names are:

nostop GDB should not stop your program when this signal happens. It may still print a message telling you that the signal has come in.

stop GDB should stop your program when this signal happens. This implies the print keyword as well.

print GDB should print a message when this signal happens.

noprint GDB should not mention the occurrence of the signal at all. This implies the nostop keyword as well.

pass

noignore GDB should allow your program to see this signal; your program can handle the signal, or else it may terminate if the signal is fatal and not handled. pass and noignore are synonyms.

nopass

ignore GDB should not allow your program to see this signal. nopass and ignore are synonyms.

When a signal stops your program, the signal is not visible to the program until you continue. Your program sees the signal then, if pass is in effect for the signal in question at that time. In other words, after GDB reports a signal, you can use the handle command with pass or nopass to control whether your program sees that signal when you continue.

The default is set to nostop, noprint, pass for non-erroneous signals such as SIGALRM, SIGWINCH and SIGCHLD, and to stop, print, pass for the erroneous signals.

You can also use the **signal** command to prevent your program from seeing a signal, or cause it to see a signal it normally would not see, or to give it any signal at any time. For example, if your program stopped due to some sort of memory reference error, you might store correct values into the erroneous variables and continue, hoping to see more execution; but your program would probably terminate immediately as a result of the fatal signal once

it saw the signal. To prevent this, you can continue with 'signal 0'. See Section 17.3 [Giving your Program a Signal], page 183.

On some targets, GDB can inspect extra signal information associated with the intercepted signal, before it is actually delivered to the program being debugged. This information is exported by the convenience variable <code>\$_siginfo</code>, and consists of data that is passed by the kernel to the signal handler at the time of the receipt of a signal. The data type of the information itself is target dependent. You can see the data type using the <code>ptype \$_siginfo</code> command. On Unix systems, it typically corresponds to the standard <code>siginfo_t</code> type, as defined in the 'signal.h' system header.

Here's an example, on a GNU/Linux system, printing the stray referenced address that raised a segmentation fault.

```
(gdb) continue
Program received signal SIGSEGV, Segmentation fault.
0x0000000000400766 in main ()
          *(int *)p = 0;
(gdb) ptype $_siginfo
type = struct {
   int si_signo;
   int si_errno;
   int si_code;
   union {
        int _pad[28];
        struct {...} _kill;
        struct {...} _timer;
        struct {...} _rt;
        struct {...} _sigchld;
        struct {...} _sigfault;
        struct {...} _sigpoll;
   } _sifields;
(gdb) ptype $_siginfo._sifields._sigfault
type = struct {
   void *si_addr;
(gdb) p $_siginfo._sifields._sigfault.si_addr
$1 = (void *) 0x7ffff7ff7000
```

Depending on target support, \$_siginfo may also be writable.

5.4 Stopping and Starting Multi-thread Programs

GDB supports debugging programs with multiple threads (see Section 4.10 [Debugging Programs with Multiple Threads], page 35). There are two modes of controlling execution of your program within the debugger. In the default mode, referred to as all-stop mode, when any thread in your program stops (for example, at a breakpoint or while being stepped), all other threads in the program are also stopped by GDB. On some targets, GDB also supports non-stop mode, in which other threads can continue to run freely while you examine the stopped thread in the debugger.

5.4.1 All-Stop Mode

In all-stop mode, whenever your program stops under GDB for any reason, all threads of execution stop, not just the current thread. This allows you to examine the overall state

of the program, including switching between threads, without worrying that things may change underfoot.

Conversely, whenever you restart the program, all threads start executing. This is true even when single-stepping with commands like step or next.

In particular, GDB cannot single-step all threads in lockstep. Since thread scheduling is up to your debugging target's operating system (not controlled by GDB), other threads may execute more than one statement while the current thread completes a single step. Moreover, in general other threads stop in the middle of a statement, rather than at a clean statement boundary, when the program stops.

You might even find your program stopped in another thread after continuing or even single-stepping. This happens whenever some other thread runs into a breakpoint, a signal, or an exception before the first thread completes whatever you requested.

Whenever GDB stops your program, due to a breakpoint or a signal, it automatically selects the thread where that breakpoint or signal happened. GDB alerts you to the context switch with a message such as '[Switching to Thread n]' to identify the thread.

On some OSes, you can modify GDB's default behavior by locking the OS scheduler to allow only a single thread to run.

set scheduler-locking mode

Set the scheduler locking mode. If it is off, then there is no locking and any thread may run at any time. If on, then only the current thread may run when the inferior is resumed. The step mode optimizes for single-stepping; it prevents other threads from preempting the current thread while you are stepping, so that the focus of debugging does not change unexpectedly. Other threads only rarely (or never) get a chance to run when you step. They are more likely to run when you 'next' over a function call, and they are completely free to run when you use commands like 'continue', 'until', or 'finish'. However, unless another thread hits a breakpoint during its timeslice, GDB does not change the current thread away from the thread that you are debugging.

show scheduler-locking

Display the current scheduler locking mode.

By default, when you issue one of the execution commands such as continue, next or step, GDB allows only threads of the current inferior to run. For example, if GDB is attached to two inferiors, each with two threads, the continue command resumes only the two threads of the current inferior. This is useful, for example, when you debug a program that forks and you want to hold the parent stopped (so that, for instance, it doesn't run to exit), while you debug the child. In other situations, you may not be interested in inspecting the current state of any of the processes GDB is attached to, and you may want to resume them all until some breakpoint is hit. In the latter case, you can instruct GDB to allow all threads of all the inferiors to run with the set schedule-multiple command.

set schedule-multiple

Set the mode for allowing threads of multiple processes to be resumed when an execution command is issued. When on, all threads of all processes are allowed to run. When off, only the threads of the current process are resumed. The

default is off. The scheduler-locking mode takes precedence when set to on, or while you are stepping and set to step.

show schedule-multiple

Display the current mode for resuming the execution of threads of multiple processes.

5.4.2 Non-Stop Mode

For some multi-threaded targets, GDB supports an optional mode of operation in which you can examine stopped program threads in the debugger while other threads continue to execute freely. This minimizes intrusion when debugging live systems, such as programs where some threads have real-time constraints or must continue to respond to external events. This is referred to as *non-stop* mode.

In non-stop mode, when a thread stops to report a debugging event, *only* that thread is stopped; GDB does not stop other threads as well, in contrast to the all-stop mode behavior. Additionally, execution commands such as **continue** and **step** apply by default only to the current thread in non-stop mode, rather than all threads as in all-stop mode. This allows you to control threads explicitly in ways that are not possible in all-stop mode — for example, stepping one thread while allowing others to run freely, stepping one thread while holding all others stopped, or stepping several threads independently and simultaneously.

To enter non-stop mode, use this sequence of commands before you run or attach to your program:

```
# Enable the async interface.
set target-async 1
# If using the CLI, pagination breaks non-stop.
set pagination off
# Finally, turn it on!
set non-stop on
```

You can use these commands to manipulate the non-stop mode setting:

set non-stop on

Enable selection of non-stop mode.

set non-stop off

Disable selection of non-stop mode.

show non-stop

Show the current non-stop enablement setting.

Note these commands only reflect whether non-stop mode is enabled, not whether the currently-executing program is being run in non-stop mode. In particular, the set non-stop preference is only consulted when GDB starts or connects to the target program, and it is generally not possible to switch modes once debugging has started. Furthermore, since not all targets support non-stop mode, even when you have enabled non-stop mode, GDB may still fall back to all-stop operation by default.

In non-stop mode, all execution commands apply only to the current thread by default. That is, continue only continues one thread. To continue all threads, issue continue -a or c -a.

You can use GDB's background execution commands (see Section 5.4.3 [Background Execution], page 68) to run some threads in the background while you continue to examine or step others from GDB. The MI execution commands (see Section 27.13 [GDB/MI Program Execution], page 312) are always executed asynchronously in non-stop mode.

Suspending execution is done with the interrupt command when running in the background, or Ctrl-c during foreground execution. In all-stop mode, this stops the whole process; but in non-stop mode the interrupt applies only to the current thread. To stop the whole program, use interrupt -a.

Other execution commands do not currently support the -a option.

In non-stop mode, when a thread stops, GDB doesn't automatically make that thread current, as it does in all-stop mode. This is because the thread stop notifications are asynchronous with respect to GDB's command interpreter, and it would be confusing if GDB unexpectedly changed to a different thread just as you entered a command to operate on the previously current thread.

5.4.3 Background Execution

GDB's execution commands have two variants: the normal foreground (synchronous) behavior, and a background (asynchronous) behavior. In foreground execution, GDB waits for the program to report that some thread has stopped before prompting for another command. In background execution, GDB immediately gives a command prompt so that you can issue other commands while your program runs.

You need to explicitly enable asynchronous mode before you can use background execution commands. You can use these commands to manipulate the asynchronous mode setting:

```
set target-async on
```

Enable asynchronous mode.

```
set target-async off
```

Disable asynchronous mode.

```
show target-async
```

Show the current target-async setting.

If the target doesn't support async mode, GDB issues an error message if you attempt to use the background execution commands.

To specify background execution, add a & to the command. For example, the background form of the continue command is continue&, or just c&. The execution commands that accept background execution are:

```
run See Section 4.2 [Starting your Program], page 25.

attach See Section 4.7 [Debugging an Already-running Process], page 31.

step See Section 5.2 [Continuing and Stepping], page 60.

stepi See Section 5.2 [Continuing and Stepping], page 60.

next See Section 5.2 [Continuing and Stepping], page 60.

nexti See Section 5.2 [Continuing and Stepping], page 60.
```

```
continue See Section 5.2 [Continuing and Stepping], page 60.

finish See Section 5.2 [Continuing and Stepping], page 60.

until See Section 5.2 [Continuing and Stepping], page 60.
```

Background execution is especially useful in conjunction with non-stop mode for debugging programs with multiple threads; see Section 5.4.2 [Non-Stop Mode], page 67. However, you can also use these commands in the normal all-stop mode with the restriction that you cannot issue another execution command until the previous one finishes. Examples of commands that are valid in all-stop mode while the program is running include help and info break.

You can interrupt your program while it is running in the background by using the interrupt command.

```
interrupt
interrupt -a
```

Suspend execution of the running program. In all-stop mode, interrupt stops the whole process, but in non-stop mode, it stops only the current thread. To stop the whole program in non-stop mode, use interrupt -a.

5.4.4 Thread-Specific Breakpoints

When your program has multiple threads (see Section 4.10 [Debugging Programs with Multiple Threads], page 35), you can choose whether to set breakpoints on all threads, or on a particular thread.

```
break linespec thread threadno break linespec thread threadno if ...
```

linespec specifies source lines; there are several ways of writing them (see Section 9.2 [Specify Location], page 84), but the effect is always to specify some source line.

Use the qualifier 'thread threadno' with a breakpoint command to specify that you only want GDB to stop the program when a particular thread reaches this breakpoint. threadno is one of the numeric thread identifiers assigned by GDB, shown in the first column of the 'info threads' display.

If you do not specify 'thread threadno' when you set a breakpoint, the breakpoint applies to *all* threads of your program.

You can use the thread qualifier on conditional breakpoints as well; in this case, place 'thread threadno' before or after the breakpoint condition, like this:

```
(gdb) break frik.c:13 thread 28 if bartab > lim
```

5.4.5 Interrupted System Calls

There is an unfortunate side effect when using GDB to debug multi-threaded programs. If one thread stops for a breakpoint, or for some other reason, and another thread is blocked in a system call, then the system call may return prematurely. This is a consequence of the interaction between multiple threads and the signals that GDB uses to implement breakpoints and other events that stop execution.

To handle this problem, your program should check the return value of each system call and react appropriately. This is good programming style anyways.

```
For example, do not write code like this: sleep (10);
```

The call to sleep will return early if a different thread stops at a breakpoint or for some other reason.

```
Instead, write this:
   int unslept = 10;
   while (unslept > 0)
      unslept = sleep (unslept);
```

A system call is allowed to return early, so the system is still conforming to its specification. But gdb does cause your multi-threaded program to behave differently than it would without gdb.

Also, GDB uses internal breakpoints in the thread library to monitor certain events such as thread creation and thread destruction. When such an event happens, a system call in another thread may return prematurely, even though your program does not appear to stop.

6 Running programs backward

When you are debugging a program, it is not unusual to realize that you have gone too far, and some event of interest has already happened. If the target environment supports it, GDB can allow you to "rewind" the program by running it backward.

A target environment that supports reverse execution should be able to "undo" the changes in machine state that have taken place as the program was executing normally. Variables, registers etc. should revert to their previous values. Obviously this requires a great deal of sophistication on the part of the target environment; not all target environments can support reverse execution.

When a program is executed in reverse, the instructions that have most recently been executed are "un-executed", in reverse order. The program counter runs backward, following the previous thread of execution in reverse. As each instruction is "un-executed", the values of memory and/or registers that were changed by that instruction are reverted to their previous states. After executing a piece of source code in reverse, all side effects of that code should be "undone", and all variables should be returned to their prior values¹.

If you are debugging in a target environment that supports reverse execution, GDB provides the following commands.

```
reverse-continue [ignore-count]
rc [ignore-count]
```

Beginning at the point where your program last stopped, start executing in reverse. Reverse execution will stop for breakpoints and synchronous exceptions (signals), just like normal execution. Behavior of asynchronous signals depends on the target environment.

reverse-step [count]

Run the program backward until control reaches the start of a different source line; then stop it, and return control to GDB.

Like the step command, reverse-step will only stop at the beginning of a source line. It "un-executes" the previously executed source line. If the previous source line included calls to debuggable functions, reverse-step will step (backward) into the called function, stopping at the beginning of the *last* statement in the called function (typically a return statement).

Also, as with the step command, if non-debuggable functions are called, reverse-step will run thru them backward without stopping.

reverse-stepi [count]

Reverse-execute one machine instruction. Note that the instruction to be reverse-executed is *not* the one pointed to by the program counter, but the

Note that some side effects are easier to undo than others. For instance, memory and registers are relatively easy, but device I/O is hard. Some targets may be able undo things like device I/O, and some may not.

The contract between GDB and the reverse executing target requires only that the target do something reasonable when GDB tells it to execute backwards, and then report the results back to GDB. Whatever the target reports back to GDB, GDB will report back to the user. GDB assumes that the memory and registers that the target reports are in a consistant state, but GDB accepts whatever it is given.

instruction executed prior to that one. For instance, if the last instruction was a jump, reverse-stepi will take you back from the destination of the jump to the jump instruction itself.

reverse-next [count]

Run backward to the beginning of the previous line executed in the current (innermost) stack frame. If the line contains function calls, they will be "unexecuted" without stopping. Starting from the first line of a function, reversenext will take you back to the caller of that function, before the function was called, just as the normal next command would take you from the last line of a function back to its return to its caller.

reverse-nexti [count]

Like nexti, reverse-nexti executes a single instruction in reverse, except that called functions are "un-executed" atomically. That is, if the previously executed instruction was a return from another function, reverse-nexti will continue to execute in reverse until the call to that function (from the current stack frame) is reached.

reverse-finish

Just as the finish command takes you to the point where the current function returns, reverse-finish takes you to the point where it was called. Instead of ending up at the end of the current function invocation, you end up at the beginning.

set exec-direction

Set the direction of target execution.

set exec-direction reverse

GDB will perform all execution commands in reverse, until the exec-direction mode is changed to "forward". Affected commands include step, stepi, next, nexti, continue, and finish. The return command cannot be used in reverse mode.

set exec-direction forward

GDB will perform all execution commands in the normal fashion. This is the default.

² Unless the code is too heavily optimized.

7 Recording Inferior's Execution and Replaying It

On some platforms, GDB provides a special process record and replay target that can record a log of the process execution, and replay it later with both forward and reverse execution commands.

When this target is in use, if the execution log includes the record for the next instruction, GDB will debug in *replay mode*. In the replay mode, the inferior does not really execute code instructions. Instead, all the events that normally happen during code execution are taken from the execution log. While code is not really executed in replay mode, the values of registers (including the program counter register) and the memory of the inferior are still changed as they normally would. Their contents are taken from the execution log.

If the record for the next instruction is not in the execution log, GDB will debug in record mode. In this mode, the inferior executes normally, and GDB records the execution log for future replay.

The process record and replay target supports reverse execution (see Chapter 6 [Reverse Execution], page 71), even if the platform on which the inferior runs does not. However, the reverse execution is limited in this case by the range of the instructions recorded in the execution log. In other words, reverse execution on platforms that don't support it directly can only be done in the replay mode.

When debugging in the reverse direction, GDB will work in replay mode as long as the execution log includes the record for the previous instruction; otherwise, it will work in record mode, if the platform supports reverse execution, or stop if not.

For architecture environments that support process record and replay, GDB provides the following commands:

target record

This command starts the process record and replay target. The process record and replay target can only debug a process that is already running. Therefore, you need first to start the process with the *run* or *start* commands, and then start the recording with the *target record* command.

Both record and rec are aliases of target record.

Displaced stepping (see Appendix C [displaced stepping], page 397) will be automatically disabled when process record and replay target is started. That's because the process record and replay target doesn't support displaced stepping.

If the inferior is in the non-stop mode (see Section 5.4.2 [Non-Stop Mode], page 67) or in the asynchronous execution mode (see Section 5.4.3 [Background Execution], page 68), the process record and replay target cannot be started because it doesn't support these two modes.

record stop

Stop the process record and replay target. When process record and replay target stops, the entire execution log will be deleted and the inferior will either be terminated, or will remain in its final state.

When you stop the process record and replay target in record mode (at the end of the execution log), the inferior will be stopped at the next instruction that would have been recorded. In other words, if you record for a while and

then stop recording, the inferior process will be left in the same state as if the recording never happened.

On the other hand, if the process record and replay target is stopped while in replay mode (that is, not at the end of the execution log, but at some earlier point), the inferior process will become "live" at that earlier state, and it will then be possible to continue the usual "live" debugging of the process from that state.

When the inferior process exits, or GDB detaches from it, process record and replay target will automatically stop itself.

set record insn-number-max limit

Set the limit of instructions to be recorded. Default value is 200000.

If *limit* is a positive number, then GDB will start deleting instructions from the log once the number of the record instructions becomes greater than *limit*. For every new recorded instruction, GDB will delete the earliest recorded instruction to keep the number of recorded instructions at the limit. (Since deleting recorded instructions loses information, GDB lets you control what happens when the limit is reached, by means of the stop-at-limit option, described below.)

If *limit* is zero, GDB will never delete recorded instructions from the execution log. The number of recorded instructions is unlimited in this case.

show record insn-number-max

Show the limit of instructions to be recorded.

set record stop-at-limit

Control the behavior when the number of recorded instructions reaches the limit. If ON (the default), GDB will stop when the limit is reached for the first time and ask you whether you want to stop the inferior or continue running it and recording the execution log. If you decide to continue recording, each new recorded instruction will cause the oldest one to be deleted.

If this option is OFF, GDB will automatically delete the oldest record to make room for each new one, without asking.

show record stop-at-limit

Show the current setting of stop-at-limit.

info record

Show various statistics about the state of process record and its in-memory execution log buffer, including:

- Whether in record mode or replay mode.
- Lowest recorded instruction number (counting from when the current execution log started recording instructions).
- Highest recorded instruction number.
- Current instruction about to be replayed (if in replay mode).
- Number of instructions contained in the execution log.
- Maximum number of instructions that may be contained in the execution log.

record delete

When record target runs in replay mode ("in the past"), delete the subsequent execution log and begin to record a new execution log starting from the current address. This means you will abandon the previously recorded "future" and begin recording a new "future".

8 Examining the Stack

When your program has stopped, the first thing you need to know is where it stopped and how it got there.

Each time your program performs a function call, information about the call is generated. That information includes the location of the call in your program, the arguments of the call, and the local variables of the function being called. The information is saved in a block of data called a *stack frame*. The stack frames are allocated in a region of memory called the *call stack*.

When your program stops, the GDB commands for examining the stack allow you to see all of this information.

One of the stack frames is *selected* by gdb and many gdb commands refer implicitly to the selected frame. In particular, whenever you ask gdb for the value of a variable in your program, the value is found in the selected frame. There are special gdb commands to select whichever frame you are interested in. See Section 8.3 [Selecting a Frame], page 80.

When your program stops, GDB automatically selects the currently executing frame and describes it briefly, similar to the frame command (see Section 8.4 [Information about a Frame], page 81).

8.1 Stack Frames

The call stack is divided up into contiguous pieces called *stack frames*, or *frames* for short; each frame is the data associated with one call to one function. The frame contains the arguments given to the function, the function's local variables, and the address at which the function is executing.

When your program is started, the stack has only one frame, that of the function main. This is called the *initial* frame or the *outermost* frame. Each time a function is called, a new frame is made. Each time a function returns, the frame for that function invocation is eliminated. If a function is recursive, there can be many frames for the same function. The frame for the function in which execution is actually occurring is called the *innermost* frame. This is the most recently created of all the stack frames that still exist.

Inside your program, stack frames are identified by their addresses. A stack frame consists of many bytes, each of which has its own address; each kind of computer has a convention for choosing one byte whose address serves as the address of the frame. Usually this address is kept in a register called the *frame pointer register* (see Section 10.11 [Registers], page 110) while execution is going on in that frame.

GDB assigns numbers to all existing stack frames, starting with zero for the innermost frame, one for the frame that called it, and so on upward. These numbers do not really exist in your program; they are assigned by GDB to give you a way of designating stack frames in GDB commands.

Some compilers provide a way to compile functions so that they operate without stack frames. (For example, the GCC option

'-fomit-frame-pointer'

generates functions without a frame.) This is occasionally done with heavily used library functions to save the frame setup time. GDB has limited facilities for dealing with

these function invocations. If the innermost function invocation has no stack frame, GDB nevertheless regards it as though it had a separate frame, which is numbered zero as usual, allowing correct tracing of the function call chain. However, GDB has no provision for frameless functions elsewhere in the stack.

frame args

The frame command allows you to move from one stack frame to another, and to print the stack frame you select. args may be either the address of the frame or the stack frame number. Without an argument, frame prints the current stack frame.

select-frame

The select-frame command allows you to move from one stack frame to another without printing the frame. This is the silent version of frame.

8.2 Backtraces

A backtrace is a summary of how your program got where it is. It shows one line per frame, for many frames, starting with the currently executing frame (frame zero), followed by its caller (frame one), and on up the stack.

backtrace

bt

Print a backtrace of the entire stack: one line per frame for all frames in the stack.

You can stop the backtrace at any time by typing the system interrupt character, normally Ctrl-c.

backtrace n

bt n Similar, but print only the innermost n frames.

backtrace -n

bt -n Similar, but print only the outermost n frames.

backtrace full

bt full

bt full n

bt full -n

Print the values of the local variables also. n specifies the number of frames to print, as described above.

The names where and info stack (abbreviated info s) are additional aliases for backtrace.

In a multi-threaded program, GDB by default shows the backtrace only for the current thread. To display the backtrace for several or all of the threads, use the command thread apply (see Section 4.10 [Threads], page 35). For example, if you type thread apply all backtrace, GDB will display the backtrace for all the threads; this is handy when you debug a core dump of a multi-threaded program.

Each line in the backtrace shows the frame number and the function name. The program counter value is also shown—unless you use set print address off. The backtrace also shows the source file name and line number, as well as the arguments to the function. The program counter value is omitted if it is at the beginning of the code for that line number.

Here is an example of a backtrace. It was made with the command 'bt 3', so it shows the innermost three frames.

```
#0 m4_traceon (obs=0x24eb0, argc=1, argv=0x2b8c8)
   at builtin.c:993
#1 0x6e38 in expand_macro (sym=0x2b600, data=...) at macro.c:242
#2 0x6840 in expand_token (obs=0x0, t=177664, td=0xf7fffb08)
   at macro.c:71
(More stack frames follow...)
```

The display for frame zero does not begin with a program counter value, indicating that your program has stopped at the beginning of the code for line 993 of builtin.c.

The value of parameter data in frame 1 has been replaced by By default, gdb prints the value of a parameter only if it is a scalar (integer, pointer, enumeration, etc). See command set print frame-arguments in Section 10.8 [Print Settings], page 102 for more details on how to configure the way function parameter values are printed.

If your program was compiled with optimizations, some compilers will optimize away arguments passed to functions if those arguments are never used after the call. Such optimizations generate code that passes arguments through registers, but doesn't store those arguments in the stack frame. GDB has no way of displaying such arguments in stack frames other than the innermost one. Here's what such a backtrace might look like:

```
#0 m4_traceon (obs=0x24eb0, argc=1, argv=0x2b8c8)
   at builtin.c:993
#1 0x6e38 in expand_macro (sym=<value optimized out>) at macro.c:242
#2 0x6840 in expand_token (obs=0x0, t=<value optimized out>, td=0xf7fffb08)
   at macro.c:71
(More stack frames follow...)
```

The values of arguments that were not saved in their stack frames are shown as '<value optimized out>'.

If you need to display the values of such optimized-out arguments, either deduce that from other variables whose values depend on the one you are interested in, or recompile without optimizations.

Most programs have a standard user entry point—a place where system libraries and startup code transition into user code. For C this is main¹. When GDB finds the entry function in a backtrace it will terminate the backtrace, to avoid tracing into highly system-specific (and generally uninteresting) code.

If you need to examine the startup code, or limit the number of levels in a backtrace, you can change this behavior:

```
set backtrace past-main on
```

Backtraces will continue past the user entry point.

set backtrace past-main off

Backtraces will stop when they encounter the user entry point. This is the default.

show backtrace past-main

Display the current user entry point backtrace policy.

¹ Note that embedded programs (the so-called "free-standing" environment) are not required to have a main function as the entry point. They could even have multiple entry points.

```
set backtrace past-entry set backtrace past-entry on
```

Backtraces will continue past the internal entry point of an application. This entry point is encoded by the linker when the application is built, and is likely before the user entry point main (or equivalent) is called.

set backtrace past-entry off

Backtraces will stop when they encounter the internal entry point of an application. This is the default.

show backtrace past-entry

Display the current internal entry point backtrace policy.

```
set backtrace limit n
set backtrace limit 0
```

Limit the backtrace to *n* levels. A value of zero means unlimited.

show backtrace limit

Display the current limit on backtrace levels.

8.3 Selecting a Frame

Most commands for examining the stack and other data in your program work on whichever stack frame is selected at the moment. Here are the commands for selecting a stack frame; all of them finish by printing a brief description of the stack frame just selected.

frame n

f n Select frame number n. Recall that frame zero is the innermost (currently executing) frame, frame one is the frame that called the innermost one, and so on. The highest-numbered frame is the one for main.

frame addr

f addr Select the frame at address addr. This is useful mainly if the chaining of stack frames has been damaged by a bug, making it impossible for GDB to assign numbers properly to all frames. In addition, this can be useful when your program has multiple stacks and switches between them.

On the MIPS and Alpha architectures, frame needs two addresses to select an arbitrary frame: a stack pointer and a program counter.

up n Move n frames up the stack. For positive numbers n, this advances toward the outermost frame, to higher frame numbers, to frames that have existed longer. n defaults to one.

down n Move n frames down the stack. For positive numbers n, this advances toward the innermost frame, to lower frame numbers, to frames that were created more recently. n defaults to one. You may abbreviate down as do.

All of these commands end by printing two lines of output describing the frame. The first line shows the frame number, the function name, the arguments, and the source file and line number of execution in that frame. The second line shows the text of that source line.

For example:

```
(gdb) up
#1 0x22f0 in main (argc=1, argv=0xf7fffbf4, env=0xf7fffbfc)
    at env.c:10
10          read_input_file (argv[i]);
```

After such a printout, the list command with no arguments prints ten lines centered on the point of execution in the frame. You can also edit the program at the point of execution with your favorite editing program by typing edit. See Section 9.1 [Printing Source Lines], page 83, for details.

```
up-silently n down-silently n
```

These two commands are variants of up and down, respectively; they differ in that they do their work silently, without causing display of the new frame. They are intended primarily for use in GDB command scripts, where the output might be unnecessary and distracting.

8.4 Information About a Frame

There are several other commands to print information about the selected stack frame.

frame

f

When used without any argument, this command does not change which frame is selected, but prints a brief description of the currently selected stack frame. It can be abbreviated f. With an argument, this command is used to select a stack frame. See Section 8.3 [Selecting a Frame], page 80.

info frame

- info f This command prints a verbose description of the selected stack frame, including:
 - the address of the frame
 - the address of the next frame down (called by this frame)
 - the address of the next frame up (caller of this frame)
 - the language in which the source code corresponding to this frame is written
 - the address of the frame's arguments
 - the address of the frame's local variables
 - the program counter saved in it (the address of execution in the caller frame)
 - which registers were saved in the frame

The verbose description is useful when something has gone wrong that has made the stack format fail to fit the usual conventions.

info frame addr

Print a verbose description of the frame at address addr, without selecting that frame. The selected frame remains unchanged by this command. This requires the same kind of address (more than one for some architectures) that you specify in the frame command. See Section 8.3 [Selecting a Frame], page 80.

info args Print the arguments of the selected frame, each on a separate line.

info locals

Print the local variables of the selected frame, each on a separate line. These are all variables (declared either static or automatic) accessible at the point of execution of the selected frame.

info catch

Print a list of all the exception handlers that are active in the current stack frame at the current point of execution. To see other exception handlers, visit the associated frame (using the up, down, or frame commands); then type info catch. See Section 5.1.3 [Setting Catchpoints], page 51.

9 Examining Source Files

GDB can print parts of your program's source, since the debugging information recorded in the program tells GDB what source files were used to build it. When your program stops, GDB spontaneously prints the line where it stopped. Likewise, when you select a stack frame (see Section 8.3 [Selecting a Frame], page 80), GDB prints the line where execution in that frame has stopped. You can print other portions of source files by explicit command.

If you use GDB through its GNU Emacs interface, you may prefer to use Emacs facilities to view source; see Chapter 26 [Using GDB under GNU Emacs], page 287.

9.1 Printing Source Lines

To print lines from a source file, use the list command (abbreviated 1). By default, ten lines are printed. There are several ways to specify what part of the file you want to print; see Section 9.2 [Specify Location], page 84, for the full list.

Here are the forms of the list command most commonly used:

list linenum

Print lines centered around line number *linenum* in the current source file.

list function

Print lines centered around the beginning of function function.

Print more lines. If the last lines printed were printed with a list command, this prints lines following the last lines printed; however, if the last line printed was a solitary line printed as part of displaying a stack frame (see Chapter 8 [Examining the Stack], page 77), this prints lines centered around that line.

list - Print lines just before the lines last printed.

By default, GDB prints ten source lines with any of these forms of the list command. You can change this using set listsize:

set listsize count

Make the list command display *count* source lines (unless the list argument explicitly specifies some other number).

show listsize

Display the number of lines that list prints.

Repeating a list command with RET discards the argument, so it is equivalent to typing just list. This is more useful than listing the same lines again. An exception is made for an argument of '-'; that argument is preserved in repetition so that each repetition moves up in the source file.

In general, the list command expects you to supply zero, one or two *linespecs*. Linespecs specify source lines; there are several ways of writing them (see Section 9.2 [Specify Location], page 84), but the effect is always to specify some source line.

Here is a complete description of the possible arguments for list:

list linespec

Print lines centered around the line specified by linespec.

list first, last

Print lines from *first* to *last*. Both arguments are linespecs. When a list command has two linespecs, and the source file of the second linespec is omitted, this refers to the same source file as the first linespec.

list ,last

Print lines ending with last.

list first,

Print lines starting with first.

list + Print lines just after the lines last printed.

list - Print lines just before the lines last printed.

list As described in the preceding table.

9.2 Specifying a Location

Several GDB commands accept arguments that specify a location of your program's code. Since GDB is a source-level debugger, a location usually specifies some line in the source code; for that reason, locations are also known as *linespecs*.

Here are all the different ways of specifying a code location that gdb understands:

linenum Specifies the line number linenum of the current source file.

-offset

+offset

Specifies the line offset lines before or after the current line. For the list command, the current line is the last one printed; for the breakpoint commands, this is the line at which execution stopped in the currently selected stack frame (see Section 8.1 [Frames], page 77, for a description of stack frames.) When used as the second of the two linespecs in a list command, this specifies the line offset lines up or down from the first linespec.

filename: linenum

Specifies the line linenum in the source file filename.

function Specifies the line that begins the body of the function function. For example, in C, this is the line with the open brace.

filename:function

Specifies the line that begins the body of the function function in the file filename. You only need the file name with a function name to avoid ambiguity when there are identically named functions in different source files.

*address Specifies the program address address. For line-oriented commands, such as list and edit, this specifies a source line that contains address. For break and other breakpoint oriented commands, this can be used to set breakpoints in parts of your program which do not have debugging information or source files.

Here address may be any expression valid in the current working language (see Chapter 15 [Languages], page 147) that specifies a code address. In addition, as a convenience, GDB extends the semantics of expressions used in locations

to cover the situations that frequently happen during debugging. Here are the various forms of address:

expression

Any expression valid in the current working language.

funcaddr

An address of a function or procedure derived from its name. In C, C++, Java, Objective-C, Fortran, minimal, and assembly, this is simply the function's name function (and actually a special case of a valid expression). In Pascal and Modula-2, this is &function. In Ada, this is function'Address (although the Pascal form also works).

This form specifies the address of the function's first instruction, before the stack frame and arguments have been set up.

'filename'::funcaddr

Like funcaddr above, but also specifies the name of the source file explicitly. This is useful if the name of the function does not specify the function unambiguously, e.g., if there are several functions with identical names in different source files.

9.3 Editing Source Files

To edit the lines in a source file, use the edit command. The editing program of your choice is invoked with the current line set to the active line in the program. Alternatively, there are several ways to specify what part of the file you want to print if you want to see other parts of the program:

edit location

Edit the source file specified by location. Editing starts at that *location*, e.g., at the specified source line of the specified file. See Section 9.2 [Specify Location], page 84, for all the possible forms of the *location* argument; here are the forms of the edit command most commonly used:

edit number

Edit the current source file with *number* as the active line number.

edit function

Edit the file containing function at the beginning of its definition.

9.3.1 Choosing your Editor

You can customize GDB to use any editor you want¹. By default, it is '/bin/ex', but you can change this by setting the environment variable EDITOR before using GDB. For example, to configure GDB to use the vi editor, you could use these commands with the sh shell:

```
EDITOR=/usr/bin/vi
export EDITOR
gdb ...
```

¹ The only restriction is that your editor (say ex), recognizes the following command-line syntax: ex +number file

The optional numeric value +number specifies the number of the line in the file where to start editing.

```
or in the csh shell,
setenv EDITOR /usr/bin/vi
gdb ...
```

9.4 Searching Source Files

There are two commands for searching through the current source file for a regular expression.

```
forward-search regexp
search regexp
```

The command 'forward-search regexp' checks each line, starting with the one following the last line listed, for a match for regexp. It lists the line that is found. You can use the synonym 'search regexp' or abbreviate the command name as fo.

reverse-search regexp

The command 'reverse-search regexp' checks each line, starting with the one before the last line listed and going backward, for a match for regexp. It lists the line that is found. You can abbreviate this command as rev.

9.5 Specifying Source Directories

Executable programs sometimes do not record the directories of the source files from which they were compiled, just the names. Even when they do, the directories could be moved between the compilation and your debugging session. GDB has a list of directories to search for source files; this is called the *source path*. Each time GDB wants a source file, it tries all the directories in the list, in the order they are present in the list, until it finds a file with the desired name.

For example, suppose an executable references the file '/usr/src/foo-1.0/lib/foo.c', and our source path is '/mnt/cross'. The file is first looked up literally; if this fails, '/mnt/cross/usr/src/foo-1.0/lib/foo.c' is tried; if this fails, '/mnt/cross/foo.c' is opened; if this fails, an error message is printed. GDB does not look up the parts of the source file name, such as '/mnt/cross/src/foo-1.0/lib/foo.c'. Likewise, the subdirectories of the source path are not searched: if the source path is '/mnt/cross', and the binary refers to 'foo.c', GDB would not find it under '/mnt/cross/usr/src/foo-1.0/lib'.

Plain file names, relative file names with leading directories, file names containing dots, etc. are all treated as described above; for instance, if the source path is '/mnt/cross', and the source file is recorded as '../lib/foo.c', GDB would first try '../lib/foo.c', then '/mnt/cross/../lib/foo.c', and after that—'/mnt/cross/foo.c'.

Note that the executable search path is *not* used to locate the source files.

Whenever you reset or rearrange the source path, GDB clears out any information it has cached about where source files are found and where each line is in the file.

When you start GDB, its source path includes only 'cdir' and 'cwd', in that order. To add other directories, use the directory command.

The search path is used to find both program source files and GDB script files (read using the '-command' option and 'source' command).

In addition to the source path, GDB provides a set of commands that manage a list of source path substitution rules. A substitution rule specifies how to rewrite source directories stored in the program's debug information in case the sources were moved to a different directory between compilation and debugging. A rule is made of two strings, the first specifying what needs to be rewritten in the path, and the second specifying how it should be rewritten. In [set substitute-path], page 88, we name these two parts from and to respectively. GDB does a simple string replacement of from with to at the start of the directory part of the source file name, and uses that result instead of the original file name to look up the sources.

Using the previous example, suppose the 'foo-1.0' tree has been moved from '/usr/src' to '/mnt/cross', then you can tell GDB to replace '/usr/src' in all source path names with '/mnt/cross'. The first lookup will then be '/mnt/cross/foo-1.0/lib/foo.c' in place of the original location of '/usr/src/foo-1.0/lib/foo.c'. To define a source path substitution rule, use the set substitute-path command (see [set substitute-path], page 88).

To avoid unexpected substitution results, a rule is applied only if the from part of the directory name ends at a directory separator. For instance, a rule substituting '/usr/source' into '/mnt/cross' will be applied to '/usr/source/foo-1.0' but not to '/usr/sourceware/foo-2.0'. And because the substitution is applied only at the beginning of the directory name, this rule will not be applied to '/root/usr/source/baz.c' either.

In many cases, you can achieve the same result using the directory command. However, set substitute-path can be more efficient in the case where the sources are organized in a complex tree with multiple subdirectories. With the directory command, you need to add each subdirectory of your project. If you moved the entire tree while preserving its internal organization, then set substitute-path allows you to direct the debugger to all the sources with one single command.

set substitute-path is also more than just a shortcut command. The source path is only used if the file at the original location no longer exists. On the other hand, set substitute-path modifies the debugger behavior to look at the rewritten location instead. So, if for any reason a source file that is not relevant to your executable is located at the original location, a substitution rule is the only method available to point GDB at the new location.

You can configure a default source path substitution rule by configuring GDB with the '--with-relocated-sources=dir' option. The dir should be the name of a directory under GDB's configured prefix (set with '--prefix' or '--exec-prefix'), and directory names in debug information under dir will be adjusted automatically if the installed GDB is moved to a new location. This is useful if GDB, libraries or executables with debug information and corresponding source code are being moved together.

directory dirname ... dir dirname ...

Add directory dirname to the front of the source path. Several directory names may be given to this command, separated by ':' (';' on MS-DOS and MS-Windows, where ':' usually appears as part of absolute file names) or white-space. You may specify a directory that is already in the source path; this moves it forward, so GDB searches it sooner.

You can use the string '\$cdir' to refer to the compilation directory (if one is recorded), and '\$cwd' to refer to the current working directory. '\$cwd' is not the same as '.'—the former tracks the current working directory as it changes during your GDB session, while the latter is immediately expanded to the current directory at the time you add an entry to the source path.

directory

Reset the source path to its default value ('\$cdir:\$cwd' on Unix systems). This requires confirmation.

show directories

Print the source path: show which directories it contains.

set substitute-path from to

Define a source path substitution rule, and add it at the end of the current list of existing substitution rules. If a rule with the same *from* was already defined, then the old rule is also deleted.

For example, if the file '/foo/bar/baz.c' was moved to '/mnt/cross/baz.c', then the command

```
(gdb) set substitute-path /usr/src /mnt/cross
```

will tell GDB to replace '/usr/src' with '/mnt/cross', which will allow GDB to find the file 'baz.c' even though it was moved.

In the case when more than one substitution rule have been defined, the rules are evaluated one by one in the order where they have been defined. The first one matching, if any, is selected to perform the substitution.

For instance, if we had entered the following commands:

```
(gdb) set substitute-path /usr/src/include /mnt/include (gdb) set substitute-path /usr/src /mnt/src
```

GDB would then rewrite '/usr/src/include/defs.h' into '/mnt/include/defs.h' by using the first rule. However, it would use the second rule to rewrite '/usr/src/lib/foo.c' into '/mnt/src/lib/foo.c'.

unset substitute-path [path]

If a path is specified, search the current list of substitution rules for a rule that would rewrite that path. Delete that rule if found. A warning is emitted by the debugger if no rule could be found.

If no path is specified, then all substitution rules are deleted.

show substitute-path [path]

If a path is specified, then print the source path substitution rule which would rewrite that path, if any.

If no path is specified, then print all existing source path substitution rules.

If your source path is cluttered with directories that are no longer of interest, GDB may sometimes cause confusion by finding the wrong versions of source. You can correct the situation as follows:

- 1. Use directory with no argument to reset the source path to its default value.
- 2. Use directory with suitable arguments to reinstall the directories you want in the source path. You can add all the directories in one command.

9.6 Source and Machine Code

You can use the command info line to map source lines to program addresses (and vice versa), and the command disassemble to display a range of addresses as machine instructions. You can use the command set disassemble-next-line to set whether to disassemble next source line when execution stops. When run under GNU Emacs mode, the info line command causes the arrow to point to the line specified. Also, info line prints addresses in symbolic form as well as hex.

info line linespec

Print the starting and ending addresses of the compiled code for source line *line-spec*. You can specify source lines in any of the ways documented in Section 9.2 [Specify Location], page 84.

For example, we can use info line to discover the location of the object code for the first line of function m4_changequote:

```
(gdb) info line m4_changequote
Line 895 of "builtin.c" starts at pc 0x634c and ends at 0x6350.
```

We can also inquire (using *addr as the form for linespec) what source line covers a particular address:

```
(gdb) info line *0x63ff
Line 926 of "builtin.c" starts at pc 0x63e4 and ends at 0x6404.
```

After info line, the default address for the x command is changed to the starting address of the line, so that 'x/i' is sufficient to begin examining the machine code (see Section 10.6 [Examining Memory], page 99). Also, this address is saved as the value of the convenience variable \$_ (see Section 10.10 [Convenience Variables], page 109).

```
disassemble /m
disassemble /r
```

This specialized command dumps a range of memory as machine instructions. It can also print mixed source+disassembly by specifying the /m modifier and print the raw instructions in hex as well as in symbolic form by specifying the /r. The default memory range is the function surrounding the program counter of the selected frame. A single argument to this command is a program counter value; GDB dumps the function surrounding this value. When two arguments are given, they should be separated by a comma, possibly surrounded by whitespace. The arguments specify a range of addresses (first inclusive, second exclusive) to dump. In that case, the name of the function is also printed (since there could be several functions in the given range).

The argument(s) can be any expression yielding a numeric value, such as '0x32c4', '&main+10' or '\$pc - 8'.

If the range of memory being disassembled contains current program counter, the instruction at that location is shown with a => marker.

The following example shows the disassembly of a range of addresses of HP PA-RISC 2.0 code:

```
(gdb) disas 0x32c4, 0x32e4
Dump of assembler code from 0x32c4 to 0x32e4:
```

```
0x32c4 <main+204>:
                           addil 0,dp
   0x32c8 <main+208>:
                          ldw 0x22c(sr0,r1),r26
                          ldil 0x3000,r31
   0x32cc <main+212>:
   0x32d0 <main+216>:
                          ble 0x3f8(sr4,r31)
   0x32d4 <main+220>:
                          ldo 0(r31),rp
   0x32d8 <main+224>:
                          addil -0x800,dp
   0x32dc <main+228>:
                          ldo 0x588(r1),r26
  0x32e0 <main+232>:
                           ldil 0x3000,r31
End of assembler dump.
```

Here is an example showing mixed source+assembly for Intel x86, when the program is stopped just after function prologue:

```
(gdb) disas /m main
Dump of assembler code for function main:
   0x08048330 <+0>:
                              %ebp
                              %esp,%ebp
   0x08048331 <+1>:
                       mov
   0x08048333 <+3>:
                       sub
                              $0x8, %esp
   0x08048336 <+6>:
                       and
                              $0xfffffff0, %esp
   0x08048339 <+9>:
                              $0x10, %esp
                       sub
         printf ("Hello.\n");
=> 0x0804833c <+12>: movl
                              $0x8048440,(%esp)
   0x08048343 <+19>:
                      call
                              0x8048284 <puts@plt>
7
         return 0;
8
       }
   0x08048348 <+24>:
                       mov
                              $0x0, %eax
   0x0804834d <+29>:
                       leave
   0x0804834e <+30>:
                       ret
```

End of assembler dump.

Some architectures have more than one commonly-used set of instruction mnemonics or other syntax.

For programs that were dynamically linked and use shared libraries, instructions that call functions or branch to locations in the shared libraries might show a seemingly bogus location—it's actually a location of the relocation table. On some architectures, GDB might be able to resolve these to actual function names.

set disassembly-flavor instruction-set

Select the instruction set to use when disassembling the program via the disassemble or x/i commands.

Currently this command is only defined for the Intel x86 family. You can set instruction-set to either intel or att. The default is att, the AT&T flavor used by default by Unix assemblers for x86-based targets.

show disassembly-flavor

Show the current setting of the disassembly flavor.

set disassemble-next-line show disassemble-next-line

Control whether or not GDB will disassemble the next source line or instruction when execution stops. If ON, GDB will display disassembly of the next source line when execution of the program being debugged stops. This is *in addition*

to displaying the source line itself, which GDB always does if possible. If the next source line cannot be displayed for some reason (e.g., if GDB cannot find the source file, or there's no line info in the debug info), GDB will display disassembly of the next instruction instead of showing the next source line. If AUTO, GDB will display disassembly of next instruction only if the source line cannot be displayed. This setting causes GDB to display some feedback when you step through a function with no line info or whose source file is unavailable. The default is OFF, which means never display the disassembly of the next line or instruction.

10 Examining Data

The usual way to examine data in your program is with the **print** command (abbreviated p), or its synonym **inspect**. It evaluates and prints the value of an expression of the language your program is written in (see Chapter 15 [Using GDB with Different Languages], page 147).

```
print expr
print /f expr
```

expr is an expression (in the source language). By default the value of expr is printed in a format appropriate to its data type; you can choose a different format by specifying 'f', where f is a letter specifying the format; see Section 10.5 [Output Formats], page 97.

print

print /f If you omit expr, GDB displays the last value again (from the value history; see Section 10.9 [Value History], page 108). This allows you to conveniently inspect the same value in an alternative format.

A more low-level way of examining data is with the x command. It examines data in memory at a specified address and prints it in a specified format. See Section 10.6 [Examining Memory], page 99.

If you are interested in information about types, or about how the fields of a struct or a class are declared, use the ptype exp command rather than print. See Chapter 16 [Examining the Symbol Table], page 175.

10.1 Expressions

print and many other GDB commands accept an expression and compute its value. Any kind of constant, variable or operator defined by the programming language you are using is valid in an expression in GDB. This includes conditional expressions, function calls, casts, and string constants. It also includes preprocessor macros, if you compiled your program to include this information; see Section 4.1 [Compilation], page 25.

GDB supports array constants in expressions input by the user. The syntax is {element, element...}. For example, you can use the command print {1, 2, 3} to create an array of three integers. If you pass an array to a function or assign it to a program variable, GDB copies the array to memory that is malloced in the target program.

Because C is so widespread, most of the expressions shown in examples in this manual are in C. See Chapter 15 [Using GDB with Different Languages], page 147, for information on how to use expressions in other languages.

In this section, we discuss operators that you can use in GDB expressions regardless of your programming language.

Casts are supported in all languages, not just in C, because it is so useful to cast a number into a pointer in order to examine a structure at that address in memory.

GDB supports these operators, in addition to those common to programming languages:

© '©' is a binary operator for treating parts of memory as arrays. See Section 10.4 [Artificial Arrays], page 97, for more information.

":: allows you to specify a variable in terms of the file or function where it is defined. See Section 10.3 [Program Variables], page 95.

{type} addr

::

Refers to an object of type type stored at address addr in memory. addr may be any expression whose value is an integer or pointer (but parentheses are required around binary operators, just as in a cast). This construct is allowed regardless of what kind of data is normally supposed to reside at addr.

10.2 Ambiguous Expressions

Expressions can sometimes contain some ambiguous elements. For instance, some programming languages (notably Ada, C++ and Objective-C) permit a single function name to be defined several times, for application in different contexts. This is called *overloading*. Another example involving Ada is generics. A *generic package* is similar to C++ templates and is typically instantiated several times, resulting in the same function name being defined in different contexts.

In some cases and depending on the language, it is possible to adjust the expression to remove the ambiguity. For instance in C++, you can specify the signature of the function you want to break on, as in *break function (types)*. In Ada, using the fully qualified name of your function often makes the expression unambiguous as well.

When an ambiguity that needs to be resolved is detected, the debugger has the capability to display a menu of numbered choices for each possibility, and then waits for the selection with the prompt '>'. The first option is always '[0] cancel', and typing $0\ RET$ aborts the current command. If the command in which the expression was used allows more than one choice to be selected, the next option in the menu is '[1] all', and typing $1\ RET$ selects all possible choices.

For example, the following session excerpt shows an attempt to set a breakpoint at the overloaded symbol String::after. We choose three particular definitions of that function name:

```
(gdb) b String::after
[0] cancel
[1] all
[2] file:String.cc; line number:867
[3] file:String.cc; line number:860
[4] file:String.cc; line number:875
[5] file:String.cc; line number:853
[6] file:String.cc; line number:846
[7] file:String.cc; line number:735
> 2 4 6
Breakpoint 1 at 0xb26c: file String.cc, line 867.
Breakpoint 2 at 0xb344: file String.cc, line 875.
Breakpoint 3 at Oxafcc: file String.cc, line 846.
Multiple breakpoints were set.
Use the "delete" command to delete unwanted
breakpoints.
(gdb)
```

set multiple-symbols mode

This option allows you to adjust the debugger behavior when an expression is ambiguous.

By default, mode is set to all. If the command with which the expression is used allows more than one choice, then GDB automatically selects all possible choices. For instance, inserting a breakpoint on a function using an ambiguous name results in a breakpoint inserted on each possible match. However, if a unique choice must be made, then GDB uses the menu to help you disambiguate the expression. For instance, printing the address of an overloaded function will result in the use of the menu.

When mode is set to ask, the debugger always uses the menu when an ambiguity is detected.

Finally, when *mode* is set to cancel, the debugger reports an error due to the ambiguity and the command is aborted.

show multiple-symbols

Show the current value of the multiple-symbols setting.

10.3 Program Variables

The most common kind of expression to use is the name of a variable in your program.

Variables in expressions are understood in the selected stack frame (see Section 8.3 [Selecting a Frame], page 80); they must be either:

• global (or file-static)

or

• visible according to the scope rules of the programming language from the point of execution in that frame

This means that in the function

```
foo (a)
    int a;
{
    bar (a);
    {
     int b = test ();
     bar (b);
    }
}
```

you can examine and use the variable a whenever your program is executing within the function foo, but you can only use or examine the variable b while your program is executing inside the block where b is declared.

There is an exception: you can refer to a variable or function whose scope is a single source file even if the current execution point is not in this file. But it is possible to have more than one such variable or function with the same name (in different source files). If that happens, referring to that name has unpredictable effects. If you wish, you can specify a static variable in a particular function or file, using the colon-colon (::) notation:

```
file::variable
function::variable
```

Here file or function is the name of the context for the static variable. In the case of file names, you can use quotes to make sure GDB parses the file name as a single word—for example, to print a global value of x defined in 'f2.c':

```
(gdb) p 'f2.c'::x
```

This use of '::' is very rarely in conflict with the very similar use of the same notation in C++. GDB also supports use of the C++ scope resolution operator in GDB expressions.

Warning: Occasionally, a local variable may appear to have the wrong value at certain points in a function—just after entry to a new scope, and just before exit.

You may see this problem when you are stepping by machine instructions. This is because, on most machines, it takes more than one instruction to set up a stack frame (including local variable definitions); if you are stepping by machine instructions, variables may appear to have the wrong values until the stack frame is completely built. On exit, it usually also takes more than one machine instruction to destroy a stack frame; after you begin stepping through that group of instructions, local variable definitions may be gone.

This may also happen when the compiler does significant optimizations. To be sure of always seeing accurate values, turn off all optimization when compiling.

Another possible effect of compiler optimizations is to optimize unused variables out of existence, or assign variables to registers (as opposed to memory addresses). Depending on the support for such cases offered by the debug info format used by the compiler, gdb might not be able to display values for such local variables. If that happens, gdb will print a message like this:

```
No symbol "foo" in current context.
```

To solve such problems, either recompile without optimizations, or use a different debug info format, if the compiler supports several such formats. For example, GCC, the GNU C/C++ compiler, usually supports the '-gstabs+' option. '-gstabs+' produces debug info in a format that is superior to formats such as COFF. You may be able to use DWARF 2 ('-gdwarf-2'), which is also an effective form for debug info. See Section "Options for Debugging Your Program or GCC" in *Using the GNU Compiler Collection (GCC)*. See Section 15.4.1 [C and C++], page 151, for more information about debug info formats that are best suited to C++ programs.

If you ask to print an object whose contents are unknown to GDB, e.g., because its data type is not completely specified by the debug information, GDB will say '<incomplete type>'. See Chapter 16 [Symbols], page 175, for more about this.

Strings are identified as arrays of char values without specified signedness. Arrays of either signed char or unsigned char get printed as arrays of 1 byte sized integers. – fsigned-char or -funsigned-char GCC options have no effect as GDB defines literal string type "char" as char without a sign. For program code

```
char var0[] = "A";
signed char var1[] = "A";
You get during debugging
  (gdb) print var0
$1 = "A"
  (gdb) print var1
$2 = {65 'A', 0 '\0'}
```

10.4 Artificial Arrays

It is often useful to print out several successive objects of the same type in memory; a section of an array, or an array of dynamically determined size for which only a pointer exists in the program.

You can do this by referring to a contiguous span of memory as an artificial array, using the binary operator '@'. The left operand of '@' should be the first element of the desired array and be an individual object. The right operand should be the desired length of the array. The result is an array value whose elements are all of the type of the left argument. The first element is actually the left argument; the second element comes from bytes of memory immediately following those that hold the first element, and so on. Here is an example. If a program says

```
int *array = (int *) malloc (len * sizeof (int));
you can print the contents of array with
   p *array@len
```

The left operand of '@' must reside in memory. Array values made with '@' in this way behave just like other arrays in terms of subscripting, and are coerced to pointers when used in expressions. Artificial arrays most often appear in expressions via the value history (see Section 10.9 [Value History], page 108), after printing one out.

Another way to create an artificial array is to use a cast. This re-interprets a value as if it were an array. The value need not be in memory:

```
(gdb) p/x (short[2])0x12345678
$1 = {0x1234, 0x5678}
```

As a convenience, if you leave the array length out (as in '(type[]) value') GDB calculates the size to fill the value (as 'sizeof(value)/sizeof(type)':

```
(gdb) p/x (short[])0x12345678
$2 = {0x1234, 0x5678}
```

Sometimes the artificial array mechanism is not quite enough; in moderately complex data structures, the elements of interest may not actually be adjacent—for example, if you are interested in the values of pointers in an array. One useful work-around in this situation is to use a convenience variable (see Section 10.10 [Convenience Variables], page 109) as a counter in an expression that prints the first interesting value, and then repeat that expression via RET. For instance, suppose you have an array dtab of pointers to structures, and you are interested in the values of a field fv in each structure. Here is an example of what you might type:

```
set $i = 0
p dtab[$i++]->fv
RET
RET
...
```

10.5 Output Formats

By default, GDB prints a value according to its data type. Sometimes this is not what you want. For example, you might want to print a number in hex, or a pointer in decimal. Or you might want to view data in memory at a certain address as a character string or as an instruction. To do these things, specify an *output format* when you print a value.

The simplest use of output formats is to say how to print a value already computed. This is done by starting the arguments of the print command with a slash and a format letter. The format letters supported are:

- x Regard the bits of the value as an integer, and print the integer in hexadecimal.
- d Print as integer in signed decimal.
- u Print as integer in unsigned decimal.
- o Print as integer in octal.
- t Print as integer in binary. The letter 't' stands for "two". 1
- a Print as an address, both absolute in hexadecimal and as an offset from the nearest preceding symbol. You can use this format used to discover where (in what function) an unknown address is located:

```
(gdb) p/a 0x54320
$3 = 0x54320 <_initialize_vx+396>
```

The command info symbol 0x54320 yields similar results. See Chapter 16 [Symbols], page 175.

c Regard as an integer and print it as a character constant. This prints both the numerical value and its character representation. The character representation is replaced with the octal escape '\nnn' for characters outside the 7-bit ASCII range.

Without this format, GDB displays char, unsigned char, and signed char data as character constants. Single-byte members of vectors are displayed as integer data.

- f Regard the bits of the value as a floating point number and print using typical floating point syntax.
- s Regard as a string, if possible. With this format, pointers to single-byte data are displayed as null-terminated strings and arrays of single-byte data are displayed as fixed-length strings. Other values are displayed in their natural types.

Without this format, GDB displays pointers to and arrays of char, unsigned char, and signed char as strings. Single-byte members of a vector are displayed as an integer array.

Print using the 'raw' formatting. By default, GDB will use a type-specific pretty-printer. The 'r' format bypasses any pretty-printer which might exist for the value's type.

For example, to print the program counter in hex (see Section 10.11 [Registers], page 110), type

```
p/x $pc
```

Note that no space is required before the slash; this is because command names in GDB cannot contain a slash.

¹ 'b' cannot be used because these format letters are also used with the x command, where 'b' stands for "byte"; see Section 10.6 [Examining Memory], page 99.

To reprint the last value in the value history with a different format, you can use the print command with just a format and no expression. For example, 'p/x' reprints the last value in hex.

10.6 Examining Memory

You can use the command x (for "examine") to examine memory in any of several formats, independently of your program's data types.

x/nfu addr
x addr

x Use the x command to examine memory.

n, f, and u are all optional parameters that specify how much memory to display and how to format it; addr is an expression giving the address where you want to start displaying memory. If you use defaults for nfu, you need not type the slash '/'. Several commands set convenient defaults for addr.

n, the repeat count

The repeat count is a decimal integer; the default is 1. It specifies how much memory (counting by units u) to display.

f, the display format

The display format is one of the formats used by print ('x', 'd', 'u', 'o', 't', 'a', 'c', 'f', 's'), and in addition 'i' (for machine instructions). The default is 'x' (hexadecimal) initially. The default changes each time you use either x or print.

u, the unit size

The unit size is any of

b Bytes.

h Halfwords (two bytes).

w Words (four bytes). This is the initial default.

g Giant words (eight bytes).

Each time you specify a unit size with x, that size becomes the default unit the next time you use x. (For the 's' and 'i' formats, the unit size is ignored and is normally not written.)

addr, starting display address

addr is the address where you want GDB to begin displaying memory. The expression need not have a pointer value (though it may); it is always interpreted as an integer address of a byte of memory. See Section 10.1 [Expressions], page 93, for more information on expressions. The default for addr is usually just after the last address examined—but several other commands also set the default address: info breakpoints (to the address of the last breakpoint listed), info line (to the starting address of a line), and print (if you use it to display a value from memory).

For example, 'x/3uh 0x54320' is a request to display three halfwords (h) of memory, formatted as unsigned decimal integers ('u'), starting at address 0x54320. 'x/4xw \$sp' prints the four words ('w') of memory above the stack pointer (here, '\$sp'; see Section 10.11 [Registers], page 110) in hexadecimal ('x').

Since the letters indicating unit sizes are all distinct from the letters specifying output formats, you do not have to remember whether unit size or format comes first; either order works. The output specifications '4xw' and '4wx' mean exactly the same thing. (However, the count n must come first; 'wx4' does not work.)

Even though the unit size u is ignored for the formats 's' and 'i', you might still want to use a count n; for example, '3i' specifies that you want to see three machine instructions, including any operands. For convenience, especially when used with the display command, the 'i' format also prints branch delay slot instructions, if any, beyond the count specified, which immediately follow the last instruction that is within the count. The command disassemble gives an alternative way of inspecting machine instructions; see Section 9.6 [Source and Machine Code], page 89.

All the defaults for the arguments to x are designed to make it easy to continue scanning memory with minimal specifications each time you use x. For example, after you have inspected three machine instructions with 'x/3i addr', you can inspect the next seven with just 'x/7'. If you use RET to repeat the x command, the repeat count n is used again; the other arguments default as for successive uses of x.

When examining machine instructions, the instruction at current program counter is shown with a => marker. For example:

The addresses and contents printed by the x command are not saved in the value history because there is often too much of them and they would get in the way. Instead, GDB makes these values available for subsequent use in expressions as values of the convenience variables $_{-}$ and $_{-}$. After an x command, the last address examined is available for use in expressions in the convenience variable $_{-}$. The contents of that address, as examined, are available in the convenience variable $_{-}$.

If the x command has a repeat count, the address and contents saved are from the last memory unit printed; this is not the same as the last address printed if several units were printed on the last line of output.

When you are debugging a program running on a remote target machine (see Chapter 20 [Remote Debugging], page 203), you may wish to verify the program's image in the remote machine's memory against the executable file you downloaded to the target. The comparesections command is provided for such situations.

compare-sections [section-name]

Compare the data of a loadable section section-name in the executable file of the program being debugged with the same section in the remote machine's memory, and report any mismatches. With no arguments, compares all loadable

sections. This command's availability depends on the target's support for the "qCRC" remote request.

10.7 Automatic Display

If you find that you want to print the value of an expression frequently (to see how it changes), you might want to add it to the *automatic display list* so that GDB prints its value each time your program stops. Each expression added to the list is given a number to identify it; to remove an expression from the list, you specify that number. The automatic display looks like this:

```
2: foo = 38
3: bar[5] = (struct hack *) 0x3804
```

This display shows item numbers, expressions and their current values. As with displays you request manually using x or print, you can specify the output format you prefer; in fact, display decides whether to use print or x depending your format specification—it uses x if you specify either the 'i' or 's' format, or a unit size; otherwise it uses print.

display expr

Add the expression *expr* to the list of expressions to display each time your program stops. See Section 10.1 [Expressions], page 93.

display does not repeat if you press RET again after using it.

display/fmt expr

For *fmt* specifying only a display format and not a size or count, add the expression *expr* to the auto-display list but arrange to display it each time in the specified format *fmt*. See Section 10.5 [Output Formats], page 97.

display/fmt addr

For fmt 'i' or 's', or including a unit-size or a number of units, add the expression addr as a memory address to be examined each time your program stops. Examining means in effect doing 'x/fmt addr'. See Section 10.6 [Examining Memory], page 99.

For example, 'display/i \$pc' can be helpful, to see the machine instruction about to be executed each time execution stops ('\$pc' is a common name for the program counter; see Section 10.11 [Registers], page 110).

```
undisplay dnums...
delete display dnums...
```

Remove item numbers dnums from the list of expressions to display.

undisplay does not repeat if you press RET after using it. (Otherwise you would just get the error 'No display number ...'.)

disable display dnums...

Disable the display of item numbers *dnums*. A disabled display item is not printed automatically, but is not forgotten. It may be enabled again later.

enable display dnums...

Enable display of item numbers *dnums*. It becomes effective once again in auto display of its expression, until you specify otherwise.

display Display the current values of the expressions on the list, just as is done when your program stops.

info display

Print the list of expressions previously set up to display automatically, each one with its item number, but without showing the values. This includes disabled expressions, which are marked as such. It also includes expressions which would not be displayed right now because they refer to automatic variables not currently available.

If a display expression refers to local variables, then it does not make sense outside the lexical context for which it was set up. Such an expression is disabled when execution enters a context where one of its variables is not defined. For example, if you give the command display last_char while inside a function with an argument last_char, GDB displays this argument while your program continues to stop inside that function. When it stops elsewhere—where there is no variable last_char—the display is disabled automatically. The next time your program stops where last_char is meaningful, you can enable the display expression once again.

10.8 Print Settings

GDB provides the following ways to control how arrays, structures, and symbols are printed. These settings are useful for debugging programs in any language:

```
set print address set print address on
```

GDB prints memory addresses showing the location of stack traces, structure values, pointer values, breakpoints, and so forth, even when it also displays the contents of those addresses. The default is on. For example, this is what a stack frame display looks like with set print address on:

```
(gdb) f
#0 set_quotes (lq=0x34c78 "<<", rq=0x34c88 ">>")
    at input.c:530
530    if (lquote != def_lquote)
```

set print address off

Do not print addresses when displaying their contents. For example, this is the same stack frame displayed with set print address off:

You can use 'set print address off' to eliminate all machine dependent displays from the GDB interface. For example, with print address off, you should get the same text for backtraces on all machines—whether or not they involve pointer arguments.

show print address

Show whether or not addresses are to be printed.

When GDB prints a symbolic address, it normally prints the closest earlier symbol plus an offset. If that symbol does not uniquely identify the address (for example, it is a name whose scope is a single source file), you may need to clarify. One way to do this is with info line, for example 'info line *0x4537'. Alternately, you can set GDB to print the source file and line number when it prints a symbolic address:

set print symbol-filename on

Tell GDB to print the source file name and line number of a symbol in the symbolic form of an address.

set print symbol-filename off

Do not print source file name and line number of a symbol. This is the default.

show print symbol-filename

Show whether or not GDB will print the source file name and line number of a symbol in the symbolic form of an address.

Another situation where it is helpful to show symbol filenames and line numbers is when disassembling code; GDB shows you the line number and source file that corresponds to each instruction.

Also, you may wish to see the symbolic form only if the address being printed is reasonably close to the closest earlier symbol:

set print max-symbolic-offset max-offset

Tell GDB to only display the symbolic form of an address if the offset between the closest earlier symbol and the address is less than *max-offset*. The default is 0, which tells GDB to always print the symbolic form of an address if any symbol precedes it.

show print max-symbolic-offset

Ask how large the maximum offset is that GDB prints in a symbolic address.

If you have a pointer and you are not sure where it points, try 'set print symbol-filename on'. Then you can determine the name and source file location of the variable where it points, using 'p/a pointer'. This interprets the address in symbolic form. For example, here GDB shows that a variable ptt points at another variable t, defined in 'hi2.c':

```
(gdb) set print symbol-filename on
(gdb) p/a ptt
$4 = 0xe008 <t in hi2.c>
```

Warning: For pointers that point to a local variable, 'p/a' does not show the symbol name and filename of the referent, even with the appropriate set print options turned on.

Other settings control how different kinds of objects are printed:

```
set print array on
```

Pretty print arrays. This format is more convenient to read, but uses more space. The default is off.

set print array off

Return to compressed format for arrays.

show print array

Show whether compressed or pretty format is selected for displaying arrays.

```
set print array-indexes
set print array-indexes on
```

Print the index of each element when displaying arrays. May be more convenient to locate a given element in the array or quickly find the index of a given element in that printed array. The default is off.

set print array-indexes off

Stop printing element indexes when displaying arrays.

show print array-indexes

Show whether the index of each element is printed when displaying arrays.

set print elements number-of-elements

Set a limit on how many elements of an array GDB will print. If GDB is printing a large array, it stops printing after it has printed the number of elements set by the set print elements command. This limit also applies to the display of strings. When GDB starts, this limit is set to 200. Setting number-of-elements to zero means that the printing is unlimited.

show print elements

Display the number of elements of a large array that GDB will print. If the number is 0, then the printing is unlimited.

set print frame-arguments value

This command allows to control how the values of arguments are printed when the debugger prints a frame (see Section 8.1 [Frames], page 77). The possible values are:

all The values of all arguments are printed.

scalars Print the value of an argument only if it is a scalar. The value of more complex arguments such as arrays, structures, unions, etc, is replaced by This is the default. Here is an example where

only scalar arguments are shown:

```
#1 0x08048361 in call_me (i=3, s=..., ss=0xbf8d508c, u=..., e=green)
at frame-args.c:23
```

None of the argument values are printed. Instead, the value of each argument is replaced by In this case, the example above now becomes:

```
#1 0x08048361 in call_me (i=..., s=..., ss=..., u=..., e=...) at frame-args.c:23
```

By default, only scalar arguments are printed. This command can be used to configure the debugger to print the value of all arguments, regardless of their type. However, it is often advantageous to not print the value of more complex parameters. For instance, it reduces the amount of information printed in each frame, making the backtrace more readable. Also, it improves performance when displaying Ada frames, because the computation of large arguments can sometimes be CPU-intensive, especially in large applications. Setting print frame-arguments to scalars (the default) or none avoids this computation, thus speeding up the display of each Ada frame.

show print frame-arguments

Show how the value of arguments should be displayed when printing a frame.

set print repeats

Set the threshold for suppressing display of repeated array elements. When the number of consecutive identical elements of an array exceeds the threshold, GDB prints the string "<repeats n times>", where n is the number of identical repetitions, instead of displaying the identical elements themselves. Setting the threshold to zero will cause all elements to be individually printed. The default threshold is 10.

show print repeats

Display the current threshold for printing repeated identical elements.

set print null-stop

Cause GDB to stop printing the characters of an array when the first NULL is encountered. This is useful when large arrays actually contain only short strings. The default is off.

show print null-stop

Show whether gdb stops printing an array on the first null character.

set print pretty on

Cause GDB to print structures in an indented format with one member per line, like this:

```
$1 = {
  next = 0x0,
  flags = {
    sweet = 1,
    sour = 1
  },
  meat = 0x54 "Pork"
}
```

set print pretty off

Cause gdb to print structures in a compact format, like this:

```
1 = {next = 0x0, flags = {sweet = 1, sour = 1}, \\ meat = 0x54 "Pork"}
```

This is the default format.

show print pretty

Show which format GDB is using to print structures.

set print sevenbit-strings on

Print using only seven-bit characters; if this option is set, GDB displays any eight-bit characters (in strings or character values) using the notation \nnn. This setting is best if you are working in English (ASCII) and you use the high-order bit of characters as a marker or "meta" bit.

set print sevenbit-strings off

Print full eight-bit characters. This allows the use of more international character sets, and is the default.

show print sevenbit-strings

Show whether or not gdb is printing only seven-bit characters.

set print union on

Tell gdb to print unions which are contained in structures and other unions. This is the default setting.

set print union off

Tell GDB not to print unions which are contained in structures and other unions. GDB will print "{...}" instead.

show print union

Ask GDB whether or not it will print unions which are contained in structures and other unions.

For example, given the declarations

```
typedef enum {Tree, Bug} Species;
      typedef enum {Big_tree, Acorn, Seedling} Tree_forms;
      typedef enum {Caterpillar, Cocoon, Butterfly}
                    Bug_forms;
      struct thing {
        Species it;
        union {
          Tree_forms tree;
          Bug_forms bug;
       } form;
      };
      struct thing foo = {Tree, {Acorn}};
with set print union on in effect 'p foo' would print
      $1 = {it = Tree, form = {tree = Acorn, bug = Cocoon}}
and with set print union off in effect it would print
      1 = \{ it = Tree, form = \{...\} \}
```

set print union affects programs written in C-like languages and in Pascal.

These settings are of interest when debugging C++ programs:

```
set print demangle set print demangle on
```

Print C++ names in their source form rather than in the encoded ("mangled") form passed to the assembler and linker for type-safe linkage. The default is on.

show print demangle

Show whether C++ names are printed in mangled or demangled form.

```
set print asm-demangle set print asm-demangle on
```

Print C++ names in their source form rather than their mangled form, even in assembler code printouts such as instruction disassemblies. The default is off.

show print asm-demangle

Show whether C++ names in assembly listings are printed in mangled or demangled form.

set demangle-style style

Choose among several encoding schemes used by different compilers to represent C++ names. The choices for *style* are currently:

auto Allow GDB to choose a decoding style by inspecting your program.

gnu Decode based on the GNU C++ compiler (g++) encoding algorithm.

This is the default.

hp Decode based on the HP ANSI C++ (acc) encoding algorithm.

lucid Decode based on the Lucid C++ compiler (lcc) encoding algorithm.

Decode using the algorithm in the C++ Annotated Reference Manual. Warning: this setting alone is not sufficient to allow debugging cfront-generated executables. GDB would require further enhance-

ment to permit that.

If you omit style, you will see a list of possible formats.

show demangle-style

Display the encoding style currently in use for decoding C++ symbols.

set print object set print object on

When displaying a pointer to an object, identify the *actual* (derived) type of the object rather than the *declared* type, using the virtual function table.

set print object off

Display only the declared type of objects, without reference to the virtual function table. This is the default setting.

show print object

Show whether actual, or declared, object types are displayed.

set print static-members

set print static-members on

Print static members when displaying a C++ object. The default is on.

set print static-members off

Do not print static members when displaying a C++ object.

show print static-members

Show whether C++ static members are printed or not.

set print pascal_static-members

set print pascal_static-members on

Print static members when displaying a Pascal object. The default is on.

set print pascal_static-members off

Do not print static members when displaying a Pascal object.

show print pascal_static-members

Show whether Pascal static members are printed or not.

```
set print vtbl set print vtbl on
```

Pretty print C++ virtual function tables. The default is off. (The vtbl commands do not work on programs compiled with the HP ANSI C++ compiler (aCC).)

set print vtbl off

Do not pretty print C++ virtual function tables.

show print vtbl

Show whether C++ virtual function tables are pretty printed, or not.

10.9 Value History

Values printed by the print command are saved in the GDB value history. This allows you to refer to them in other expressions. Values are kept until the symbol table is re-read or discarded (for example with the file or symbol-file commands). When the symbol table changes, the value history is discarded, since the values may contain pointers back to the types defined in the symbol table.

The values printed are given history numbers by which you can refer to them. These are successive integers starting with one. print shows you the history number assigned to a value by printing '\$num = ' before the value; here num is the history number.

To refer to any previous value, use '\$' followed by the value's history number. The way print labels its output is designed to remind you of this. Just \$ refers to the most recent value in the history, and \$\$ refers to the value before that. \$\$n\$ refers to the nth value from the end; \$\$2 is the value just prior to \$\$, \$\$1 is equivalent to \$\$, and \$\$0 is equivalent to \$\$.

For example, suppose you have just printed a pointer to a structure and want to see the contents of the structure. It suffices to type

```
p *$
```

If you have a chain of structures where the component next points to the next one, you can print the contents of the next one with this:

```
p *$.next
```

You can print successive links in the chain by repeating this command—which you can do by just typing RET.

Note that the history records values, not expressions. If the value of ${\tt x}$ is 4 and you type these commands:

```
print x
set x=5
```

then the value recorded in the value history by the print command remains 4 even though the value of x has changed.

```
show values
```

Print the last ten values in the value history, with their item numbers. This is like 'p \$\$9' repeated ten times, except that show values does not change the history.

show values n

Print ten history values centered on history item number n.

show values +

Print ten history values just after the values last printed. If no more values are available, show values + produces no display.

Pressing RET to repeat show values n has exactly the same effect as 'show values +'.

10.10 Convenience Variables

GDB provides convenience variables that you can use within GDB to hold on to a value and refer to it later. These variables exist entirely within GDB; they are not part of your program, and setting a convenience variable has no direct effect on further execution of your program. That is why you can use them freely.

Convenience variables are prefixed with '\$'. Any name preceded by '\$' can be used for a convenience variable, unless it is one of the predefined machine-specific register names (see Section 10.11 [Registers], page 110). (Value history references, in contrast, are *numbers* preceded by '\$'. See Section 10.9 [Value History], page 108.)

You can save a value in a convenience variable with an assignment expression, just as you would set a variable in your program. For example:

```
set $foo = *object_ptr
```

would save in \$foo the value contained in the object pointed to by object_ptr.

Using a convenience variable for the first time creates it, but its value is void until you assign a new value. You can alter the value with another assignment at any time.

Convenience variables have no fixed types. You can assign a convenience variable any type of value, including structures and arrays, even if that variable already has a value of a different type. The convenience variable, when used as an expression, has the type of its current value.

show convenience

Print a list of convenience variables used so far, and their values. Abbreviated show conv.

init-if-undefined \$variable = expression

Set a convenience variable if it has not already been set. This is useful for user-defined commands that keep some state. It is similar, in concept, to using local static variables with initializers in C (except that convenience variables are global). It can also be used to allow users to override default values used in a command script.

If the variable is already defined then the expression is not evaluated so any side-effects do not occur.

One of the ways to use a convenience variable is as a counter to be incremented or a pointer to be advanced. For example, to print a field from successive elements of an array of structures:

```
set $i = 0
print bar[$i++]->contents
```

Repeat that command by typing RET.

Some convenience variables are created automatically by GDB and given values likely to be useful.

\$_ The variable \$_ is automatically set by the x command to the last address examined (see Section 10.6 [Examining Memory], page 99). Other commands which provide a default address for x to examine also set \$_ to that address; these commands include info line and info breakpoint. The type of \$_ is void * except when set by the x command, in which case it is a pointer to the type of \$__.

\$__ The variable \$__ is automatically set by the x command to the value found in the last address examined. Its type is chosen to match the format in which the data was printed.

\$_exitcode

The variable **\$_exitcode** is automatically set to the exit code when the program being debugged terminates.

\$_siginfo

The variable \$_siginfo contains extra signal information (see [extra signal information], page 65). Note that \$_siginfo could be empty, if the application has not yet received any signals. For example, it will be empty before you execute the run command.

On HP-UX systems, if you refer to a function or variable name that begins with a dollar sign, GDB searches for a user or system name first, before it searches for a convenience variable.

GDB also supplies some *convenience functions*. These have a syntax similar to convenience variables. A convenience function can be used in an expression just like an ordinary function; however, a convenience function is implemented internally to GDB.

help function

Print a list of all convenience functions.

10.11 Registers

You can refer to machine register contents, in expressions, as variables with names starting with '\$'. The names of registers are different for each machine; use info registers to see the names used on your machine.

info registers

Print the names and values of all registers except floating-point and vector registers (in the selected stack frame).

info all-registers

Print the names and values of all registers, including floating-point and vector registers (in the selected stack frame).

info registers regname ...

Print the *relativized* value of each specified register *regname*. As discussed in detail below, register values are normally relative to the selected stack frame. *regname* may be any register name valid on the machine you are using, with or without the initial '\$'.

GDB has four "standard" register names that are available (in expressions) on most machines—whenever they do not conflict with an architecture's canonical mnemonics for registers. The register names \$pc and \$sp are used for the program counter register and the stack pointer. \$fp is used for a register that contains a pointer to the current stack frame, and \$ps is used for a register that contains the processor status. For example, you could print the program counter in hex with

Whenever possible, these four standard register names are available on your machine even though the machine has different canonical mnemonics, so long as there is no conflict. The info registers command shows the canonical names. For example, on the SPARC, info registers displays the processor status register as \$psr but you can also refer to it as \$ps; and on x86-based machines \$ps is an alias for the EFLAGS register.

GDB always considers the contents of an ordinary register as an integer when the register is examined in this way. Some machines have special registers which can hold nothing but floating point; these registers are considered to have floating point values. There is no way to refer to the contents of an ordinary register as floating point value (although you can print it as a floating point value with 'print/f \$regname').

Some registers have distinct "raw" and "virtual" data formats. This means that the data format in which the register contents are saved by the operating system is not the same one that your program normally sees. For example, the registers of the 68881 floating point coprocessor are always saved in "extended" (raw) format, but all C programs expect to work with "double" (virtual) format. In such cases, GDB normally works with the virtual format only (the format that makes sense for your program), but the info registers command prints the data in both formats.

Some machines have special registers whose contents can be interpreted in several different ways. For example, modern x86-based machines have SSE and MMX registers that can hold several values packed together in several different formats. GDB refers to such registers in struct notation:

```
(gdb) print $xmm1
$1 = {
    v4_float = {0, 3.43859137e-038, 1.54142831e-044, 1.821688e-044},
    v2_double = {9.92129282474342e-303, 2.7585945287983262e-313},
    v16_int8 = "\000\000\000\000\3706;\001\v\000\000\r\000\r\000\000",
    v8_int16 = {0, 0, 14072, 315, 11, 0, 13, 0},
    v4_int32 = {0, 20657912, 11, 13},
    v2_int64 = {88725056443645952, 55834574859},
    uint128 = 0x000000000000000013b36f800000000
}
```

² This is a way of removing one word from the stack, on machines where stacks grow downward in memory (most machines, nowadays). This assumes that the innermost stack frame is selected; setting \$sp is not allowed when other stack frames are selected. To pop entire frames off the stack, regardless of machine architecture, use return; see Section 17.4 [Returning from a Function], page 183.

To set values of such registers, you need to tell GDB which view of the register you wish to change, as if you were assigning value to a struct member:

(gdb) set \$xmm1.uint128 = 0x0000000000000000000000FFFFFFFF

Normally, register values are relative to the selected stack frame (see Section 8.3 [Selecting a Frame], page 80). This means that you get the value that the register would contain if all stack frames farther in were exited and their saved registers restored. In order to see the true contents of hardware registers, you must select the innermost frame (with 'frame 0').

However, GDB must deduce where registers are saved, from the machine code generated by your compiler. If some registers are not saved, or if GDB is unable to locate the saved registers, the selected stack frame makes no difference.

10.12 Floating Point Hardware

Depending on the configuration, GDB may be able to give you more information about the status of the floating point hardware.

info float

Display hardware-dependent information about the floating point unit. The exact contents and layout vary depending on the floating point chip. Currently, 'info float' is supported on the ARM and x86 machines.

10.13 Vector Unit

Depending on the configuration, GDB may be able to give you more information about the status of the vector unit.

info vector

Display information about the vector unit. The exact contents and layout vary depending on the hardware.

10.14 Operating System Auxiliary Information

GDB provides interfaces to useful OS facilities that can help you debug your program.

When GDB runs on a *Posix system* (such as GNU or Unix machines), it interfaces with the inferior via the ptrace system call. The operating system creates a special sata structure, called struct user, for this interface. You can use the command info udot to display the contents of this data structure.

info udot Display the contents of the struct user maintained by the OS kernel for the program being debugged. GDB displays the contents of struct user as a list of hex numbers, similar to the examine command.

Some operating systems supply an *auxiliary vector* to programs at startup. This is akin to the arguments and environment that you specify for a program, but contains a system-dependent variety of binary values that tell system libraries important details about the hardware, operating system, and process. Each value's purpose is identified by an integer tag; the meanings are well-known but system-specific. Depending on the configuration and operating system facilities, GDB may be able to show you this information. For remote targets, this functionality may further depend on the remote stub's support of the 'qXfer:auxv:read' packet, see [qXfer auxiliary vector read], page 425.

info auxv Display the auxiliary vector of the inferior, which can be either a live process or a core dump file. GDB prints each tag value numerically, and also shows names and text descriptions for recognized tags. Some values in the vector are numbers, some bit masks, and some pointers to strings or other data. GDB displays each value in the most appropriate form for a recognized tag, and in hexadecimal for an unrecognized tag.

On some targets, GDB can access operating-system-specific information and display it to user, without interpretation. For remote targets, this functionality depends on the remote stub's support of the 'qXfer:osdata:read' packet, see [qXfer osdata read], page 426.

info os processes

Display the list of processes on the target. For each process, GDB prints the process identifier, the name of the user, and the command corresponding to the process.

10.15 Memory Region Attributes

Memory region attributes allow you to describe special handling required by regions of your target's memory. GDB uses attributes to determine whether to allow certain types of memory accesses; whether to use specific width accesses; and whether to cache target memory. By default the description of memory regions is fetched from the target (if the current target supports this), but the user can override the fetched regions.

Defined memory regions can be individually enabled and disabled. When a memory region is disabled, GDB uses the default attributes when accessing memory in that region. Similarly, if no memory regions have been defined, GDB uses the default attributes when accessing all memory.

When a memory region is defined, it is given a number to identify it; to enable, disable, or remove a memory region, you specify that number.

mem lower upper attributes...

Define a memory region bounded by *lower* and *upper* with attributes attributes..., and add it to the list of regions monitored by GDB. Note that upper == 0 is a special case: it is treated as the target's maximum memory address. (0xffff on 16 bit targets, 0xffffffff on 32 bit targets, etc.)

mem auto Discard any user changes to the memory regions and use target-supplied regions, if available, or no regions if the target does not support.

delete mem nums...

Remove memory regions nums... from the list of regions monitored by GDB.

disable mem nums...

Disable monitoring of memory regions *nums*.... A disabled memory region is not forgotten. It may be enabled again later.

enable mem nums...

Enable monitoring of memory regions nums....

info mem Print a table of all defined memory regions, with the following columns for each region:

Debugging with GDB

Memory Region Number

Enabled or Disabled.

Enabled memory regions are marked with 'y'. Disabled memory regions are marked with 'n'.

Lo Address

The address defining the inclusive lower bound of the memory region.

Hi Address

The address defining the exclusive upper bound of the memory region.

Attributes The list of attributes set for this memory region.

10.15.1 Attributes

10.15.1.1 Memory Access Mode

The access mode attributes set whether GDB may make read or write accesses to a memory region.

While these attributes prevent GDB from performing invalid memory accesses, they do nothing to prevent the target system, I/O DMA, etc. from accessing memory.

ro Memory is read only.

wo Memory is write only.

rw Memory is read/write. This is the default.

10.15.1.2 Memory Access Size

The access size attribute tells GDB to use specific sized accesses in the memory region. Often memory mapped device registers require specific sized accesses. If no access size attribute is specified, GDB may use accesses of any size.

8 Use 8 bit memory accesses.

16 Use 16 bit memory accesses.

32 Use 32 bit memory accesses.

64 Use 64 bit memory accesses.

10.15.1.3 Data Cache

The data cache attributes set whether GDB will cache target memory. While this generally improves performance by reducing debug protocol overhead, it can lead to incorrect results because GDB does not know about volatile variables or memory mapped device registers.

cache Enable GDB to cache target memory.

nocache Disable GDB from caching target memory. This is the default.

10.15.2 Memory Access Checking

GDB can be instructed to refuse accesses to memory that is not explicitly described. This can be useful if accessing such regions has undesired effects for a specific target, or to provide better error checking. The following commands control this behaviour.

set mem inaccessible-by-default [on|off]

If on is specified, make GDB treat memory not explicitly described by the memory ranges as non-existent and refuse accesses to such memory. The checks are only performed if there's at least one memory range defined. If off is specified, make GDB treat the memory not explicitly described by the memory ranges as RAM. The default value is on.

show mem inaccessible-by-default

Show the current handling of accesses to unknown memory.

10.16 Copy Between Memory and a File

You can use the commands dump, append, and restore to copy data between target memory and a file. The dump and append commands write data to a file, and the restore command reads data from a file back into the inferior's memory. Files may be in binary, Motorola S-record, Intel hex, or Tektronix Hex format; however, GDB can only append to binary files.

dump [format] memory filename start_addr end_addr dump [format] value filename expr

Dump the contents of memory from $start_addr$ to end_addr , or the value of expr, to filename in the given format.

The format parameter may be any one of:

binary Raw binary form.

ihex Intel hex format.

srec Motorola S-record format.

tekhex Tektronix Hex format.

GDB uses the same definitions of these formats as the GNU binary utilities, like 'objdump' and 'objcopy'. If format is omitted, GDB dumps the data in raw binary form.

append [binary] memory filename start_addr end_addr append [binary] value filename expr

Append the contents of memory from $start_addr$ to end_addr , or the value of expr, to the file filename, in raw binary form. (GDB can only append data to files in raw binary form.)

restore filename [binary] bias start end

Restore the contents of file *filename* into memory. The restore command can automatically recognize any known BFD file format, except for raw binary. To restore a raw binary file you must specify the optional keyword binary after the filename.

If bias is non-zero, its value will be added to the addresses contained in the file. Binary files always start at address zero, so they will be restored at address

bias. Other bfd files have a built-in location; they will be restored at offset bias from that location.

If start and/or end are non-zero, then only data between file offset start and file offset end will be restored. These offsets are relative to the addresses in the file, before the bias argument is applied.

10.17 How to Produce a Core File from Your Program

A core file or core dump is a file that records the memory image of a running process and its process status (register values etc.). Its primary use is post-mortem debugging of a program that crashed while it ran outside a debugger. A program that crashes automatically produces a core file, unless this feature is disabled by the user. See Section 18.1 [Files], page 187, for information on invoking GDB in the post-mortem debugging mode.

Occasionally, you may wish to produce a core file of the program you are debugging in order to preserve a snapshot of its state. GDB has a special command for that.

generate-core-file [file] gcore [file]

Produce a core dump of the inferior process. The optional argument file specifies the file name where to put the core dump. If not specified, the file name defaults to 'core.pid', where pid is the inferior process ID.

Note that this command is implemented only for some systems (as of this writing, GNU/Linux, FreeBSD, Solaris, Unixware, and S390).

10.18 Character Sets

If the program you are debugging uses a different character set to represent characters and strings than the one GDB uses itself, GDB can automatically translate between the character sets for you. The character set GDB uses we call the *host character set*; the one the inferior program uses we call the *target character set*.

For example, if you are running GDB on a GNU/Linux system, which uses the ISO Latin 1 character set, but you are using GDB's remote protocol (see Chapter 20 [Remote Debugging], page 203) to debug a program running on an IBM mainframe, which uses the EBCDIC character set, then the host character set is Latin-1, and the target character set is EBCDIC. If you give GDB the command set target-charset EBCDIC-US, then GDB translates between EBCDIC and Latin 1 as you print character or string values, or use character and string literals in expressions.

GDB has no way to automatically recognize which character set the inferior program uses; you must tell it, using the set target-charset command, described below.

Here are the commands for controlling GDB's character set support:

set target-charset charset

Set the current target character set to *charset*. To display the list of supported target character sets, type *set target-charset TABTAB*.

set host-charset charset

Set the current host character set to charset.

By default, GDB uses a host character set appropriate to the system it is running on; you can override that default using the set host-charset command. On some systems, GDB cannot automatically determine the appropriate host character set. In this case, GDB uses 'UTF-8'.

GDB can only use certain character sets as its host character set. If you type set target-charset TABTAB, GDB will list the host character sets it supports.

set charset charset

Set the current host and target character sets to *charset*. As above, if you type **set charset TABTAB**, GDB will list the names of the character sets that can be used for both host and target.

show charset

Show the names of the current host and target character sets.

show host-charset

Show the name of the current host character set.

show target-charset

Show the name of the current target character set.

set target-wide-charset charset

Set the current target's wide character set to *charset*. This is the character set used by the target's wchar_t type. To display the list of supported wide character sets, type set target-wide-charset TABTAB.

show target-wide-charset

Show the name of the current target's wide character set.

Here is an example of GDB's character set support in action. Assume that the following source code has been placed in the file 'charset-test.c':

#include <stdio.h>

In this program, ascii_hello and ibm1047_hello are arrays containing the string 'Hello, world!' followed by a newline, encoded in the ASCII and IBM1047 character sets.

We compile the program, and invoke the debugger on it:

```
$ gcc -g charset-test.c -o charset-test
$ gdb -nw charset-test
GNU gdb 2001-12-19-cvs
Copyright 2001 Free Software Foundation, Inc.
...
(gdb)
```

We can use the **show charset** command to see what character sets GDB is currently using to interpret and display characters and strings:

```
(gdb) show charset
The current host and target character set is 'ISO-8859-1'.
(gdb)
```

For the sake of printing this manual, let's use ASCII as our initial character set:

```
(gdb) set charset ASCII
(gdb) show charset
The current host and target character set is 'ASCII'.
(gdb)
```

Let's assume that ASCII is indeed the correct character set for our host system — in other words, let's assume that if GDB prints characters using the ASCII character set, our terminal will display them properly. Since our current target character set is also ASCII, the contents of ascii_hello print legibly:

```
(gdb) print ascii_hello
$1 = 0x401698 "Hello, world!\n"
(gdb) print ascii_hello[0]
$2 = 72 'H'
(gdb)
```

GDB uses the target character set for character and string literals you use in expressions:

```
(gdb) print '+'
$3 = 43 '+'
(gdb)
```

The ASCII character set uses the number 43 to encode the '+' character.

GDB relies on the user to tell it which character set the target program uses. If we print ibm1047_hello while our target character set is still ASCII, we get jibberish:

```
(gdb) print ibm1047_hello
$4 = 0x4016a8 "\310\205\223\223\226k@\246\226\231\223\204Z%"
(gdb) print ibm1047_hello[0]
$5 = 200 '\310'
(gdb)
```

If we invoke the set target-charset followed by TABTAB, GDB tells us the character sets it supports:

```
(gdb) set target-charset
ASCII EBCDIC-US IBM1047 ISO-8859-1
(gdb) set target-charset
```

We can select IBM1047 as our target character set, and examine the program's strings again. Now the ASCII string is wrong, but GDB translates the contents of ibm1047_hello from the target character set, IBM1047, to the host character set, ASCII, and they display correctly:

```
(gdb) set target-charset IBM1047
(gdb) show charset
The current host character set is 'ASCII'.
The current target character set is 'IBM1047'.
(gdb) print ascii_hello
$6 = 0x401698 "\110\145%%?\054\040\167?\162%\144\041\012"
(gdb) print ascii_hello[0]
$7 = 72 '\110'
(gdb) print ibm1047_hello
$8 = 0x4016a8 "Hello, world!\n"
(gdb) print ibm1047_hello[0]
```

```
$9 = 200 'H'
(gdb)
```

As above, GDB uses the target character set for character and string literals you use in expressions:

```
(gdb) print '+'
$10 = 78 '+'
(gdb)
```

The IBM1047 character set uses the number 78 to encode the '+' character.

10.19 Caching Data of Remote Targets

GDB caches data exchanged between the debugger and a remote target (see Chapter 20 [Remote Debugging], page 203). Such caching generally improves performance, because it reduces the overhead of the remote protocol by bundling memory reads and writes into large chunks. Unfortunately, simply caching everything would lead to incorrect results, since GDB does not necessarily know anything about volatile values, memory-mapped I/O addresses, etc. Furthermore, in non-stop mode (see Section 5.4.2 [Non-Stop Mode], page 67) memory can be changed while a gdb command is executing. Therefore, by default, GDB only caches data known to be on the stack³. Other regions of memory can be explicitly marked as cacheable; see see Section 10.15 [Memory Region Attributes], page 113.

```
set remotecache on set remotecache off
```

This option no longer does anything; it exists for compatibility with old scripts.

show remotecache

Show the current state of the obsolete remotecache flag.

```
set stack-cache on set stack-cache off
```

Enable or disable caching of stack accesses. When $\mathtt{ON},$ use caching. By default, this option is $\mathtt{ON}.$

show stack-cache

Show the current state of data caching for memory accesses.

info dcache [line]

Print the information about the data cache performance. The information displayed includes the dcache width and depth, and for each cache line, its number, address, and how many times it was referenced. This command is useful for debugging the data cache operation.

If a line number is specified, the contents of that line will be printed in hex.

10.20 Search Memory

Memory can be searched for a particular sequence of bytes with the find command.

³ In non-stop mode, it is moderately rare for a running thread to modify the stack of a stopped thread in a way that would interfere with a backtrace, and caching of stack reads provides a significant speed up of remote backtraces.

```
find [/sn] start_addr, +len, val1 [, val2, ...]
find [/sn] start_addr, end_addr, val1 [, val2, ...]
```

Search memory for the sequence of bytes specified by val1, val2, etc. The search begins at address start_addr and continues for either len bytes or through to end_addr inclusive.

s and n are optional parameters. They may be specified in either order, apart or together.

s, search query size

The size of each search query value.

```
b bytes
h halfwords (two bytes)
w words (four bytes)
g giant words (eight bytes)
```

All values are interpreted in the current language. This means, for example, that if the current source language is C/C++ then searching for the string "hello" includes the trailing '\0'.

If the value size is not specified, it is taken from the value's type in the current language. This is useful when one wants to specify the search pattern as a mixture of types. Note that this means, for example, that in the case of C-like languages a search for an untyped 0x42 will search for '(int) 0x42' which is typically four bytes.

n, maximum number of finds

The maximum number of matches to print. The default is to print all finds.

You can use strings as search values. Quote them with double-quotes ("). The string value is copied into the search pattern byte by byte, regardless of the endianness of the target and the size specification.

The address of each match found is printed as well as a count of the number of matches found.

The address of the last value found is stored in convenience variable '\$_'. A count of the number of matches is stored in '\$numfound'.

For example, if stopped at the printf in this function:

```
void
hello ()
{
    static char hello[] = "hello-hello";
    static struct { char c; short s; int i; }
        __attribute__ ((packed)) mixed
        = { 'c', 0x1234, 0x87654321 };
    printf ("%s\n", hello);
}
you get during debugging:
    (gdb) find &hello[0], +sizeof(hello), "hello"
    0x804956d <hello.1620+6>
    1 pattern found
    (gdb) find &hello[0], +sizeof(hello), 'h', 'e', 'l', 'l', 'o'
```

```
0x8049567 <hello.1620>
0x804956d <hello.1620+6>
2 patterns found
(gdb) find /b1 &hello[0], +sizeof(hello), 'h', 0x65, 'l'
0x8049567 <hello.1620>
1 pattern found
(gdb) find &mixed, +sizeof(mixed), (char) 'c', (short) 0x1234, (int) 0x87654321
0x8049560 <mixed.1625>
1 pattern found
(gdb) print $numfound
$1 = 1
(gdb) print $_
$2 = (void *) 0x8049560
```

11 Debugging Optimized Code

Almost all compilers support optimization. With optimization disabled, the compiler generates assembly code that corresponds directly to your source code, in a simplistic way. As the compiler applies more powerful optimizations, the generated assembly code diverges from your original source code. With help from debugging information generated by the compiler, GDB can map from the running program back to constructs from your original source.

GDB is more accurate with optimization disabled. If you can recompile without optimization, it is easier to follow the progress of your program during debugging. But, there are many cases where you may need to debug an optimized version.

When you debug a program compiled with '-g -0', remember that the optimizer has rearranged your code; the debugger shows you what is really there. Do not be too surprised when the execution path does not exactly match your source file! An extreme example: if you define a variable, but never use it, GDB never sees that variable—because the compiler optimizes it out of existence.

Some things do not work as well with '-g -0' as with just '-g', particularly on machines with instruction scheduling. If in doubt, recompile with '-g' alone, and if this fixes the problem, please report it to us as a bug (including a test case!). See Section 10.3 [Variables], page 95, for more information about debugging optimized code.

11.1 Inline Functions

Inlining is an optimization that inserts a copy of the function body directly at each call site, instead of jumping to a shared routine. GDB displays inlined functions just like non-inlined functions. They appear in backtraces. You can view their arguments and local variables, step into them with step, skip them with next, and escape from them with finish. You can check whether a function was inlined by using the info frame command.

For GDB to support inlined functions, the compiler must record information about inlining in the debug information — GCC using the DWARF 2 format does this, and several other compilers do also. GDB only supports inlined functions when using DWARF 2. Versions of GCC before 4.1 do not emit two required attributes ('DW_AT_call_file' and 'DW_AT_call_line'); GDB does not display inlined function calls with earlier versions of GCC. It instead displays the arguments and local variables of inlined functions as local variables in the caller.

The body of an inlined function is directly included at its call site; unlike a non-inlined function, there are no instructions devoted to the call. GDB still pretends that the call site and the start of the inlined function are different instructions. Stepping to the call site shows the call site, and then stepping again shows the first line of the inlined function, even though no additional instructions are executed.

This makes source-level debugging much clearer; you can see both the context of the call and then the effect of the call. Only stepping by a single instruction using stepi or nexti does not do this; single instruction steps always show the inlined body.

There are some ways that GDB does not pretend that inlined function calls are the same as normal calls:

• You cannot set breakpoints on inlined functions. GDB either reports that there is no symbol with that name, or else sets the breakpoint only on non-inlined copies of the function. This limitation will be removed in a future version of GDB; until then, set a breakpoint by line number on the first line of the inlined function instead.

- Setting breakpoints at the call site of an inlined function may not work, because the call site does not contain any code. GDB may incorrectly move the breakpoint to the next line of the enclosing function, after the call. This limitation will be removed in a future version of GDB; until then, set a breakpoint on an earlier line or inside the inlined function instead.
- GDB cannot locate the return value of inlined calls after using the finish command. This is a limitation of compiler-generated debugging information; after finish, you can step to the next line and print a variable where your program stored the return value.

12 C Preprocessor Macros

Some languages, such as C and C++, provide a way to define and invoke "preprocessor macros" which expand into strings of tokens. GDB can evaluate expressions containing macro invocations, show the result of macro expansion, and show a macro's definition, including where it was defined.

You may need to compile your program specially to provide GDB with information about preprocessor macros. Most compilers do not include macros in their debugging information, even when you compile with the '-g' flag. See Section 4.1 [Compilation], page 25.

A program may define a macro at one point, remove that definition later, and then provide a different definition after that. Thus, at different points in the program, a macro may have different definitions, or have no definition at all. If there is a current stack frame, GDB uses the macros in scope at that frame's source code line. Otherwise, GDB uses the macros in scope at the current listing location; see Section 9.1 [List], page 83.

Whenever gdb evaluates an expression, it always expands any macro invocations present in the expression. gdb also provides the following commands for working with macros explicitly.

macro expand expression macro exp expression

Show the results of expanding all preprocessor macro invocations in expression. Since GDB simply expands macros, but does not parse the result, expression need not be a valid expression; it can be any string of tokens.

macro expand-once expression macro exp1 expression

(This command is not yet implemented.) Show the results of expanding those preprocessor macro invocations that appear explicitly in expression. Macro invocations appearing in that expansion are left unchanged. This command allows you to see the effect of a particular macro more clearly, without being confused by further expansions. Since GDB simply expands macros, but does not parse the result, expression need not be a valid expression; it can be any string of tokens.

info macro macro

Show the definition of the macro named *macro*, and describe the source location or compiler command-line where that definition was established.

macro define macro replacement-list macro define macro (arglist) replacement-list

Introduce a definition for a preprocessor macro named *macro*, invocations of which are replaced by the tokens given in *replacement-list*. The first form of this command defines an "object-like" macro, which takes no arguments; the second form defines a "function-like" macro, which takes the arguments given in *arglist*.

A definition introduced by this command is in scope in every expression evaluated in GDB, until it is removed with the macro undef command, described below. The definition overrides all definitions for macro present in the program being debugged, as well as any previous user-supplied definition.

macro undef macro

Remove any user-supplied definition for the macro named *macro*. This command only affects definitions provided with the macro define command, described above; it cannot remove definitions present in the program being debugged.

macro list

List all the macros defined using the macro define command.

Here is a transcript showing the above commands in action. First, we show our source files:

```
$ cat sample.c
#include <stdio.h>
#include "sample.h"

#define M 42
#define ADD(x) (M + x)

main ()
{
    #define N 28
        printf ("Hello, world!\n");
#undef N
        printf ("We're so creative.\n");
#define N 1729
        printf ("Goodbye, world!\n");
}
$ cat sample.h
#define Q <
$\frac{4}{2}$</pre>
```

Now, we compile the program using the GNU C compiler, GCC. We pass the '-gdwarf-2' and '-g3' flags to ensure the compiler includes information about preprocessor macros in the debugging information.

```
$ gcc -gdwarf-2 -g3 sample.c -o sample
$
Now, we start GDB on our sample program:
$ gdb -nw sample
GNU gdb 2002-05-06-cvs
Copyright 2002 Free Software Foundation, Inc.
GDB is free software, ...
(gdb)
```

We can expand macros and examine their definitions, even when the program is not running. GDB uses the current listing position to decide which macro definitions are in scope:

```
(gdb) list main
3
4     #define M 42
5     #define ADD(x) (M + x)
6
7     main ()
8     {
9     #define N 28
10     printf ("Hello, world!\n");
11     #undef N
```

```
12     printf ("We're so creative.\n");
(gdb) info macro ADD
Defined at /home/jimb/gdb/macros/play/sample.c:5
#define ADD(x) (M + x)
(gdb) info macro Q
Defined at /home/jimb/gdb/macros/play/sample.h:1
    included at /home/jimb/gdb/macros/play/sample.c:2
#define Q <
(gdb) macro expand ADD(1)
expands to: (42 + 1)
(gdb) macro expand-once ADD(1)
expands to: once (M + 1)
(gdb)</pre>
```

In the example above, note that macro expand-once expands only the macro invocation explicit in the original text — the invocation of ADD — but does not expand the invocation of the macro M, which was introduced by ADD.

Once the program is running, GDB uses the macro definitions in force at the source line of the current stack frame:

```
(gdb) break main
   Breakpoint 1 at 0x8048370: file sample.c, line 10.
   Starting program: /home/jimb/gdb/macros/play/sample
  Breakpoint 1, main () at sample.c:10
   10
            printf ("Hello, world!\n");
   (gdb)
At line 10, the definition of the macro N at line 9 is in force:
   (gdb) info macro N
   Defined at /home/jimb/gdb/macros/play/sample.c:9
   #define N 28
   (gdb) macro expand N Q M
   expands to: 28 < 42
   (gdb) print N Q M
   $1 = 1
   (gdb)
```

As we step over directives that remove N's definition, and then give it a new definition, GDB finds the definition (or lack thereof) in force at each point:

```
(gdb) next
Hello, world!
         printf ("We're so creative.\n");
(gdb) info macro {\tt N}
The symbol 'N' has no definition as a C/C++ preprocessor macro
at /home/jimb/gdb/macros/play/sample.c:12
(gdb) next
We're so creative.
         printf ("Goodbye, world!\n");
(gdb) info macro N
Defined at /home/jimb/gdb/macros/play/sample.c:13
#define N 1729
(gdb) macro expand N Q M
expands to: 1729 < 42
(gdb) print N Q M
$2 = 0
(gdb)
```

In addition to source files, macros can be defined on the compilation command line using the '-Dname=value' syntax. For macros defined in such a way, GDB displays the location of their definition as line zero of the source file submitted to the compiler.

```
(gdb) info macro __STDC__
Defined at /home/jimb/gdb/macros/play/sample.c:0
-D__STDC__=1
(gdb)
```

13 Tracepoints

In some applications, it is not feasible for the debugger to interrupt the program's execution long enough for the developer to learn anything helpful about its behavior. If the program's correctness depends on its real-time behavior, delays introduced by a debugger might cause the program to change its behavior drastically, or perhaps fail, even when the code itself is correct. It is useful to be able to observe the program's behavior without interrupting it.

Using GDB's trace and collect commands, you can specify locations in the program, called *tracepoints*, and arbitrary expressions to evaluate when those tracepoints are reached. Later, using the tfind command, you can examine the values those expressions had when the program hit the tracepoints. The expressions may also denote objects in memory—structures or arrays, for example—whose values GDB should record; while visiting a particular tracepoint, you may inspect those objects as if they were in memory at that moment. However, because GDB records these values without interacting with you, it can do so quickly and unobtrusively, hopefully not disturbing the program's behavior.

The tracepoint facility is currently available only for remote targets. See Chapter 19 [Targets], page 199. In addition, your remote target must know how to collect trace data. This functionality is implemented in the remote stub; however, none of the stubs distributed with gdb support tracepoints as of this writing. The format of the remote packets used to implement tracepoints are described in Section D.7 [Tracepoint Packets], page 429.

It is also possible to get trace data from a file, in a manner reminiscent of corefiles; you specify the filename, and use tfind to search through the file. See Section 13.4 [Trace Files], page 139, for more details.

This chapter describes the tracepoint commands and features.

13.1 Commands to Set Tracepoints

Before running such a trace experiment, an arbitrary number of tracepoints can be set. A tracepoint is actually a special type of breakpoint (see Section 5.1.1 [Set Breaks], page 44), so you can manipulate it using standard breakpoint commands. For instance, as with breakpoints, tracepoint numbers are successive integers starting from one, and many of the commands associated with tracepoints take the tracepoint number as their argument, to identify which tracepoint to work on.

For each tracepoint, you can specify, in advance, some arbitrary set of data that you want the target to collect in the trace buffer when it hits that tracepoint. The collected data can include registers, local variables, or global data. Later, you can use GDB commands to examine the values these data had at the time the tracepoint was hit.

Tracepoints do not support every breakpoint feature. Conditional expressions and ignore counts on tracepoints have no effect, and tracepoints cannot run GDB commands when they are hit. Tracepoints may not be thread-specific either.

Some targets may support fast tracepoints, which are inserted in a different way (such as with a jump instead of a trap), that is faster but possibly restricted in where they may be installed.

This section describes commands to set tracepoints and associated conditions and actions.

13.1.1 Create and Delete Tracepoints

trace location

The trace command is very similar to the break command. Its argument location can be a source line, a function name, or an address in the target program. See Section 9.2 [Specify Location], page 84. The trace command defines a tracepoint, which is a point in the target program where the debugger will briefly stop, collect some data, and then allow the program to continue. Setting a tracepoint or changing its actions doesn't take effect until the next tstart command, and once a trace experiment is running, further changes will not have any effect until the next trace experiment starts.

Here are some examples of using the trace command:

You can abbreviate trace as tr.

trace location if cond

Set a tracepoint with condition *cond*; evaluate the expression *cond* each time the tracepoint is reached, and collect data only if the value is nonzero—that is, if *cond* evaluates as true. See Section 13.1.4 [Tracepoint Conditions], page 131, for more information on tracepoint conditions.

ftrace location [if cond]

The ftrace command sets a fast tracepoint. For targets that support them, fast tracepoints will use a more efficient but possibly less general technique to trigger data collection, such as a jump instruction instead of a trap, or some sort of hardware support. It may not be possible to create a fast tracepoint at the desired location, in which case the command will exit with an explanatory message.

GDB handles arguments to ftrace exactly as for trace.

The convenience variable **\$tpnum** records the tracepoint number of the most recently set tracepoint.

delete tracepoint [num]

Permanently delete one or more tracepoints. With no argument, the default is to delete all tracepoints. Note that the regular delete command can remove tracepoints also.

Examples:

```
(gdb) delete trace 1 2 3 // remove three tracepoints (gdb) delete trace // remove all tracepoints
```

You can abbreviate this command as del tr.

13.1.2 Enable and Disable Tracepoints

These commands are deprecated; they are equivalent to plain disable and enable.

disable tracepoint [num]

Disable tracepoint *num*, or all tracepoints if no argument *num* is given. A disabled tracepoint will have no effect during the next trace experiment, but it is not forgotten. You can re-enable a disabled tracepoint using the enable tracepoint command.

enable tracepoint [num]

Enable tracepoint *num*, or all tracepoints. The enabled tracepoints will become effective the next time a trace experiment is run.

13.1.3 Tracepoint Passcounts

passcount [n [num]]

Set the passcount of a tracepoint. The passcount is a way to automatically stop a trace experiment. If a tracepoint's passcount is n, then the trace experiment will be automatically stopped on the n'th time that tracepoint is hit. If the tracepoint number num is not specified, the passcount command sets the passcount of the most recently defined tracepoint. If no passcount is given, the trace experiment will run until stopped explicitly by the user.

Examples:

```
(gdb) passcount 5 2 // Stop on the 5th execution of // tracepoint 2

(gdb) passcount 12 // Stop on the 12th execution of the // most recently defined tracepoint.

(gdb) trace foo (gdb) pass 3 (gdb) trace bar (gdb) pass 2 (gdb) trace baz (gdb) pass 1 // Stop tracing when foo has been // executed 3 times OR when bar has // been executed 2 times // OR when baz has been executed 1 time.
```

13.1.4 Tracepoint Conditions

The simplest sort of tracepoint collects data every time your program reaches a specified place. You can also specify a *condition* for a tracepoint. A condition is just a Boolean expression in your programming language (see Section 10.1 [Expressions], page 93). A tracepoint with a condition evaluates the expression each time your program reaches it, and data collection happens only if the condition is true.

Tracepoint conditions can be specified when a tracepoint is set, by using 'if' in the arguments to the trace command. See Section 13.1.1 [Setting Tracepoints], page 130. They can also be set or changed at any time with the condition command, just as with breakpoints.

Unlike breakpoint conditions, GDB does not actually evaluate the conditional expression itself. Instead, GDB encodes the expression into an agent expression (see Appendix E [Agent

Expressions], page 453 suitable for execution on the target, independently of GDB. Global variables become raw memory locations, locals become stack accesses, and so forth.

For instance, suppose you have a function that is usually called frequently, but should not be called after an error has occurred. You could use the following tracepoint command to collect data about calls of that function that happen while the error code is propagating through the program; an unconditional tracepoint could end up collecting thousands of useless trace frames that you would have to search through.

(gdb) trace normal_operation if errcode > 0

13.1.5 Trace State Variables

A trace state variable is a special type of variable that is created and managed by target-side code. The syntax is the same as that for GDB's convenience variables (a string prefixed with "\$"), but they are stored on the target. They must be created explicitly, using a tvariable command. They are always 64-bit signed integers.

Trace state variables are remembered by GDB, and downloaded to the target along with tracepoint information when the trace experiment starts. There are no intrinsic limits on the number of trace state variables, beyond memory limitations of the target.

Although trace state variables are managed by the target, you can use them in print commands and expressions as if they were convenience variables; GDB will get the current value from the target while the trace experiment is running. Trace state variables share the same namespace as other "\$" variables, which means that you cannot have trace state variables with names like \$23 or \$pc, nor can you have a trace state variable and a convenience variable with the same name.

tvariable \$name [= expression]

The tvariable command creates a new trace state variable named \$name, and optionally gives it an initial value of expression. expression is evaluated when this command is entered; the result will be converted to an integer if possible, otherwise GDB will report an error. A subsequent tvariable command specifying the same name does not create a variable, but instead assigns the supplied initial value to the existing variable of that name, overwriting any previous initial value. The default initial value is 0.

info tvariables

List all the trace state variables along with their initial values. Their current values may also be displayed, if the trace experiment is currently running.

delete tvariable [\$name ...]

Delete the given trace state variables, or all of them if no arguments are specified.

13.1.6 Tracepoint Action Lists

actions [num]

This command will prompt for a list of actions to be taken when the tracepoint is hit. If the tracepoint number *num* is not specified, this command sets the actions for the one that was most recently defined (so that you can define a tracepoint and then say actions without bothering about its number). You

specify the actions themselves on the following lines, one action at a time, and terminate the actions list with a line containing just end. So far, the only defined actions are collect and while-stepping.

To remove all actions from a tracepoint, type 'actions num' and follow it immediately with 'end'.

```
(gdb) collect data // collect some data
(gdb) while-stepping 5 // single-step 5 times, collect data
(gdb) end // signals the end of actions.
```

In the following example, the action list begins with collect commands indicating the things to be collected when the tracepoint is hit. Then, in order to single-step and collect additional data following the tracepoint, a while-stepping command is used, followed by the list of things to be collected while stepping. The while-stepping command is terminated by its own separate end command. Lastly, the action list is terminated by an end command.

collect expr1, expr2, ...

Collect values of the given expressions when the tracepoint is hit. This command accepts a comma-separated list of any valid expressions. In addition to global, static, or local variables, the following special arguments are supported:

\$regs collect all registers

\$args collect all function arguments

\$locals collect all local variables.

You can give several consecutive collect commands, each one with a single argument, or one collect command with several arguments separated by commas: the effect is the same.

The command info scope (see Chapter 16 [Symbols], page 175) is particularly useful for figuring out what data to collect.

teval expr1, expr2, ...

Evaluate the given expressions when the tracepoint is hit. This command accepts a comma-separated list of expressions. The results are discarded, so this is mainly useful for assigning values to trace state variables (see Section 13.1.5 [Trace State Variables], page 132) without adding those values to the trace buffer, as would be the case if the collect action were used.

while-stepping n

Perform n single-step traces after the tracepoint, collecting new data at each step. The while-stepping command is followed by the list of what to collect while stepping (followed by its own end command):

```
> while-stepping 12
> collect $regs, myglobal
> end
>
```

You may abbreviate while-stepping as ws or stepping.

```
set default-collect expr1, expr2, ...
```

This variable is a list of expressions to collect at each tracepoint hit. It is effectively an additional collect action prepended to every tracepoint action list. The expressions are parsed individually for each tracepoint, so for instance a variable named xyz may be interpreted as a global for one tracepoint, and a local for another, as appropriate to the tracepoint's location.

show default-collect

Show the list of expressions that are collected by default at each tracepoint hit.

13.1.7 Listing Tracepoints

info tracepoints [num]

Display information about the tracepoint num. If you don't specify a tracepoint number, displays information about all the tracepoints defined so far. The format is similar to that used for info breakpoints; in fact, info tracepoints is the same command, simply restricting itself to tracepoints.

A tracepoint's listing may include additional information specific to tracing:

- its passcount as given by the passcount n command
- its step count as given by the while-stepping n command
- its action list as given by the actions command. The actions are prefixed with an 'A' so as to distinguish them from commands.

```
(gdb) info trace

Num Type Disp Enb Address What

1 tracepoint keep y 0x0804ab57 in foo() at main.cxx:7
    pass count 1200
    step count 20
    A while-stepping 20
    A collect globfoo, $regs
    A end
    A collect globfoo2
    A end
(gdb)
```

This command can be abbreviated info tp.

13.1.8 Starting and Stopping Trace Experiments

This command takes no arguments. It starts the trace experiment, and begins collecting data. This has the side effect of discarding all the data collected in the trace buffer during the previous trace experiment.

tstop

This command takes no arguments. It ends the trace experiment, and stops collecting data.

Note: a trace experiment and data collection may stop automatically if any tracepoint's passcount is reached (see Section 13.1.3 [Tracepoint Passcounts], page 131), or if the trace buffer becomes full.

tstatus This command displays the status of the current trace data collection.

Here is an example of the commands we described so far:

You can choose to continue running the trace experiment even if GDB disconnects from the target, voluntarily or involuntarily. For commands such as detach, the debugger will ask what you want to do with the trace. But for unexpected terminations (GDB crash, network outage), it would be unfortunate to lose hard-won trace data, so the variable disconnected-tracing lets you decide whether the trace should continue running without GDB.

```
set disconnected-tracing on set disconnected-tracing off
```

Choose whether a tracing run should continue to run if GDB has disconnected from the target. Note that detach or quit will ask you directly what to do about a running trace no matter what this variable's setting, so the variable is mainly useful for handling unexpected situations, such as loss of the network.

show disconnected-tracing

Show the current choice for disconnected tracing.

When you reconnect to the target, the trace experiment may or may not still be running; it might have filled the trace buffer in the meantime, or stopped for one of the other reasons. If it is running, it will continue after reconnection.

Upon reconnection, the target will upload information about the tracepoints in effect. GDB will then compare that information to the set of tracepoints currently defined, and attempt to match them up, allowing for the possibility that the numbers may have changed due to creation and deletion in the meantime. If one of the target's tracepoints does not match any in GDB, the debugger will create a new tracepoint, so that you have a number with which to specify that tracepoint. This matching-up process is necessarily heuristic, and it may result in useless tracepoints being created; you may simply delete them if they are of no use.

13.2 Using the Collected Data

After the tracepoint experiment ends, you use GDB commands for examining the trace data. The basic idea is that each tracepoint collects a trace snapshot every time it is hit and another snapshot every time it single-steps. All these snapshots are consecutively numbered from zero and go into a buffer, and you can examine them later. The way you examine them is to focus on a specific trace snapshot. When the remote stub is focused on a trace snapshot, it will respond to all GDB requests for memory and registers by reading from the buffer which belongs to that snapshot, rather than from real memory or registers of the program being debugged. This means that all GDB commands (print, info registers, backtrace, etc.) will behave as if we were currently debugging the program state as it was when the tracepoint occurred. Any requests for data that are not in the buffer will fail.

13.2.1 tfind n

The basic command for selecting a trace snapshot from the buffer is tfind n, which finds trace snapshot number n, counting from zero. If no argument n is given, the next snapshot is selected.

Here are the various forms of using the tfind command.

tfind start

Find the first snapshot in the buffer. This is a synonym for tfind 0 (since 0 is the number of the first snapshot).

tfind none

Stop debugging trace snapshots, resume live debugging.

tfind end Same as 'tfind none'.

tfind No argument means find the next trace snapshot.

tfind - Find the previous trace snapshot before the current one. This permits retracing earlier steps.

tfind tracepoint num

Find the next snapshot associated with tracepoint *num*. Search proceeds forward from the last examined trace snapshot. If no argument *num* is given, it means find the next snapshot collected for the same tracepoint as the current snapshot.

tfind pc addr

Find the next snapshot associated with the value addr of the program counter. Search proceeds forward from the last examined trace snapshot. If no argument addr is given, it means find the next snapshot with the same value of PC as the current snapshot.

tfind outside addr1, addr2

Find the next snapshot whose PC is outside the given range of addresses (exclusive).

tfind range addr1, addr2

Find the next snapshot whose PC is between addr1 and addr2 (inclusive).

tfind line [file:]n

Find the next snapshot associated with the source line n. If the optional argument file is given, refer to line n in that source file. Search proceeds forward from the last examined trace snapshot. If no argument n is given, it means find the next line other than the one currently being examined; thus saying tfind line repeatedly can appear to have the same effect as stepping from line to line in a live debugging session.

The default arguments for the tfind commands are specifically designed to make it easy to scan through the trace buffer. For instance, tfind with no argument selects the next trace snapshot, and tfind – with no argument selects the previous trace snapshot. So, by giving one tfind command, and then simply hitting RET repeatedly you can examine all the trace snapshots in order. Or, by saying tfind – and then hitting RET repeatedly you can examine the snapshots in reverse order. The tfind line command with no argument selects the snapshot for the next source line executed. The tfind pc command with no argument selects the next snapshot with the same program counter (PC) as the current frame. The tfind tracepoint command with no argument selects the next trace snapshot collected by the same tracepoint as the current one.

In addition to letting you scan through the trace buffer manually, these commands make it easy to construct GDB scripts that scan through the trace buffer and print out whatever collected data you are interested in. Thus, if we want to examine the PC, FP, and SP registers from each trace frame in the buffer, we can say this:

```
(gdb) tfind start
(gdb) while ($trace_frame != -1)
> printf "Frame %d, PC = %08X, SP = %08X, FP = %08X\n", \
          $trace_frame, $pc, $sp, $fp
> tfind
> end
Frame 0, PC = 0020DC64, SP = 0030BF3C, FP = 0030BF44
Frame 1, PC = 0020DC6C, SP = 0030BF38, FP = 0030BF44
Frame 2, PC = 0020DC70, SP = 0030BF34, FP = 0030BF44
Frame 3, PC = 0020DC74, SP = 0030BF30, FP = 0030BF44
Frame 4, PC = 0020DC78, SP = 0030BF2C, FP = 0030BF44
Frame 5, PC = 0020DC7C, SP = 0030BF28, FP = 0030BF44
Frame 6, PC = 0020DC80, SP = 0030BF24, FP = 0030BF44
Frame 7, PC = 0020DC84, SP = 0030BF20, FP = 0030BF44
Frame 8, PC = 0020DC88, SP = 0030BF1C, FP = 0030BF44
Frame 9, PC = 0020DC8E, SP = 0030BF18, FP = 0030BF44
Frame 10, PC = 00203F6C, SP = 0030BE3C, FP = 0030BF14
```

Or, if we want to examine the variable X at each source line in the buffer:

```
(gdb) tfind start
(gdb) while ($trace_frame != -1)
> printf "Frame %d, X == %d\n", $trace_frame, X
> tfind line
> end

Frame 0, X = 1
Frame 7, X = 2
Frame 13, X = 255
```

13.2.2 tdump

This command takes no arguments. It prints all the data collected at the current trace snapshot.

```
(gdb) trace 444
(gdb) actions
Enter actions for tracepoint #2, one per line:
> collect $regs, $locals, $args, gdb_long_test
> end
(gdb) tstart
(gdb) tfind line 444
#0 gdb_test (p1=0x11, p2=0x22, p3=0x33, p4=0x44, p5=0x55, p6=0x66)
at gdb_test.c:444
           printp( "%s: arguments = 0x\%X 0x\%X 0x\%X 0x\%X 0x\%X 0x\%X \n", )
(gdb) tdump
Data collected at tracepoint 2, trace frame 1:
d0
               0xc4aa0085
                                 -995491707
d1
               0x18
                         24
               08x0
d2
                         128
               0x33
d3
                         51
               0x71aea3d
                                 119204413
d4
d5
               0x22
                         34
d6
               0xe0
                         224
               0x380035 3670069
d7
               0x19e24a 1696330
a0
               0x3000668
                                 50333288
a1
               0x100
                         256
a2
               0x322000 3284992
a3
               0x3000698
                                 50333336
a4
               0x1ad3cc 1758156
a5
               0x30bf3c 0x30bf3c
fp
               0x30bf34 0x30bf34
sp
ps
               0x20b2c8 0x20b2c8
рс
fpcontrol
               0x0
                         0
fpstatus
               0x0
                         0
fpiaddr
               0x0
p = 0x20e5b4 "gdb-test"
p1 = (void *) 0x11
p2 = (void *) 0x22
p3 = (void *) 0x33
p4 = (void *) 0x44
p5 = (void *) 0x55
p6 = (void *) 0x66
gdb_long_test = 17 '\021'
(gdb)
```

13.2.3 save-tracepoints filename

This command saves all current tracepoint definitions together with their actions and pass-counts, into a file 'filename' suitable for use in a later debugging session. To read the saved tracepoint definitions, use the source command (see Section 23.1.3 [Command Files], page 258).

13.3 Convenience Variables for Tracepoints

(int) \$trace_frame

The current trace snapshot (a.k.a. frame) number, or -1 if no snapshot is selected.

(int) \$tracepoint

The tracepoint for the current trace snapshot.

(int) \$trace_line

The line number for the current trace snapshot.

(char []) \$trace_file

The source file for the current trace snapshot.

(char []) \$trace_func

The name of the function containing \$tracepoint.

Note: \$trace_file is not suitable for use in printf, use output instead.

Here's a simple example of using these convenience variables for stepping through all the trace snapshots and printing some of their data. Note that these are not the same as trace state variables, which are managed by the target.

(gdb) tfind start

```
(gdb) while $trace_frame != -1
> output $trace_file
> printf ", line %d (tracepoint #%d)\n", $trace_line, $tracepoint
> tfind
> end
```

13.4 Using Trace Files

In some situations, the target running a trace experiment may no longer be available; perhaps it crashed, or the hardware was needed for a different activity. To handle these cases, you can arrange to dump the trace data into a file, and later use that file as a source of trace data, via the target tfile command.

tsave [-r] filename

Save the trace data to filename. By default, this command assumes that filename refers to the host filesystem, so if necessary GDB will copy raw trace data up from the target and then save it. If the target supports it, you can also supply the optional argument -r ("remote") to direct the target to save the data directly into filename in its own filesystem, which may be more efficient if the trace buffer is very large. (Note, however, that target tfile can only read from files accessible to the host.)

target tfile filename

Use the file named *filename* as a source of trace data. Commands that examine data work as they do with a live target, but it is not possible to run any new trace experiments. **tstatus** will report the state of the trace run at the moment the data was saved, as well as the current trace frame you are examining. *filename* must be on a filesystem accessible to the host.

14 Debugging Programs That Use Overlays

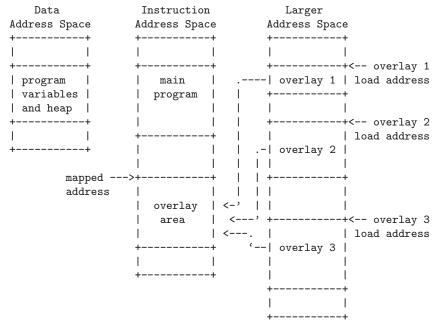
If your program is too large to fit completely in your target system's memory, you can sometimes use *overlays* to work around this problem. GDB provides some support for debugging programs that use overlays.

14.1 How Overlays Work

Suppose you have a computer whose instruction address space is only 64 kilobytes long, but which has much more memory which can be accessed by other means: special instructions, segment registers, or memory management hardware, for example. Suppose further that you want to adapt a program which is larger than 64 kilobytes to run on this system.

One solution is to identify modules of your program which are relatively independent, and need not call each other directly; call these modules overlays. Separate the overlays from the main program, and place their machine code in the larger memory. Place your main program in instruction memory, but leave at least enough space there to hold the largest overlay as well.

Now, to call a function located in an overlay, you must first copy that overlay's machine code from the large memory into the space set aside for it in the instruction memory, and then jump to its entry point there.



A code overlay

The diagram (see [A code overlay], page 141) shows a system with separate data and instruction address spaces. To map an overlay, the program copies its code from the larger address space to the instruction address space. Since the overlays shown here all use the same mapped address, only one may be mapped at a time. For a system with a single address space for data and instructions, the diagram would be similar, except that the program variables and heap would share an address space with the main program and the overlay area.

An overlay loaded into instruction memory and ready for use is called a *mapped* overlay; its *mapped address* is its address in the instruction memory. An overlay not present (or only partially present) in instruction memory is called *unmapped*; its *load address* is its address in the larger memory. The mapped address is also called the *virtual memory address*, or *VMA*; the load address is also called the *load memory address*, or *LMA*.

Unfortunately, overlays are not a completely transparent way to adapt a program to limited instruction memory. They introduce a new set of global constraints you must keep in mind as you design your program:

- Before calling or returning to a function in an overlay, your program must make sure that overlay is actually mapped. Otherwise, the call or return will transfer control to the right address, but in the wrong overlay, and your program will probably crash.
- If the process of mapping an overlay is expensive on your system, you will need to choose your overlays carefully to minimize their effect on your program's performance.
- The executable file you load onto your system must contain each overlay's instructions, appearing at the overlay's load address, not its mapped address. However, each overlay's instructions must be relocated and its symbols defined as if the overlay were at its mapped address. You can use GNU linker scripts to specify different load and relocation addresses for pieces of your program; see Section "Overlay Description" in Using Id: the GNU linker.
- The procedure for loading executable files onto your system must be able to load their contents into the larger address space as well as the instruction and data spaces.

The overlay system described above is rather simple, and could be improved in many ways:

- If your system has suitable bank switch registers or memory management hardware, you could use those facilities to make an overlay's load area contents simply appear at their mapped address in instruction space. This would probably be faster than copying the overlay to its mapped area in the usual way.
- If your overlays are small enough, you could set aside more than one overlay area, and have more than one overlay mapped at a time.
- You can use overlays to manage data, as well as instructions. In general, data overlays
 are even less transparent to your design than code overlays: whereas code overlays only
 require care when you call or return to functions, data overlays require care every time
 you access the data. Also, if you change the contents of a data overlay, you must copy
 its contents back out to its load address before you can copy a different data overlay
 into the same mapped area.

14.2 Overlay Commands

To use GDB's overlay support, each overlay in your program must correspond to a separate section of the executable file. The section's virtual memory address and load memory address must be the overlay's mapped and load addresses. Identifying overlays with sections allows GDB to determine the appropriate address of a function or variable, depending on whether the overlay is mapped or not.

GDB's overlay commands all start with the word overlay; you can abbreviate this as ov or ovly. The commands are:

overlay off

Disable GDB's overlay support. When overlay support is disabled, GDB assumes that all functions and variables are always present at their mapped addresses. By default, GDB's overlay support is disabled.

overlay manual

Enable manual overlay debugging. In this mode, GDB relies on you to tell it which overlays are mapped, and which are not, using the overlay map-overlay and overlay unmap-overlay commands described below.

overlay map-overlay overlay overlay

Tell GDB that overlay is now mapped; overlay must be the name of the object file section containing the overlay. When an overlay is mapped, GDB assumes it can find the overlay's functions and variables at their mapped addresses. GDB assumes that any other overlays whose mapped ranges overlap that of overlay are now unmapped.

overlay unmap-overlay overlay overlay

Tell GDB that overlay is no longer mapped; overlay must be the name of the object file section containing the overlay. When an overlay is unmapped, GDB assumes it can find the overlay's functions and variables at their load addresses.

overlay auto

Enable automatic overlay debugging. In this mode, GDB consults a data structure the overlay manager maintains in the inferior to see which overlays are mapped. For details, see Section 14.3 [Automatic Overlay Debugging], page 144.

overlay load-target overlay load

Re-read the overlay table from the inferior. Normally, GDB re-reads the table GDB automatically each time the inferior stops, so this command should only be necessary if you have changed the overlay mapping yourself using GDB. This command is only useful when using automatic overlay debugging.

overlay list-overlays overlay list

Display a list of the overlays currently mapped, along with their mapped addresses, load addresses, and sizes.

Normally, when gdb prints a code address, it includes the name of the function the address falls in:

```
(gdb) print main
$3 = {int ()} 0x11a0 <main>
```

When overlay debugging is enabled, GDB recognizes code in unmapped overlays, and prints the names of unmapped functions with asterisks around them. For example, if foo is a function in an unmapped overlay, GDB prints it this way:

```
(gdb) overlay list
No sections are mapped.
(gdb) print foo
```

Debugging with GDB

```
$5 = {int (int)} 0x100000 <*foo*>
```

When foo's overlay is mapped, GDB prints the function's name normally:

When overlay debugging is enabled, GDB can find the correct address for functions and variables in an overlay, whether or not the overlay is mapped. This allows most GDB commands, like break and disassemble, to work normally, even on unmapped code. However, GDB's breakpoint support has some limitations:

- You can set breakpoints in functions in unmapped overlays, as long as GDB can write to the overlay at its load address.
- GDB can not set hardware or simulator-based breakpoints in unmapped overlays. However, if you set a breakpoint at the end of your overlay manager (and tell GDB which overlays are now mapped, if you are using manual overlay management), GDB will re-set its breakpoints properly.

14.3 Automatic Overlay Debugging

GDB can automatically track which overlays are mapped and which are not, given some simple co-operation from the overlay manager in the inferior. If you enable automatic overlay debugging with the overlay auto command (see Section 14.2 [Overlay Commands], page 142), GDB looks in the inferior's memory for certain variables describing the current state of the overlays.

Here are the variables your overlay manager must define to support GDB's automatic overlay debugging:

_ovly_table:

This variable must be an array of the following structures:

```
struct
{
    /* The overlay's mapped address. */
    unsigned long vma;

    /* The size of the overlay, in bytes. */
    unsigned long size;

    /* The overlay's load address. */
    unsigned long lma;

    /* Non-zero if the overlay is currently mapped;
    zero otherwise. */
    unsigned long mapped;
}
```

_novlys: This variable must be a four-byte signed integer, holding the total number of elements in _ovly_table.

To decide whether a particular overlay is mapped or not, GDB looks for an entry in _ovly_table whose vma and lma members equal the VMA and LMA of the overlay's section

in the executable file. When gdb finds a matching entry, it consults the entry's mapped member to determine whether the overlay is currently mapped.

In addition, your overlay manager may define a function called <code>_ovly_debug_event</code>. If this function is defined, GDB will silently set a breakpoint there. If the overlay manager then calls this function whenever it has changed the overlay table, this will enable GDB to accurately keep track of which overlays are in program memory, and update any breakpoints that may be set in overlays. This will allow breakpoints to work even if the overlays are kept in ROM or other non-writable memory while they are not being executed.

14.4 Overlay Sample Program

When linking a program which uses overlays, you must place the overlays at their load addresses, while relocating them to run at their mapped addresses. To do this, you must write a linker script (see Section "Overlay Description" in *Using Id: the GNU linker*). Unfortunately, since linker scripts are specific to a particular host system, target architecture, and target memory layout, this manual cannot provide portable sample code demonstrating GDB's overlay support.

However, the GDB source distribution does contain an overlaid program, with linker scripts for a few systems, as part of its test suite. The program consists of the following files from 'gdb/testsuite/gdb.base':

```
'overlays.c'
           The main program file.
'ovlymgr.c'
           A simple overlay manager, used by 'overlays.c'.
'foo.c'
'bar.c'
'baz.c'
           Overlay modules, loaded and used by 'overlays.c'.
'grbx.c'
'd10v.ld'
'm32r.1d'
           Linker scripts for linking the test program on the d10v-elf and m32r-elf
           targets.
   You can build the test program using the d10v-elf GCC cross-compiler like this:
      $ d10v-elf-gcc -g -c overlays.c
      $ d10v-elf-gcc -g -c ovlymgr.c
      $ d10v-elf-gcc -g -c foo.c
      $ d10v-elf-gcc -g -c bar.c
      $ d10v-elf-gcc -g -c baz.c
      $ d10v-elf-gcc -g -c grbx.c
```

The build process is identical for any other architecture, except that you must substitute the appropriate compiler and linker script for the target system for d10v-elf-gcc and d10v.ld.

baz.o grbx.o -Wl,-Td10v.ld -o overlays

\$ d10v-elf-gcc -g overlays.o ovlymgr.o foo.o bar.o \

15 Using gdb with Different Languages

Although programming languages generally have common aspects, they are rarely expressed in the same manner. For instance, in ANSI C, dereferencing a pointer p is accomplished by *p, but in Modula-2, it is accomplished by p^. Values can also be represented (and displayed) differently. Hex numbers in C appear as '0x1ae', while in Modula-2 they appear as '1AEH'.

Language-specific information is built into GDB for some languages, allowing you to express operations like the above in your program's native language, and allowing GDB to output values in a manner consistent with the syntax of your program's native language. The language you use to build expressions is called the working language.

15.1 Switching Between Source Languages

There are two ways to control the working language—either have GDB set it automatically, or select it manually yourself. You can use the set language command for either purpose. On startup, GDB defaults to setting the language automatically. The working language is used to determine how expressions you type are interpreted, how values are printed, etc.

In addition to the working language, every source file that GDB knows about has its own working language. For some object file formats, the compiler might indicate which language a particular source file is in. However, most of the time GDB infers the language from the name of the file. The language of a source file controls whether C++ names are demangled—this way backtrace can show each frame appropriately for its own language. There is no way to set the language of a source file from within GDB, but you can set the language associated with a filename extension. See Section 15.2 [Displaying the Language], page 148.

This is most commonly a problem when you use a program, such as cfront or f2c, that generates C but is written in another language. In that case, make the program use #line directives in its C output; that way GDB will know the correct language of the source code of the original program, and will display that source code, not the generated C code.

15.1.1 List of Filename Extensions and Languages

If a source file name ends in one of the following extensions, then GDB infers that its language is the one indicated.

```
'.ada'
'.ads'
'.adb'
'.a' Ada source file.
'.c' C source file
'.C'
'.cc'
'.cc'
'.cp'
'.cpp'
'.cxx'
'.c++' C++ source file
```

```
'.m' Objective-C source file
'.f'
'.F' Fortran source file
'.mod' Modula-2 source file
'.s'
'.S' Assembler source file. This actually behaves almost like C, but GDB does not skip over function prologues when stepping.
```

In addition, you may set the language associated with a filename extension. See Section 15.2 [Displaying the Language], page 148.

15.1.2 Setting the Working Language

If you allow GDB to set the language automatically, expressions are interpreted the same way in your debugging session and your program.

If you wish, you may set the language manually. To do this, issue the command 'set language lang', where lang is the name of a language, such as c or modula-2. For a list of the supported languages, type 'set language'.

Setting the language manually prevents GDB from updating the working language automatically. This can lead to confusion if you try to debug a program when the working language is not the same as the source language, when an expression is acceptable to both languages—but means different things. For instance, if the current source file were written in C, and GDB was parsing Modula-2, a command such as:

```
print a = b + c
```

might not have the effect you intended. In C, this means to add b and c and place the result in a. The result printed would be the value of a. In Modula-2, this means to compare a to the result of b+c, yielding a BOOLEAN value.

15.1.3 Having GDB Infer the Source Language

To have GDB set the working language automatically, use 'set language local' or 'set language auto'. GDB then infers the working language. That is, when your program stops in a frame (usually by encountering a breakpoint), GDB sets the working language to the language recorded for the function in that frame. If the language for a frame is unknown (that is, if the function or block corresponding to the frame was defined in a source file that does not have a recognized extension), the current working language is not changed, and GDB issues a warning.

This may not seem necessary for most programs, which are written entirely in one source language. However, program modules and libraries written in one source language can be used by a main program written in a different source language. Using 'set language auto' in this case frees you from having to set the working language manually.

15.2 Displaying the Language

The following commands help you find out which language is the working language, and also what language source files were written in.

show language

Display the current working language. This is the language you can use with commands such as print to build and compute expressions that may involve variables in your program.

info frame

Display the source language for this frame. This language becomes the working language if you use an identifier from this frame. See Section 8.4 [Information about a Frame], page 81, to identify the other information listed here.

info source

Display the source language of this source file. See Chapter 16 [Examining the Symbol Table], page 175, to identify the other information listed here.

In unusual circumstances, you may have source files with extensions not in the standard list. You can then set the extension associated with a language explicitly:

set extension-language ext language

Tell GDB that source files with extension ext are to be assumed as written in the source language language.

info extensions

List all the filename extensions and the associated languages.

15.3 Type and Range Checking

Warning: In this release, the GDB commands for type and range checking are included, but they do not yet have any effect. This section documents the intended facilities.

Some languages are designed to guard you against making seemingly common errors through a series of compile- and run-time checks. These include checking the type of arguments to functions and operators, and making sure mathematical overflows are caught at run time. Checks such as these help to ensure a program's correctness once it has been compiled by eliminating type mismatches, and providing active checks for range errors when your program is running.

GDB can check for conditions like the above if you wish. Although GDB does not check the statements in your program, it can check expressions entered directly into GDB for evaluation via the print command, for example. As with the working language, GDB can also decide whether or not to check automatically based on your program's source language. See Section 15.4 [Supported Languages], page 151, for the default settings of supported languages.

15.3.1 An Overview of Type Checking

Some languages, such as Modula-2, are strongly typed, meaning that the arguments to operators and functions have to be of the correct type, otherwise an error occurs. These checks prevent type mismatch errors from ever causing any run-time problems. For example,

$$\begin{array}{c} 1 + 2 \Rightarrow 3 \\ \text{but} \\ \hline \text{error} \quad 1 + 2.3 \end{array}$$

The second example fails because the CARDINAL 1 is not type-compatible with the REAL 2.3.

For the expressions you use in GDB commands, you can tell the GDB type checker to skip checking; to treat any mismatches as errors and abandon the expression; or to only issue warnings when type mismatches occur, but evaluate the expression anyway. When you choose the last of these, GDB evaluates expressions like the second example above, but also issues a warning.

Even if you turn type checking off, there may be other reasons related to type that prevent GDB from evaluating an expression. For instance, GDB does not know how to add an int and a struct foo. These particular type errors have nothing to do with the language in use, and usually arise from expressions, such as the one described above, which make little sense to evaluate anyway.

Each language defines to what degree it is strict about type. For instance, both Modula-2 and C require the arguments to arithmetical operators to be numbers. In C, enumerated types and pointers can be represented as numbers, so that they are valid arguments to mathematical operators. See Section 15.4 [Supported Languages], page 151, for further details on specific languages.

GDB provides some additional commands for controlling the type checker:

set check type auto

Set type checking on or off based on the current working language. See Section 15.4 [Supported Languages], page 151, for the default settings for each language.

set check type on set check type off

Set type checking on or off, overriding the default setting for the current working language. Issue a warning if the setting does not match the language default. If any type mismatches occur in evaluating an expression while type checking is on, GDB prints a message and aborts evaluation of the expression.

set check type warn

Cause the type checker to issue warnings, but to always attempt to evaluate the expression. Evaluating the expression may still be impossible for other reasons. For example, GDB cannot add numbers and structures.

show type Show the current setting of the type checker, and whether or not gdb is setting it automatically.

15.3.2 An Overview of Range Checking

In some languages (such as Modula-2), it is an error to exceed the bounds of a type; this is enforced with run-time checks. Such range checking is meant to ensure program correctness by making sure computations do not overflow, or indices on an array element access do not exceed the bounds of the array.

For expressions you use in GDB commands, you can tell GDB to treat range errors in one of three ways: ignore them, always treat them as errors and abandon the expression, or issue warnings but evaluate the expression anyway.

A range error can result from numerical overflow, from exceeding an array index bound, or when you type a constant that is not a member of any type. Some languages, however, do not treat overflows as an error. In many implementations of C, mathematical overflow causes the result to "wrap around" to lower values—for example, if m is the largest integer value, and s is the smallest, then

```
m + 1 \Rightarrow s
```

This, too, is specific to individual languages, and in some cases specific to individual compilers or machines. See Section 15.4 [Supported Languages], page 151, for further details on specific languages.

GDB provides some additional commands for controlling the range checker:

set check range auto

Set range checking on or off based on the current working language. See Section 15.4 [Supported Languages], page 151, for the default settings for each language.

```
set check range on set check range off
```

Set range checking on or off, overriding the default setting for the current working language. A warning is issued if the setting does not match the language default. If a range error occurs and range checking is on, then a message is printed and evaluation of the expression is aborted.

set check range warn

Output messages when the GDB range checker detects a range error, but attempt to evaluate the expression anyway. Evaluating the expression may still be impossible for other reasons, such as accessing memory that the process does not own (a typical example from many Unix systems).

show range

Show the current setting of the range checker, and whether or not it is being set automatically by GDB.

15.4 Supported Languages

GDB supports C, C++, Objective-C, Fortran, Java, Pascal, assembly, Modula-2, and Ada. Some GDB features may be used in expressions regardless of the language you use: the GDB @ and :: operators, and the '{type}addr' construct (see Section 10.1 [Expressions], page 93) can be used with the constructs of any supported language.

The following sections detail to what degree each source language is supported by GDB. These sections are not meant to be language tutorials or references, but serve only as a reference guide to what the GDB expression parser accepts, and what input and output formats should look like for different languages. There are many good books written on each of these languages; please look to these for a language reference or tutorial.

15.4.1 C and C++

Since C and C++ are so closely related, many features of GDB apply to both languages. Whenever this is the case, we discuss those languages together.

The C++ debugging facilities are jointly implemented by the C++ compiler and GDB. Therefore, to debug your C++ code effectively, you must compile your C++ programs with a supported C++ compiler, such as GNU g++, or the HP ANSI C++ compiler (aCC).

For best results when using GNU C++, use the DWARF 2 debugging format; if it doesn't work on your system, try the stabs+ debugging format. You can select those formats explicitly with the g++ command-line options '-gdwarf-2' and '-gstabs+'. See Section "Options for Debugging Your Program or GCC" in *Using the GNU Compiler Collection (GCC)*.

15.4.1.1 C and C++ Operators

Operators must be defined on values of specific types. For instance, + is defined on numbers, but not on structures. Operators are often defined on groups of types.

For the purposes of C and C++, the following definitions hold:

- Integral types include int with any of its storage-class specifiers; char; enum; and, for C++, bool.
- Floating-point types include float, double, and long double (if supported by the target platform).
- Pointer types include all types defined as (type *).
- Scalar types include all of the above.

The following operators are supported. They are listed here in order of increasing precedence:

- , The comma or sequencing operator. Expressions in a comma-separated list are evaluated from left to right, with the result of the entire expression being the last expression evaluated.
- = Assignment. The value of an assignment expression is the value assigned. Defined on scalar types.
- op= Used in an expression of the form a op=b, and translated to a=a op b. op= and = have the same precedence. op is any one of the operators $|, \hat{,} \&, <<, >>, +, -, *, /, %.$
- ?: The ternary operator. a ? b : c can be thought of as: if a then b else c. a should be of an integral type.
- Logical OR. Defined on integral types.
- && Logical AND. Defined on integral types.
- Bitwise OR. Defined on integral types.
- [^] Bitwise exclusive-OR. Defined on integral types.
- & Bitwise AND. Defined on integral types.
- ==, != Equality and inequality. Defined on scalar types. The value of these expressions is 0 for false and non-zero for true.
- <, >, <=, >=

Less than, greater than, less than or equal, greater than or equal. Defined on scalar types. The value of these expressions is 0 for false and non-zero for true.

- <<, >> left shift, and right shift. Defined on integral types.
- © The GDB "artificial array" operator (see Section 10.1 [Expressions], page 93).
- +, Addition and subtraction. Defined on integral types, floating-point types and pointer types.
- *, /, % Multiplication, division, and modulus. Multiplication and division are defined on integral and floating-point types. Modulus is defined on integral types.
- ++, -- Increment and decrement. When appearing before a variable, the operation is performed before the variable is used in an expression; when appearing after it, the variable's value is used before the operation takes place.
- * Pointer dereferencing. Defined on pointer types. Same precedence as ++.
- & Address operator. Defined on variables. Same precedence as ++.

 For debugging C++, GDB implements a use of '&' beyond what is allowed in the C++ language itself: you can use '&(&ref)' to examine the address where a C++ reference variable (declared with '&ref') is stored.
- Negative. Defined on integral and floating-point types. Same precedence as ++.
- ! Logical negation. Defined on integral types. Same precedence as ++.
- Bitwise complement operator. Defined on integral types. Same precedence as ++.
- ., -> Structure member, and pointer-to-structure member. For convenience, GDB regards the two as equivalent, choosing whether to dereference a pointer based on the stored type information. Defined on struct and union data.
- .*, ->* Dereferences of pointers to members.
- [] Array indexing. a[i] is defined as *(a+i). Same precedence as ->.
- () Function parameter list. Same precedence as ->.
- :: C++ scope resolution operator. Defined on struct, union, and class types.
- :: Doubled colons also represent the GDB scope operator (see Section 10.1 [Expressions], page 93). Same precedence as ::, above.

If an operator is redefined in the user code, GDB usually attempts to invoke the redefined version instead of using the operator's predefined meaning.

15.4.1.2 C and C++ Constants

GDB allows you to express the constants of C and C++ in the following ways:

- Integer constants are a sequence of digits. Octal constants are specified by a leading '0' (i.e. zero), and hexadecimal constants by a leading '0x' or '0X'. Constants may also end with a letter '1', specifying that the constant should be treated as a long value.
- Floating point constants are a sequence of digits, followed by a decimal point, followed by a sequence of digits, and optionally followed by an exponent. An exponent is of the form: 'e[[+]|-]nnn', where nnn is another sequence of digits. The '+' is optional for positive exponents. A floating-point constant may also end with a letter 'f' or 'F', specifying that the constant should be treated as being of the float (as opposed to the default double) type; or with a letter 'l' or 'L', which specifies a long double constant.

• Enumerated constants consist of enumerated identifiers, or their integral equivalents.

- Character constants are a single character surrounded by single quotes ('), or a number—the ordinal value of the corresponding character (usually its ASCII value). Within quotes, the single character may be represented by a letter or by escape sequences, which are of the form '\nnn', where nnn is the octal representation of the character's ordinal value; or of the form '\x', where 'x' is a predefined special character—for example, '\n' for newline.
- String constants are a sequence of character constants surrounded by double quotes ("). Any valid character constant (as described above) may appear. Double quotes within the string must be preceded by a backslash, so for instance '"a\"b'c"' is a string of five characters.
- Pointer constants are an integral value. You can also write pointers to constants using the C operator '&'.
- Array constants are comma-separated lists surrounded by braces '{' and '}'; for example, '{1,2,3}' is a three-element array of integers, '{{1,2}, {3,4}, {5,6}}' is a three-by-two array, and '{&"hi", &"there", &"fred"}' is a three-element array of pointers.

15.4.1.3 C++ Expressions

GDB expression handling can interpret most C++ expressions.

Warning: GDB can only debug C++ code if you use the proper compiler and the proper debug format. Currently, GDB works best when debugging C++ code that is compiled with GCC 2.95.3 or with GCC 3.1 or newer, using the options '-gdwarf-2' or '-gstabs+'. DWARF 2 is preferred over stabs+. Most configurations of GCC emit either DWARF 2 or stabs+ as their default debug format, so you usually don't need to specify a debug format explicitly. Other compilers and/or debug formats are likely to work badly or not at all when using GDB to debug C++ code.

- Member function calls are allowed; you can use expressions like count = aml->GetOriginal(x, y)
- 2. While a member function is active (in the selected stack frame), your expressions have the same namespace available as the member function; that is, GDB allows implicit references to the class instance pointer this following the same rules as C++.
- 3. You can call overloaded functions; GDB resolves the function call to the right definition, with some restrictions. GDB does not perform overload resolution involving user-defined type conversions, calls to constructors, or instantiations of templates that do not exist in the program. It also cannot handle ellipsis argument lists or default arguments.
 - It does perform integral conversions and promotions, floating-point promotions, arithmetic conversions, pointer conversions, conversions of class objects to base classes, and standard conversions such as those of functions or arrays to pointers; it requires an exact match on the number of function arguments.

Overload resolution is always performed, unless you have specified set overload-resolution off. See Section 15.4.1.7 [GDB Features for C++], page 156.

You must specify set overload-resolution off in order to use an explicit function signature to call an overloaded function, as in

```
p 'foo(char,int)'('x', 13)
```

The GDB command-completion facility can simplify this; see Section 3.2 [Command Completion], page 19.

- 4. GDB understands variables declared as C++ references; you can use them in expressions just as you do in C++ source—they are automatically dereferenced.
 - In the parameter list shown when GDB displays a frame, the values of reference variables are not displayed (unlike other variables); this avoids clutter, since references are often used for large structures. The *address* of a reference variable is always shown, unless you have specified 'set print address off'.
- 5. GDB supports the C++ name resolution operator ::—your expressions can use it just as expressions in your program do. Since one scope may be defined in another, you can use :: repeatedly if necessary, for example in an expression like 'scope1::scope2::name'. GDB also allows resolving name scope by reference to source files, in both C and C++ debugging (see Section 10.3 [Program Variables], page 95).

In addition, when used with HP's C++ compiler, GDB supports calling virtual functions correctly, printing out virtual bases of objects, calling functions in a base subobject, casting objects, and invoking user-defined operators.

15.4.1.4 C and C++ Defaults

If you allow GDB to set type and range checking automatically, they both default to off whenever the working language changes to C or C++. This happens regardless of whether you or GDB selects the working language.

If you allow GDB to set the language automatically, it recognizes source files whose names end with '.c', '.C', or '.cc', etc, and when GDB enters code compiled from one of these files, it sets the working language to C or C++. See Section 15.1.3 [Having GDB Infer the Source Language], page 148, for further details.

15.4.1.5 C and C++ Type and Range Checks

By default, when gdb parses C or C++ expressions, type checking is not used. However, if you turn type checking on, gdb considers two variables type equivalent if:

- The two variables are structured and have the same structure, union, or enumerated tag.
- The two variables have the same type name, or types that have been declared equivalent through typedef.

Range checking, if turned on, is done on mathematical operations. Array indices are not checked, since they are often used to index a pointer that is not itself an array.

15.4.1.6 GDB and C

The set print union and show print union commands apply to the union type. When set to 'on', any union that is inside a struct or class is also printed. Otherwise, it appears as '{...}'.

The @ operator aids in the debugging of dynamic arrays, formed with pointers and a memory allocation function. See Section 10.1 [Expressions], page 93.

15.4.1.7 GDB Features for C++

Some GDB commands are particularly useful with C++, and some are designed specifically for use with C++. Here is a summary:

breakpoint menus

When you want a breakpoint in a function whose name is overloaded, GDB has the capability to display a menu of possible breakpoint locations to help you specify which function definition you want. See Section 10.2 [Ambiguous Expressions], page 94.

rbreak regex

Setting breakpoints using regular expressions is helpful for setting breakpoints on overloaded functions that are not members of any special classes. See Section 5.1.1 [Setting Breakpoints], page 44.

catch throw catch catch

Debug C++ exception handling using these commands. See Section 5.1.3 [Setting Catchpoints], page 51.

ptype typename

Print inheritance relationships as well as other information for type typename. See Chapter 16 [Examining the Symbol Table], page 175.

set print demangle show print demangle set print asm-demangle show print asm-demangle

Control whether C++ symbols display in their source form, both when displaying code as C++ source and when displaying disassemblies. See Section 10.8 [Print Settings], page 102.

set print object show print object

Choose whether to print derived (actual) or declared types of objects. See Section 10.8 [Print Settings], page 102.

set print vtbl show print vtbl

Control the format for printing virtual function tables. See Section 10.8 [Print Settings], page 102. (The vtbl commands do not work on programs compiled with the HP ANSI C++ compiler (aCC).)

set overload-resolution on

Enable overload resolution for C++ expression evaluation. The default is on. For overloaded functions, GDB evaluates the arguments and searches for a function whose signature matches the argument types, using the standard C++ conversion rules (see Section 15.4.1.3 [C++ Expressions], page 154, for details). If it cannot find a match, it emits a message.

set overload-resolution off

Disable overload resolution for C++ expression evaluation. For overloaded functions that are not class member functions, GDB chooses the first function of the specified name that it finds in the symbol table, whether or not its arguments are of the correct type. For overloaded functions that are class member functions, GDB searches for a function whose signature *exactly* matches the argument types.

show overload-resolution

Show the current setting of overload resolution.

Overloaded symbol names

You can specify a particular definition of an overloaded symbol, using the same notation that is used to declare such symbols in C++: type symbol (types) rather than just symbol. You can also use the GDB command-line word completion facilities to list the available choices, or to finish the type list for you. See Section 3.2 [Command Completion], page 19, for details on how to do this.

15.4.1.8 Decimal Floating Point format

GDB can examine, set and perform computations with numbers in decimal floating point format, which in the C language correspond to the _Decimal32, _Decimal64 and _Decimal128 types as specified by the extension to support decimal floating-point arithmetic.

There are two encodings in use, depending on the architecture: BID (Binary Integer Decimal) for x86 and x86-64, and DPD (Densely Packed Decimal) for PowerPC. GDB will use the appropriate encoding for the configured target.

Because of a limitation in 'libdecnumber', the library used by GDB to manipulate decimal floating point numbers, it is not possible to convert (using a cast, for example) integers wider than 32-bit to decimal float.

In addition, in order to imitate GDB's behaviour with binary floating point computations, error checking in decimal float operations ignores underflow, overflow and divide by zero exceptions.

In the PowerPC architecture, GDB provides a set of pseudo-registers to inspect _Decimal128 values stored in floating point registers. See Section 21.4.7 [PowerPC], page 244 for more details.

15.4.2 Objective-C

This section provides information about some commands and command options that are useful for debugging Objective-C code. See also Chapter 16 [Symbols], page 175, and Chapter 16 [Symbols], page 175, for a few more commands specific to Objective-C support.

15.4.2.1 Method Names in Commands

The following commands have been extended to accept Objective-C method names as line specifications:

- clear
- break
- info line

- jump
- list

A fully qualified Objective-C method name is specified as

-[Class methodName]

where the minus sign is used to indicate an instance method and a plus sign (not shown) is used to indicate a class method. The class name *Class* and method name *methodName* are enclosed in brackets, similar to the way messages are specified in Objective-C source code. For example, to set a breakpoint at the create instance method of class Fruit in the program currently being debugged, enter:

```
break -[Fruit create]
```

To list ten program lines around the initialize class method, enter:

```
list +[NSText initialize]
```

In the current version of GDB, the plus or minus sign is required. In future versions of GDB, the plus or minus sign will be optional, but you can use it to narrow the search. It is also possible to specify just a method name:

```
break create
```

You must specify the complete method name, including any colons. If your program's source files contain more than one **create** method, you'll be presented with a numbered list of classes that implement that method. Indicate your choice by number, or type '0' to exit if none apply.

As another example, to clear a breakpoint established at the makeKeyAndOrderFront: method of the NSWindow class, enter:

clear -[NSWindow makeKeyAndOrderFront:]

15.4.2.2 The Print Command With Objective-C

The print command has also been extended to accept methods. For example:

```
print -[object hash]
```

will tell GDB to send the hash message to *object* and print the result. Also, an additional command has been added, print-object or po for short, which is meant to print the description of an object. However, this command may only work with certain Objective-C libraries that have a particular hook function, _NSPrintForDebugger, defined.

15.4.3 Fortran

GDB can be used to debug programs written in Fortran, but it currently supports only the features of Fortran 77 language.

Some Fortran compilers (GNU Fortran 77 and Fortran 95 compilers among them) append an underscore to the names of variables and functions. When you debug programs compiled by those compilers, you will need to refer to variables and functions with a trailing underscore.

15.4.3.1 Fortran Operators and Expressions

Operators must be defined on values of specific types. For instance, + is defined on numbers, but not on characters or other non- arithmetic types. Operators are often defined on groups of types.

- ** The exponentiation operator. It raises the first operand to the power of the second one.
- : The range operator. Normally used in the form of array(low:high) to represent a section of array.
- The access component operator. Normally used to access elements in derived types. Also suitable for unions. As unions aren't part of regular Fortran, this can only happen when accessing a register that uses a gdbarch-defined union type.

15.4.3.2 Fortran Defaults

Fortran symbols are usually case-insensitive, so GDB by default uses case-insensitive matches for Fortran symbols. You can change that with the 'set case-insensitive' command, see Chapter 16 [Symbols], page 175, for the details.

15.4.3.3 Special Fortran Commands

GDB has some commands to support Fortran-specific features, such as displaying common blocks.

info common [common-name]

This command prints the values contained in the Fortran COMMON block whose name is *common-name*. With no argument, the names of all COMMON blocks visible at the current program location are printed.

15.4.4 Pascal

Debugging Pascal programs which use sets, subranges, file variables, or nested functions does not currently work. GDB does not support entering expressions, printing values, or similar features using Pascal syntax.

The Pascal-specific command set print pascal_static-members controls whether static members of Pascal objects are displayed. See Section 10.8 [Print Settings], page 102.

15.4.5 Modula-2

The extensions made to GDB to support Modula-2 only support output from the GNU Modula-2 compiler (which is currently being developed). Other Modula-2 compilers are not currently supported, and attempting to debug executables produced by them is most likely to give an error as GDB reads in the executable's symbol table.

15.4.5.1 Operators

Operators must be defined on values of specific types. For instance, + is defined on numbers, but not on structures. Operators are often defined on groups of types. For the purposes of Modula-2, the following definitions hold:

- Integral types consist of INTEGER, CARDINAL, and their subranges.
- Character types consist of CHAR and its subranges.
- Floating-point types consist of REAL.
- Pointer types consist of anything declared as POINTER TO type.
- Scalar types consist of all of the above.

- Set types consist of SET and BITSET types.
- Boolean types consist of BOOLEAN.

The following operators are supported, and appear in order of increasing precedence:

- Function argument or array index separator.
- := Assignment. The value of var := value is value.
- <, > Less than, greater than on integral, floating-point, or enumerated types.
- <=, >= Less than or equal to, greater than or equal to on integral, floating-point and enumerated types, or set inclusion on set types. Same precedence as <.
- =, <>, # Equality and two ways of expressing inequality, valid on scalar types. Same precedence as <. In GDB scripts, only <> is available for inequality, since # conflicts with the script comment character.
- IN Set membership. Defined on set types and the types of their members. Same precedence as <.
- OR Boolean disjunction. Defined on boolean types.
- AND, & Boolean conjunction. Defined on boolean types.
- © The GDB "artificial array" operator (see Section 10.1 [Expressions], page 93).
- +, Addition and subtraction on integral and floating-point types, or union and difference on set types.
- * Multiplication on integral and floating-point types, or set intersection on set types.
- / Division on floating-point types, or symmetric set difference on set types. Same precedence as *.
- DIV, MOD Integer division and remainder. Defined on integral types. Same precedence as *.
- Negative. Defined on INTEGER and REAL data.
- [^] Pointer dereferencing. Defined on pointer types.
- NOT Boolean negation. Defined on boolean types. Same precedence as ^.
- . RECORD field selector. Defined on RECORD data. Same precedence as ^.
- [] Array indexing. Defined on ARRAY data. Same precedence as ^.
- () Procedure argument list. Defined on PROCEDURE objects. Same precedence as .
- ::, . GDB and Modula-2 scope operators.

Warning: Set expressions and their operations are not yet supported, so GDB treats the use of the operator IN, or the use of operators +, -, *, /, =, , <>, #, <=, and >= on sets as an error.

15.4.5.2 Built-in Functions and Procedures

Modula-2 also makes available several built-in procedures and functions. In describing these, the following metavariables are used:

- a represents an ARRAY variable.
- c represents a CHAR constant or variable.
- *i* represents a variable or constant of integral type.
- m represents an identifier that belongs to a set. Generally used in the same function with the metavariable s. The type of s should be SET OF mtype (where mtype is the type of m).
- n represents a variable or constant of integral or floating-point type.
- r represents a variable or constant of floating-point type.
- t represents a type.
- v represents a variable.
- x represents a variable or constant of one of many types. See the explanation of the function for details.

All Modula-2 built-in procedures also return a result, described below.

- ABS(n) Returns the absolute value of n.
- CAP(c) If c is a lower case letter, it returns its upper case equivalent, otherwise it returns its argument.
- CHR(i) Returns the character whose ordinal value is i.
- DEC(v) Decrements the value in the variable v by one. Returns the new value.
- DEC(v, i) Decrements the value in the variable v by i. Returns the new value.
- EXCL(m,s)

Removes the element m from the set s. Returns the new set.

- FLOAT(i) Returns the floating point equivalent of the integer i.
- HIGH(a) Returns the index of the last member of a.
- INC(v) Increments the value in the variable v by one. Returns the new value.
- INC(v, i) Increments the value in the variable v by i. Returns the new value.
- INCL(m,s)

Adds the element m to the set s if it is not already there. Returns the new set.

- MAX(t) Returns the maximum value of the type t.
- MIN(t) Returns the minimum value of the type t.
- ODD(i) Returns boolean TRUE if i is an odd number.
- ORD(x) Returns the ordinal value of its argument. For example, the ordinal value of a character is its ASCII value (on machines supporting the ASCII character set). x must be of an ordered type, which include integral, character and enumerated types.

- SIZE(x) Returns the size of its argument. x can be a variable or a type.
- TRUNC(r) Returns the integral part of r.
- TSIZE(x) Returns the size of its argument. x can be a variable or a type.
- VAL(t,i) Returns the member of the type t whose ordinal value is i.

Warning: Sets and their operations are not yet supported, so GDB treats the use of procedures INCL and EXCL as an error.

15.4.5.3 Constants

GDB allows you to express the constants of Modula-2 in the following ways:

- Integer constants are simply a sequence of digits. When used in an expression, a constant is interpreted to be type-compatible with the rest of the expression. Hexadecimal integers are specified by a trailing 'H', and octal integers by a trailing 'B'.
- Floating point constants appear as a sequence of digits, followed by a decimal point and another sequence of digits. An optional exponent can then be specified, in the form 'E[+|-]nnn', where '[+|-]nnn' is the desired exponent. All of the digits of the floating point constant must be valid decimal (base 10) digits.
- Character constants consist of a single character enclosed by a pair of like quotes, either single (') or double ("). They may also be expressed by their ordinal value (their ASCII value, usually) followed by a 'C'.
- String constants consist of a sequence of characters enclosed by a pair of like quotes, either single (') or double ("). Escape sequences in the style of C are also allowed. See Section 15.4.1.2 [C and C++ Constants], page 153, for a brief explanation of escape sequences.
- Enumerated constants consist of an enumerated identifier.
- Boolean constants consist of the identifiers TRUE and FALSE.
- Pointer constants consist of integral values only.
- Set constants are not yet supported.

15.4.5.4 Modula-2 Types

Currently GDB can print the following data types in Modula-2 syntax: array types, record types, set types, pointer types, procedure types, enumerated types, subrange types and base types. You can also print the contents of variables declared using these type. This section gives a number of simple source code examples together with sample GDB sessions.

The first example contains the following section of code:

```
VAR

s: SET OF CHAR;

r: [20..40];

and you can request GDB to interrogate the type and value of r and s.

(gdb) print s
{'A'...'C', 'Z'}

(gdb) ptype s
SET OF CHAR

(gdb) print r
```

Note that at present you cannot interactively manipulate set expressions using the debugger.

The following example shows how you might declare an array in Modula-2 and how you can interact with GDB to print its type and contents:

```
VAR
s: ARRAY [-10..10] OF CHAR;
(gdb) ptype s
ARRAY [-10..10] OF CHAR
```

Note that the array handling is not yet complete and although the type is printed correctly, expression handling still assumes that all arrays have a lower bound of zero and not -10 as in the example above.

Here are some more type related Modula-2 examples:

```
TYPE
  colour = (blue, red, yellow, green);
  t = [blue..yellow];
VAR
  s: t;
BEGIN
  s := blue;
```

The GDB interaction shows how you can query the data type and value of a variable.

```
(gdb) print s
$1 = blue
(gdb) ptype t
type = [blue..yellow]
```

In this example a Modula-2 array is declared and its contents displayed. Observe that the contents are written in the same way as their C counterparts.

```
VAR
s: ARRAY [1..5] OF CARDINAL;
BEGIN
s[1] := 1;
(gdb) print s
$1 = {1, 0, 0, 0, 0}
(gdb) ptype s
type = ARRAY [1..5] OF CARDINAL
```

The Modula-2 language interface to GDB also understands pointer types as shown in this example:

```
s: POINTER TO ARRAY [1..5] OF CARDINAL;

BEGIN

NEW(s);

s^[1] := 1;
```

and you can request that GDB describes the type of s.

```
(gdb) ptype s
type = POINTER TO ARRAY [1..5] OF CARDINAL
```

GDB handles compound types as we can see in this example. Here we combine array types, record types, pointer types and subrange types:

```
TYPE
         foo = RECORD
                 f1: CARDINAL;
                  f2: CHAR;
                 f3: myarray;
               END ;
         myarray = ARRAY myrange OF CARDINAL ;
        myrange = [-2..2];
      VΔR
         s: POINTER TO ARRAY myrange OF foo;
and you can ask GDB to describe the type of s as shown below.
      (gdb) ptype s
      type = POINTER TO ARRAY [-2..2] OF foo = RECORD
         f1 : CARDINAL;
         f2 : CHAR;
         f3 : ARRAY [-2..2] OF CARDINAL;
      F.ND
```

15.4.5.5 Modula-2 Defaults

If type and range checking are set automatically by GDB, they both default to on whenever the working language changes to Modula-2. This happens regardless of whether you or GDB selected the working language.

If you allow GDB to set the language automatically, then entering code compiled from a file whose name ends with '.mod' sets the working language to Modula-2. See Section 15.1.3 [Having GDB Infer the Source Language], page 148, for further details.

15.4.5.6 Deviations from Standard Modula-2

A few changes have been made to make Modula-2 programs easier to debug. This is done primarily via loosening its type strictness:

- Unlike in standard Modula-2, pointer constants can be formed by integers. This allows you to modify pointer variables during debugging. (In standard Modula-2, the actual address contained in a pointer variable is hidden from you; it can only be modified through direct assignment to another pointer variable or expression that returned a pointer.)
- C escape sequences can be used in strings and characters to represent non-printable characters. GDB prints out strings with these escape sequences embedded. Single non-printable characters are printed using the 'CHR(nnn)' format.
- The assignment operator (:=) returns the value of its right-hand argument.
- All built-in procedures both modify and return their argument.

15.4.5.7 Modula-2 Type and Range Checks

Warning: in this release, GDB does not yet perform type or range checking. GDB considers two Modula-2 variables type equivalent if:

- They are of types that have been declared equivalent via a TYPE t1 = t2 statement
- They have been declared on the same line. (Note: This is true of the GNU Modula-2 compiler, but it may not be true of other compilers.)

As long as type checking is enabled, any attempt to combine variables whose types are not equivalent is an error.

Range checking is done on all mathematical operations, assignment, array index bounds, and all built-in functions and procedures.

15.4.5.8 The Scope Operators :: and .

There are a few subtle differences between the Modula-2 scope operator (.) and the GDB scope operator (::). The two have similar syntax:

```
module . id
scope :: id
```

where *scope* is the name of a module or a procedure, *module* the name of a module, and *id* is any declared identifier within your program, except another module.

Using the :: operator makes GDB search the scope specified by *scope* for the identifier *id*. If it is not found in the specified scope, then GDB searches all scopes enclosing the one specified by *scope*.

Using the . operator makes GDB search the current scope for the identifier specified by *id* that was imported from the definition module specified by *module*. With this operator, it is an error if the identifier *id* was not imported from definition module *module*, or if *id* is not an identifier in *module*.

15.4.5.9 GDB and Modula-2

Some GDB commands have little use when debugging Modula-2 programs. Five subcommands of set print and show print apply specifically to C and C++: 'vtbl', 'demangle', 'asm-demangle', 'object', and 'union'. The first four apply to C++, and the last to the C union type, which has no direct analogue in Modula-2.

The @ operator (see Section 10.1 [Expressions], page 93), while available with any language, is not useful with Modula-2. Its intent is to aid the debugging of dynamic arrays, which cannot be created in Modula-2 as they can in C or C++. However, because an address can be specified by an integral constant, the construct '{type}adrexp' is still useful.

In GDB scripts, the Modula-2 inequality operator # is interpreted as the beginning of a comment. Use <> instead.

15.4.6 Ada

The extensions made to gdb for Ada only support output from the gnu Ada (GNAT) compiler. Other Ada compilers are not currently supported, and attempting to debug executables produced by them is most likely to be difficult.

15.4.6.1 Introduction

The Ada mode of GDB supports a fairly large subset of Ada expression syntax, with some extensions. The philosophy behind the design of this subset is

• That GDB should provide basic literals and access to operations for arithmetic, dereferencing, field selection, indexing, and subprogram calls, leaving more sophisticated computations to subprograms written into the program (which therefore may be called from GDB).

- That type safety and strict adherence to Ada language restrictions are not particularly important to the GDB user.
- That brevity is important to the GDB user.

Thus, for brevity, the debugger acts as if all names declared in user-written packages are directly visible, even if they are not visible according to Ada rules, thus making it unnecessary to fully qualify most names with their packages, regardless of context. Where this causes ambiguity, GDB asks the user's intent.

The debugger will start in Ada mode if it detects an Ada main program. As for other languages, it will enter Ada mode when stopped in a program that was translated from an Ada source file.

While in Ada mode, you may use '--' for comments. This is useful mostly for documenting command files. The standard gdb comment ('#') still works at the beginning of a line in Ada mode, but not in the middle (to allow based literals).

The debugger supports limited overloading. Given a subprogram call in which the function symbol has multiple definitions, it will use the number of actual parameters and some information about their types to attempt to narrow the set of definitions. It also makes very limited use of context, preferring procedures to functions in the context of the call command, and functions to procedures elsewhere.

15.4.6.2 Omissions from Ada

Here are the notable omissions from the subset:

- Only a subset of the attributes are supported:
 - 'First, 'Last, and 'Length on array objects (not on types and subtypes).
 - 'Min and 'Max.
 - 'Pos and 'Val.
 - 'Tag.
 - 'Range on array objects (not subtypes), but only as the right operand of the membership (in) operator.
 - 'Access, 'Unchecked_Access, and 'Unrestricted_Access (a GNAT extension).
 - 'Address.
- The names in Characters.Latin_1 are not available and concatenation is not implemented. Thus, escape characters in strings are not currently available.
- Equality tests ('=' and '/=') on arrays test for bitwise equality of representations. They will generally work correctly for strings and arrays whose elements have integer or enumeration types. They may not work correctly for arrays whose element types have user-defined equality, for arrays of real values (in particular, IEEE-conformant floating point, because of negative zeroes and NaNs), and for arrays whose elements contain unused bits with indeterminate values.

- The other component-by-component array operations (and, or, xor, not, and relational tests other than equality) are not implemented.
- There is limited support for array and record aggregates. They are permitted only on the right sides of assignments, as in these examples:

```
(gdb) set An_Array := (1, 2, 3, 4, 5, 6)

(gdb) set An_Array := (1, others => 0)

(gdb) set An_Array := (0|4 => 1, 1..3 => 2, 5 => 6)

(gdb) set A_2D_Array := ((1, 2, 3), (4, 5, 6), (7, 8, 9))

(gdb) set A_Record := (1, "Peter", True);

(gdb) set A_Record := (Name => "Peter", Id => 1, Alive => True)
```

Changing a discriminant's value by assigning an aggregate has an undefined effect if that discriminant is used within the record. However, you can first modify discriminants by directly assigning to them (which normally would not be allowed in Ada), and then performing an aggregate assignment. For example, given a variable A_Rec declared to have a type such as:

```
type Rec (Len : Small_Integer := 0) is record
   Id : Integer;
   Vals : IntArray (1 .. Len);
end record;
```

you can assign a value with a different size of Vals with two assignments:

```
(gdb) set A_Rec.Len := 4
(gdb) set A_Rec := (Id => 42, Vals => (1, 2, 3, 4))
```

As this example also illustrates, GDB is very loose about the usual rules concerning aggregates. You may leave out some of the components of an array or record aggregate (such as the Len component in the assignment to A_Rec above); they will retain their original values upon assignment. You may freely use dynamic values as indices in component associations. You may even use overlapping or redundant component associations, although which component values are assigned in such cases is not defined.

- Calls to dispatching subprograms are not implemented.
- The overloading algorithm is much more limited (i.e., less selective) than that of real Ada. It makes only limited use of the context in which a subexpression appears to resolve its meaning, and it is much looser in its rules for allowing type matches. As a result, some function calls will be ambiguous, and the user will be asked to choose the proper resolution.
- The new operator is not implemented.
- Entry calls are not implemented.
- Aside from printing, arithmetic operations on the native VAX floating-point formats are not supported.
- It is not possible to slice a packed array.
- The names True and False, when not part of a qualified name, are interpreted as if implicitly prefixed by Standard, regardless of context. Should your program redefine these names in a package or procedure (at best a dubious practice), you will have to use fully qualified names to access their new definitions.

15.4.6.3 Additions to Ada

As it does for other languages, GDB makes certain generic extensions to Ada (see Section 10.1 [Expressions], page 93):

• If the expression E is a variable residing in memory (typically a local variable or array element) and N is a positive integer, then EQN displays the values of E and the N-1 adjacent variables following it in memory as an array. In Ada, this operator is generally not necessary, since its prime use is in displaying parts of an array, and slicing will usually do this in Ada. However, there are occasional uses when debugging programs in which certain debugging information has been optimized away.

- B::var means "the variable named var that appears in function or file B." When B is a file name, you must typically surround it in single quotes.
- The expression {type} addr means "the variable of type type that appears at address addr."
- A name starting with '\$' is a convenience variable (see Section 10.10 [Convenience Vars], page 109) or a machine register (see Section 10.11 [Registers], page 110).

In addition, GDB provides a few other shortcuts and outright additions specific to Ada:

• The assignment statement is allowed as an expression, returning its right-hand operand as its value. Thus, you may enter

```
(gdb) set x := y + 3
(gdb) print A(tmp := y + 1)
```

• The semicolon is allowed as an "operator," returning as its value the value of its right-hand operand. This allows, for example, complex conditional breaks:

```
(gdb) break f
(gdb) condition 1 (report(i); k += 1; A(k) > 100)
```

• Rather than use catenation and symbolic character names to introduce special characters into strings, one may instead use a special bracket notation, which is also used to print strings. A sequence of characters of the form '["XX"]' within a string or character literal denotes the (single) character whose numeric encoding is XX in hexadecimal. The sequence of characters '["""]' also denotes a single quotation mark in strings. For example,

```
"One line.["Oa"]Next line.["Oa"]"
```

contains an ASCII newline character (Ada. Characters. Latin_1.LF) after each period.

• The subtype used as a prefix for the attributes 'Pos, 'Min, and 'Max is optional (and is ignored in any case). For example, it is valid to write

```
(gdb) print 'max(x, y)
```

• When printing arrays, GDB uses positional notation when the array has a lower bound of 1, and uses a modified named notation otherwise. For example, a one-dimensional array of three integers with a lower bound of 3 might print as

```
(3 \Rightarrow 10, 17, 1)
```

That is, in contrast to valid Ada, only the first component has a => clause.

- You may abbreviate attributes in expressions with any unique, multi-character subsequence of their names (an exact match gets preference). For example, you may use a'len, a'gth, or a'lh in place of a'length.
- Since Ada is case-insensitive, the debugger normally maps identifiers you type to lower case. The GNAT compiler uses upper-case characters for some of its internal identifiers, which are normally of no interest to users. For the rare occasions when you actually have to look at them, enclose them in angle brackets to avoid the lower-case mapping. For example,

```
(gdb) print <JMPBUF_SAVE>[0]
```

• Printing an object of class-wide type or dereferencing an access-to-class-wide value will display all the components of the object's specific type (as indicated by its run-time tag). Likewise, component selection on such a value will operate on the specific type of the object.

15.4.6.4 Stopping at the Very Beginning

It is sometimes necessary to debug the program during elaboration, and before reaching the main procedure. As defined in the Ada Reference Manual, the elaboration code is invoked from a procedure called adainit. To run your program up to the beginning of elaboration, simply use the following two commands: tbreak adainit and run.

15.4.6.5 Extensions for Ada Tasks

Support for Ada tasks is analogous to that for threads (see Section 4.10 [Threads], page 35). GDB provides the following task-related commands:

info tasks

This command shows a list of current Ada tasks, as in the following example:

```
(gdb) info tasks

ID TID P-ID Pri State Name

1 8088000 0 15 Child Activation Wait main_task
2 80a4000 1 15 Accept Statement b
3 809a800 1 15 Child Activation Wait a

* 4 80ae800 3 15 Runnable c
```

In this listing, the asterisk before the last task indicates it to be the task currently being inspected.

ID Represents GDB's internal task number.

TID The Ada task ID.

P-ID The parent's task ID (GDB's internal task number).

Pri The base priority of the task.

State Current state of the task.

Unactivated

The task has been created but has not been activated. It cannot be executing.

Runnable The task is not blocked for any reason known to Ada. (It may be waiting for a mutex, though.) It is conceptually "executing" in normal mode.

Terminated

The task is terminated, in the sense of ARM 9.3 (5). Any dependents that were waiting on terminate alternatives have been awakened and have terminated themselves.

Child Activation Wait

The task is waiting for created tasks to complete activation.

Accept Statement

The task is waiting on an accept or selective wait statement.

Waiting on entry call

The task is waiting on an entry call.

Async Select Wait

The task is waiting to start the abortable part of an asynchronous select statement.

Delay Sleep

The task is waiting on a select statement with only a delay alternative open.

Child Termination Wait

The task is sleeping having completed a master within itself, and is waiting for the tasks dependent on that master to become terminated or waiting on a terminate Phase.

Wait Child in Term Alt

The task is sleeping waiting for tasks on terminate alternatives to finish terminating.

Accepting RV with taskno

The task is accepting a rendez-vous with the task taskno.

Name Name of the task in the program.

info task taskno

This command shows detailled informations on the specified task, as in the following example:

```
(gdb) info tasks

ID TID P-ID Pri State Name

1 8077880 0 15 Child Activation Wait main_task

* 2 807c468 1 15 Runnable task_1
(gdb) info task 2
Ada Task: 0x807c468
Name: task_1
Thread: 0x807f378
Parent: 1 (main_task)
Base Priority: 15
State: Runnable
```

task This command prints the ID of the current task.

```
(gdb) info tasks

ID TID P-ID Pri State Name

1 8077870 0 15 Child Activation Wait main_task

* 2 807c458 1 15 Runnable t

(gdb) task

[Current task is 2]
```

task taskno

This command is like the thread threadno command (see Section 4.10 [Threads], page 35). It switches the context of debugging from the current task to the given task.

break linespec task taskno break linespec task taskno if ...

These commands are like the break ... thread ... command (see Section 5.4 [Thread Stops], page 65). linespec specifies source lines, as described in Section 9.2 [Specify Location], page 84.

Use the qualifier 'task taskno' with a breakpoint command to specify that you only want GDB to stop the program when a particular Ada task reaches this breakpoint. taskno is one of the numeric task identifiers assigned by GDB, shown in the first column of the 'info tasks' display.

If you do not specify 'task taskno' when you set a breakpoint, the breakpoint applies to *all* tasks of your program.

You can use the task qualifier on conditional breakpoints as well; in this case, place 'task taskno' before the breakpoint condition (before the if).

For example,

```
(gdb) info tasks

ID TID P-ID Pri State Name

1 140022020 0 15 Child Activation Wait main_task
2 140045060 1 15 Accept/Select Wait t2
3 140044840 1 15 Runnable t1

* 4 140056040 1 15 Runnable t3
(gdb) b 15 task 2

Breakpoint 5 at 0x120044cb0: file test_task_debug.adb, line 15.
(gdb) cont
Continuing.
```

Debugging with GDB

```
task # 1 running
task # 2 running
Breakpoint 5, test_task_debug () at test_task_debug.adb:15
                flush;
(gdb) info tasks
          TID P-ID Pri State
 ID
                                             Name
  1 140022020 0 15 Child Activation Wait main_task
  2 140045060
                                             t2
                    15 Runnable
                1
  3 140044840
                    15 Runnable
                1
                                             t1
  4 140056040
                    15 Delay Sleep
                                             t3
               1
```

15.4.6.6 Tasking Support when Debugging Core Files

When inspecting a core file, as opposed to debugging a live program, tasking support may be limited or even unavailable, depending on the platform being used. For instance, on x86-linux, the list of tasks is available, but task switching is not supported. On Tru64, however, task switching will work as usual.

On certain platforms, including Tru64, the debugger needs to perform some memory writes in order to provide Ada tasking support. When inspecting a core file, this means that the core file must be opened with read-write privileges, using the command '"set write on"' (see Section 17.6 [Patching], page 185). Under these circumstances, you should make a backup copy of the core file before inspecting it with GDB.

15.4.6.7 Known Peculiarities of Ada Mode

Besides the omissions listed previously (see Section 15.4.6.2 [Omissions from Ada], page 166), we know of several problems with and limitations of Ada mode in GDB, some of which will be fixed with planned future releases of the debugger and the GNU Ada compiler.

- Currently, the debugger has insufficient information to determine whether certain pointers represent pointers to objects or the objects themselves. Thus, the user may have to tack an extra .all after an expression to get it printed properly.
- Static constants that the compiler chooses not to materialize as objects in storage are invisible to the debugger.
- Named parameter associations in function argument lists are ignored (the argument lists are treated as positional).
- Many useful library packages are currently invisible to the debugger.
- Fixed-point arithmetic, conversions, input, and output is carried out using floating-point arithmetic, and may give results that only approximate those on the host machine.
- The GNAT compiler never generates the prefix Standard for any of the standard symbols defined by the Ada language. GDB knows about this: it will strip the prefix from names when you use it, and will never look for a name you have so qualified among local symbols, nor match against symbols in other packages or subprograms. If you have defined entities anywhere in your program other than parameters and local variables whose simple names match names in Standard, GNAT's lack of qualification here can cause confusion. When this happens, you can usually resolve the confusion by qualifying the problematic names with package Standard explicitly.

Older versions of the compiler sometimes generate erroneous debugging information, resulting in the debugger incorrectly printing the value of affected entities. In some cases, the debugger is able to work around an issue automatically. In other cases, the debugger is able to work around the issue, but the work-around has to be specifically enabled.

set ada trust-PAD-over-XVS on

Configure GDB to strictly follow the GNAT encoding when computing the value of Ada entities, particularly when PAD and PAD___XVS types are involved (see ada/exp_dbug.ads in the GCC sources for a complete description of the encoding used by the GNAT compiler). This is the default.

set ada trust-PAD-over-XVS off

This is related to the encoding using by the GNAT compiler. If GDB sometimes prints the wrong value for certain entities, changing ada trust-PAD-over-XVS to off activates a work-around which may fix the issue. It is always safe to set ada trust-PAD-over-XVS to off, but this incurs a slight performance penalty, so it is recommended to leave this setting to on unless necessary.

15.5 Unsupported Languages

In addition to the other fully-supported programming languages, GDB also provides a pseudo-language, called minimal. It does not represent a real programming language, but provides a set of capabilities close to what the C or assembly languages provide. This should allow most simple operations to be performed while debugging an application that uses a language currently not supported by GDB.

If the language is set to auto, GDB will automatically select this language if the current frame corresponds to an unsupported language.

16 Examining the Symbol Table

The commands described in this chapter allow you to inquire about the symbols (names of variables, functions and types) defined in your program. This information is inherent in the text of your program and does not change as your program executes. GDB finds it in your program's symbol table, in the file indicated when you started GDB (see Section 2.1.1 [Choosing Files], page 12), or by one of the file-management commands (see Section 18.1 [Commands to Specify Files], page 187).

Occasionally, you may need to refer to symbols that contain unusual characters, which GDB ordinarily treats as word delimiters. The most frequent case is in referring to static variables in other source files (see Section 10.3 [Program Variables], page 95). File names are recorded in object files as debugging symbols, but GDB would ordinarily parse a typical file name, like 'foo.c', as the three words 'foo' '.' 'c'. To allow GDB to recognize 'foo.c' as a single symbol, enclose it in single quotes; for example,

```
p 'foo.c'::x
```

looks up the value of x in the scope of the file 'foo.c'.

```
set case-sensitive on
set case-sensitive off
set case-sensitive auto
```

Normally, when GDB looks up symbols, it matches their names with case sensitivity determined by the current source language. Occasionally, you may wish to control that. The command set case-sensitive lets you do that by specifying on for case-sensitive matches or off for case-insensitive ones. If you specify auto, case sensitivity is reset to the default suitable for the source language. The default is case-sensitive matches for all languages except for Fortran, for which the default is case-insensitive matches.

show case-sensitive

This command shows the current setting of case sensitivity for symbols lookups.

info address symbol

Describe where the data for *symbol* is stored. For a register variable, this says which register it is kept in. For a non-register local variable, this prints the stack-frame offset at which the variable is always stored.

Note the contrast with 'print &symbol', which does not work at all for a register variable, and for a stack local variable prints the exact address of the current instantiation of the variable.

info symbol addr

Print the name of a symbol which is stored at the address addr. If no symbol is stored exactly at addr, GDB prints the nearest symbol and an offset from it:

```
(gdb) info symbol 0x54320
_initialize_vx + 396 in section .text
```

This is the opposite of the info address command. You can use it to find out the name of a variable or a function given its address.

For dynamically linked executables, the name of executable or shared library containing the symbol is also printed:

```
(gdb) info symbol 0x400225
_start + 5 in section .text of /tmp/a.out
(gdb) info symbol 0x2aaaac2811cf
__read_nocancel + 6 in section .text of /usr/lib64/libc.so.6
```

whatis [arg]

Print the data type of arg, which can be either an expression or a data type. With no argument, print the data type of \$, the last value in the value history. If arg is an expression, it is not actually evaluated, and any side-effecting operations (such as assignments or function calls) inside it do not take place. If arg is a type name, it may be the name of a type or typedef, or for C code it may have the form 'class class-name', 'struct struct-tag', 'union union-tag' or 'enum enum-tag'. See Section 10.1 [Expressions], page 93.

ptype [arg]

ptype accepts the same arguments as whatis, but prints a detailed description of the type, instead of just the name of the type. See Section 10.1 [Expressions], page 93.

For example, for this variable declaration:

```
struct complex {double real; double imag;} v;
the two commands give this output:
    (gdb) whatis v
    type = struct complex
    (gdb) ptype v
    type = struct complex {
        double real;
        double imag;
}
```

As with whatis, using ptype without an argument refers to the type of \$, the last value in the value history.

Sometimes, programs use opaque data types or incomplete specifications of complex data structure. If the debug information included in the program does not allow GDB to display a full declaration of the data type, it will say '<incomplete type>'. For example, given these declarations:

\$1 = <incomplete type>

"Incomplete type" is C terminology for data types that are not completely specified.

```
info types regexp
info types
```

Print a brief description of all types whose names match the regular expression regexp (or all types in your program, if you supply no argument). Each complete typename is matched as though it were a complete line; thus, 'i type value' gives information on all types in your program whose names include the string value, but 'i type ^value\$' gives information only on types whose complete name is value.

This command differs from ptype in two ways: first, like whatis, it does not print a detailed description; second, it lists all source files where a type is defined.

info scope location

List all the variables local to a particular scope. This command accepts a location argument—a function name, a source line, or an address preceded by a '*', and prints all the variables local to the scope defined by that location. (See Section 9.2 [Specify Location], page 84, for details about supported forms of location.) For example:

```
(gdb) info scope command_line_handler

Scope for command_line_handler:

Symbol rl is an argument at stack/frame offset 8, length 4.

Symbol linebuffer is in static storage at address 0x150a18, length 4.

Symbol linelength is in static storage at address 0x150a1c, length 4.

Symbol p is a local variable in register $esi, length 4.

Symbol p1 is a local variable in register $ebx, length 4.

Symbol nline is a local variable in register $edx, length 4.

Symbol repeat is a local variable at frame offset -8, length 4.
```

This command is especially useful for determining what data to collect during a trace experiment, see Section 13.1.6 [Tracepoint Actions], page 132.

info source

Show information about the current source file—that is, the source file for the function containing the current point of execution:

- the name of the source file, and the directory containing it,
- the directory it was compiled in,
- its length, in lines,
- which programming language it is written in,
- whether the executable includes debugging information for that file, and if so, what format the information is in (e.g., STABS, Dwarf 2, etc.), and
- whether the debugging information includes information about preprocessor macros.

info sources

Print the names of all source files in your program for which there is debugging information, organized into two lists: files whose symbols have already been read, and files whose symbols will be read when needed.

info functions

Print the names and data types of all defined functions.

info functions regexp

Print the names and data types of all defined functions whose names contain a match for regular expression regexp. Thus, 'info fun step' finds all functions whose names include step; 'info fun step' finds those whose names start with step. If a function name contains characters that conflict with the regular expression language (e.g. 'operator*()'), they may be quoted with a backslash.

info variables

Print the names and data types of all variables that are defined outside of functions (i.e. excluding local variables).

info variables regexp

Print the names and data types of all variables (except for local variables) whose names contain a match for regular expression regexp.

info classes

info classes regexp

Display all Objective-C classes in your program, or (with the *regexp* argument) all those matching a particular regular expression.

info selectors

info selectors regexp

Display all Objective-C selectors in your program, or (with the *regexp* argument) all those matching a particular regular expression.

Some systems allow individual object files that make up your program to be replaced without stopping and restarting your program. For example, in Vx-Works you can simply recompile a defective object file and keep on running. If you are running on one of these systems, you can allow GDB to reload the symbols for automatically relinked modules:

set symbol-reloading on

Replace symbol definitions for the corresponding source file when an object file with a particular name is seen again.

set symbol-reloading off

Do not replace symbol definitions when encountering object files of the same name more than once. This is the default state; if you are not running on a system that permits automatic relinking of modules, you should leave symbol-reloading off, since otherwise GDB may discard symbols when linking large programs, that may contain several modules (from different directories or libraries) with the same name.

show symbol-reloading

Show the current on or off setting.

set opaque-type-resolution on

Tell GDB to resolve opaque types. An opaque type is a type declared as a pointer to a struct, class, or union—for example, struct MyType *—that is used in one source file although the full declaration of struct MyType is in another source file. The default is on.

A change in the setting of this subcommand will not take effect until the next time symbols for a file are loaded.

set opaque-type-resolution off

Tell GDB not to resolve opaque types. In this case, the type is printed as follows: {<no data fields>}

```
show opaque-type-resolution
```

Show whether opaque types are resolved or not.

```
maint print symbols filename
maint print psymbols filename
maint print msymbols filename
```

Write a dump of debugging symbol data into the file filename. These commands are used to debug the GDB symbol-reading code. Only symbols with debugging data are included. If you use 'maint print symbols', GDB includes all the symbols for which it has already collected full details: that is, filename reflects symbols for only those files whose symbols GDB has read. You can use the command info sources to find out which files these are. If you use 'maint print psymbols' instead, the dump shows information about symbols that GDB only knows partially—that is, symbols defined in files that GDB has skimmed, but not yet read completely. Finally, 'maint print msymbols' dumps just the minimal symbol information required for each object file from which GDB has read some symbols. See Section 18.1 [Commands to Specify Files], page 187, for a discussion of how GDB reads symbols (in the description of symbol-file).

```
maint info symtabs [ regexp ]
maint info psymtabs [ regexp ]
```

List the struct symtab or struct partial_symtab structures whose names match regexp. If regexp is not given, list them all. The output includes expressions which you can copy into a GDB debugging this one to examine a particular structure in more detail. For example:

```
(gdb) maint info psymtabs dwarf2read
{ objfile /home/gnu/build/gdb/gdb
  ((struct objfile *) 0x82e69d0)
  { psymtab /home/gnu/src/gdb/dwarf2read.c
        ((struct partial_symtab *) 0x8474b10)
        readin no
        fullname (null)
        text addresses 0x814d3c8 -- 0x8158074
        globals (* (struct partial_symbol **) 0x8507a08 @ 9)
        statics (* (struct partial_symbol **) 0x40e95b78 @ 2882)
        dependencies (none)
    }
}
(gdb) maint info symtabs
(gdb)
```

We see that there is one partial symbol table whose filename contains the string 'dwarf2read', belonging to the 'gdb' executable; and we see that GDB has not read in any symtabs yet at all. If we set a breakpoint on a function, that will cause GDB to read the symtab for the compilation unit containing that function:

Debugging with GDB

```
dirname (null)
fullname (null)
blockvector ((struct blockvector *) 0x86c1bd0) (primary)
linetable ((struct linetable *) 0x8370fa0)
debugformat DWARF 2
}
{
gdb)
```

17 Altering Execution

Once you think you have found an error in your program, you might want to find out for certain whether correcting the apparent error would lead to correct results in the rest of the run. You can find the answer by experiment, using the GDB features for altering execution of the program.

For example, you can store new values into variables or memory locations, give your program a signal, restart it at a different address, or even return prematurely from a function.

17.1 Assignment to Variables

To alter the value of a variable, evaluate an assignment expression. See Section 10.1 [Expressions], page 93. For example,

```
print x=4
```

stores the value 4 into the variable x, and then prints the value of the assignment expression (which is 4). See Chapter 15 [Using GDB with Different Languages], page 147, for more information on operators in supported languages.

If you are not interested in seeing the value of the assignment, use the set command instead of the print command. set is really the same as print except that the expression's value is not printed and is not put in the value history (see Section 10.9 [Value History], page 108). The expression is evaluated only for its effects.

If the beginning of the argument string of the set command appears identical to a set subcommand, use the set variable command instead of just set. This command is identical to set except for its lack of subcommands. For example, if your program has a variable width, you get an error if you try to set a new value with just 'set width=13', because GDB has the command set width:

```
(gdb) whatis width
type = double
(gdb) p width
$4 = 13
(gdb) set width=47
Invalid syntax in expression.
```

The invalid expression, of course, is '=47'. In order to actually set the program's variable width, use

```
(gdb) set var width=47
```

Because the set command has many subcommands that can conflict with the names of program variables, it is a good idea to use the set variable command instead of just set. For example, if your program has a variable g, you run into problems if you try to set a new value with just 'set g=4', because GDB has the command set gnutarget, abbreviated set g:

The program variable g did not change, and you silently set the gnutarget to an invalid value. In order to set the variable g, use

```
(gdb) set var g=4
```

GDB allows more implicit conversions in assignments than C; you can freely store an integer value into a pointer variable or vice versa, and you can convert any structure to any other structure that is the same length or shorter.

To store values into arbitrary places in memory, use the '{...}' construct to generate a value of specified type at a specified address (see Section 10.1 [Expressions], page 93). For example, {int}0x83040 refers to memory location 0x83040 as an integer (which implies a certain size and representation in memory), and

```
set {int}0x83040 = 4
```

stores the value 4 into that memory location.

17.2 Continuing at a Different Address

Ordinarily, when you continue your program, you do so at the place where it stopped, with the continue command. You can instead continue at an address of your own choosing, with the following commands:

```
jump linespec
jump location
```

Resume execution at line *linespec* or at address given by *location*. Execution stops again immediately if there is a breakpoint there. See Section 9.2 [Specify Location], page 84, for a description of the different forms of *linespec* and *location*. It is common practice to use the tbreak command in conjunction with jump. See Section 5.1.1 [Setting Breakpoints], page 44.

The jump command does not change the current stack frame, or the stack pointer, or the contents of any memory location or any register other than the program counter. If line *linespec* is in a different function from the one currently executing, the results may be bizarre if the two functions expect different patterns of arguments or of local variables. For this reason, the jump command requests confirmation if the specified line is not in the function currently executing. However, even bizarre results are predictable if you are well acquainted with the machine-language code of your program.

On many systems, you can get much the same effect as the jump command by storing a new value into the register \$pc. The difference is that this does not start your program running; it only changes the address of where it will run when you continue. For example,

```
set $pc = 0x485
```

makes the next continue command or stepping command execute at address 0x485, rather than at the address where your program stopped. See Section 5.2 [Continuing and Stepping], page 60.

The most common occasion to use the jump command is to back up—perhaps with more breakpoints set—over a portion of a program that has already executed, in order to examine its execution in more detail.

17.3 Giving your Program a Signal

signal signal

Resume execution where your program stopped, but immediately give it the signal signal. signal can be the name or the number of a signal. For example, on many systems signal 2 and signal SIGINT are both ways of sending an interrupt signal.

Alternatively, if *signal* is zero, continue execution without giving a signal. This is useful when your program stopped on account of a signal and would ordinary see the signal when resumed with the **continue** command; 'signal 0' causes it to resume without a signal.

signal does not repeat when you press RET a second time after executing the command.

Invoking the signal command is not the same as invoking the kill utility from the shell. Sending a signal with kill causes GDB to decide what to do with the signal depending on the signal handling tables (see Section 5.3 [Signals], page 63). The signal command passes the signal directly to your program.

17.4 Returning from a Function

return

return expression

You can cancel execution of a function call with the **return** command. If you give an *expression* argument, its value is used as the function's return value.

When you use return, GDB discards the selected stack frame (and all frames within it). You can think of this as making the discarded frame return prematurely. If you wish to specify a value to be returned, give that value as the argument to return.

This pops the selected stack frame (see Section 8.3 [Selecting a Frame], page 80), and any other frames inside of it, leaving its caller as the innermost remaining frame. That frame becomes selected. The specified value is stored in the registers used for returning values of functions.

The return command does not resume execution; it leaves the program stopped in the state that would exist if the function had just returned. In contrast, the finish command

(see Section 5.2 [Continuing and Stepping], page 60) resumes execution until the selected stack frame returns naturally.

GDB needs to know how the expression argument should be set for the inferior. The concrete registers assignment depends on the OS ABI and the type being returned by the selected stack frame. For example it is common for OS ABI to return floating point values in FPU registers while integer values in CPU registers. Still some ABIs return even floating point values in CPU registers. Larger integer widths (such as long long int) also have specific placement rules. GDB already knows the OS ABI from its current target so it needs to find out also the type being returned to make the assignment into the right register(s).

Normally, the selected stack frame has debug info. GDB will always use the debug info instead of the implicit type of expression when the debug info is available. For example, if you type return -1, and the function in the current stack frame is declared to return a long long int, GDB transparently converts the implicit int value of -1 into a long long int:

```
Breakpoint 1, func () at gdb.base/return-nodebug.c:29
29 return 31;
(gdb) return -1
Make func return now? (y or n) y
#0 0x004004f6 in main () at gdb.base/return-nodebug.c:43
43 printf ("result=%lld\n", func ());
(gdb)
```

However, if the selected stack frame does not have a debug info, e.g., if the function was compiled without debug info, GDB has to find out the type to return from user. Specifying a different type by mistake may set the value in different inferior registers than the caller code expects. For example, typing return -1 with its implicit type int would set only a part of a long long int result for a debug info less function (on 32-bit architectures). Therefore the user is required to specify the return type by an appropriate cast explicitly:

```
Breakpoint 2, 0x0040050b in func () (gdb) return -1
Return value type not available for selected stack frame.
Please use an explicit cast of the value to return.
(gdb) return (long long int) -1
Make selected stack frame return now? (y or n) y
#0 0x00400526 in main ()
(gdb)
```

17.5 Calling Program Functions

print expr

Evaluate the expression expr and display the resulting value. expr may include calls to functions in the program being debugged.

call expr

Evaluate the expression expr without displaying void returned values.

You can use this variant of the print command if you want to execute a function from your program that does not return anything (a.k.a. a void function), but without cluttering the output with void returned values that GDB will otherwise print. If the result is not void, it is printed and saved in the value history.

It is possible for the function you call via the print or call command to generate a signal (e.g., if there's a bug in the function, or if you passed it incorrect arguments). What happens in that case is controlled by the set unwindonsignal command.

Similarly, with a C++ program it is possible for the function you call via the print or call command to generate an exception that is not handled due to the constraints of the dummy frame. In this case, any exception that is raised in the frame, but has an out-of-frame exception handler will not be found. GDB builds a dummy-frame for the inferior function call, and the unwinder cannot seek for exception handlers outside of this dummy-frame. What happens in that case is controlled by the set unwind-on-terminating-exception command.

set unwindonsignal

Set unwinding of the stack if a signal is received while in a function that GDB called in the program being debugged. If set to on, GDB unwinds the stack it created for the call and restores the context to what it was before the call. If set to off (the default), GDB stops in the frame where the signal was received.

show unwindonsignal

Show the current setting of stack unwinding in the functions called by GDB.

set unwind-on-terminating-exception

Set unwinding of the stack if a C++ exception is raised, but left unhandled while in a function that GDB called in the program being debugged. If set to on (the default), GDB unwinds the stack it created for the call and restores the context to what it was before the call. If set to off, GDB the exception is delivered to the default C++ exception handler and the inferior terminated.

show unwind-on-terminating-exception

Show the current setting of stack unwinding in the functions called by GDB.

Sometimes, a function you wish to call is actually a weak alias for another function. In such case, GDB might not pick up the type information, including the types of the function arguments, which causes GDB to call the inferior function incorrectly. As a result, the called function will function erroneously and may even crash. A solution to that is to use the name of the aliased function instead.

17.6 Patching Programs

By default, GDB opens the file containing your program's executable code (or the corefile) read-only. This prevents accidental alterations to machine code; but it also prevents you from intentionally patching your program's binary.

If you'd like to be able to patch the binary, you can specify that explicitly with the set write command. For example, you might want to turn on internal debugging flags, or even to make emergency repairs.

set write on set write off

If you specify 'set write on', GDB opens executable and core files for both reading and writing; if you specify set write off (the default), GDB opens them read-only.

If you have already loaded a file, you must load it again (using the exec-file or core-file command) after changing set write, for your new setting to take effect.

show write

Display whether executable files and core files are opened for writing as well as reading.

18 GDB Files

GDB needs to know the file name of the program to be debugged, both in order to read its symbol table and in order to start your program. To debug a core dump of a previous run, you must also tell GDB the name of the core dump file.

18.1 Commands to Specify Files

You may want to specify executable and core dump file names. The usual way to do this is at start-up time, using the arguments to GDB's start-up commands (see Chapter 2 [Getting In and Out of GDB], page 11).

Occasionally it is necessary to change to a different file during a GDB session. Or you may run GDB and forget to specify a file you want to use. Or you are debugging a remote target via gdbserver (see Section 20.3 [Using the gdbserver Program], page 205). In these situations the GDB commands to specify new files are useful.

file filename

Use filename as the program to be debugged. It is read for its symbols and for the contents of pure memory. It is also the program executed when you use the run command. If you do not specify a directory and the file is not found in the gdb working directory, gdb uses the environment variable PATH as a list of directories to search, just as the shell does when looking for a program to run. You can change the value of this variable, for both gdb and your program, using the path command.

You can load unlinked object '.o' files into GDB using the file command. You will not be able to "run" an object file, but you can disassemble functions and inspect variables. Also, if the underlying BFD functionality supports it, you could use <code>gdb-write</code> to patch object files using this technique. Note that GDB can neither interpret nor modify relocations in this case, so branches and some initialized variables will appear to go to the wrong place. But this feature is still handy from time to time.

file with no argument makes GDB discard any information it has on both executable file and the symbol table.

exec-file [filename]

Specify that the program to be run (but not the symbol table) is found in *file-name*. GDB searches the environment variable PATH if necessary to locate your program. Omitting *filename* means to discard information on the executable file.

symbol-file [filename]

Read symbol table information from file *filename*. PATH is searched when necessary. Use the **file** command to get both symbol table and program to run from the same file.

symbol-file with no argument clears out GDB information on your program's symbol table.

The symbol-file command causes GDB to forget the contents of some breakpoints and auto-display expressions. This is because they may contain pointers

to the internal data recording symbols and data types, which are part of the old symbol table data being discarded inside GDB.

symbol-file does not repeat if you press RET again after executing it once.

When GDB is configured for a particular environment, it understands debugging information in whatever format is the standard generated for that environment; you may use either a GNU compiler, or other compilers that adhere to the local conventions. Best results are usually obtained from GNU compilers; for example, using GCC you can generate debugging information for optimized code.

For most kinds of object files, with the exception of old SVR3 systems using COFF, the symbol-file command does not normally read the symbol table in full right away. Instead, it scans the symbol table quickly to find which source files and which symbols are present. The details are read later, one source file at a time, as they are needed.

The purpose of this two-stage reading strategy is to make GDB start up faster. For the most part, it is invisible except for occasional pauses while the symbol table details for a particular source file are being read. (The set verbose command can turn these pauses into messages if desired. See Section 22.7 [Optional Warnings and Messages], page 249.)

We have not implemented the two-stage strategy for COFF yet. When the symbol table is stored in COFF format, symbol-file reads the symbol table data in full right away. Note that "stabs-in-COFF" still does the two-stage strategy, since the debug info is actually in stabs format.

```
symbol-file [ -readnow ] filename
file [ -readnow ] filename
```

You can override the GDB two-stage strategy for reading symbol tables by using the '-readnow' option with any of the commands that load symbol table information, if you want to be sure GDB has the entire symbol table available.

```
core-file [filename]
```

core

Specify the whereabouts of a core dump file to be used as the "contents of memory". Traditionally, core files contain only some parts of the address space of the process that generated them; GDB can access the executable file itself for other parts.

core-file with no argument specifies that no core file is to be used.

Note that the core file is ignored when your program is actually running under GDB. So, if you have been running your program and you wish to debug a core file instead, you must kill the subprocess in which the program is running. To do this, use the kill command (see Section 4.8 [Killing the Child Process], page 32).

```
add-symbol-file filename address
add-symbol-file filename address [ -readnow ]
add-symbol-file filename -ssection address ...
```

The add-symbol-file command reads additional symbol table information from the file *filename*. You would use this command when *filename* has been dynamically loaded (by some other means) into the program that is running.

address should be the memory address at which the file has been loaded; GDB cannot figure this out for itself. You can additionally specify an arbitrary number of '-ssection address' pairs, to give an explicit section name and base address for that section. You can specify any address as an expression.

The symbol table of the file *filename* is added to the symbol table originally read with the <code>symbol-file</code> command. You can use the <code>add-symbol-file</code> command any number of times; the new symbol data thus read keeps adding to the old. To discard all old symbol data instead, use the <code>symbol-file</code> command without any arguments.

Although filename is typically a shared library file, an executable file, or some other object file which has been fully relocated for loading into a process, you can also load symbolic information from relocatable '.o' files, as long as:

- the file's symbolic information refers only to linker symbols defined in that file, not to symbols defined by other object files,
- every section the file's symbolic information refers to has actually been loaded into the inferior, as it appears in the file, and
- you can determine the address at which every section was loaded, and provide these to the add-symbol-file command.

Some embedded operating systems, like Sun Chorus and VxWorks, can load relocatable files into an already running program; such systems typically make the requirements above easy to meet. However, it's important to recognize that many native systems use complex link procedures (.linkonce section factoring and C++ constructor table assembly, for example) that make the requirements difficult to meet. In general, one cannot assume that using add-symbol-file to read a relocatable object file's symbolic information will have the same effect as linking the relocatable object file into the program in the normal way.

add-symbol-file does not repeat if you press RET after using it.

add-symbol-file-from-memory address

Load symbols from the given address in a dynamically loaded object file whose image is mapped directly into the inferior's memory. For example, the Linux kernel maps a syscall DSO into each process's address space; this DSO provides kernel-specific code for some system calls. The argument can be any expression whose evaluation yields the address of the file's shared object file header. For this command to work, you must have used symbol-file or exec-file commands in advance.

add-shared-symbol-files library-file assf library-file

The add-shared-symbol-files command can currently be used only in the Cygwin build of GDB on MS-Windows OS, where it is an alias for the dll-symbols command (see Section 21.1.5 [Cygwin Native], page 221). GDB automatically looks for shared libraries, however if GDB does not find yours, you can invoke add-shared-symbol-files. It takes one argument: the shared library's file name. assf is a shorthand alias for add-shared-symbol-files.

section section addr

The section command changes the base address of the named section of the exec file to addr. This can be used if the exec file does not contain section addresses, (such as in the a.out format), or when the addresses specified in the file itself are wrong. Each section must be changed separately. The info files command, described below, lists all the sections and their addresses.

info files info target

info files and info target are synonymous; both print the current target (see Chapter 19 [Specifying a Debugging Target], page 199), including the names of the executable and core dump files currently in use by GDB, and the files from which symbols were loaded. The command help target lists all possible targets rather than current ones.

maint info sections

Another command that can give you extra information about program sections is maint info sections. In addition to the section information displayed by info files, this command displays the flags and file offset of each section in the executable and core dump files. In addition, maint info sections provides the following command options (which may be arbitrarily combined):

ALLOBJ Display sections for all loaded object files, including shared libraries.

sections Display info only for named sections.

section-flags

Display info only for sections for which section-flags are true. The section flags that GDB currently knows about are:

ALLOC Section will have space allocated in the process when loaded. Set for all sections except those containing debug information.

LOAD Section will be loaded from the file into the child process memory. Set for pre-initialized code and data, clear for .bss sections.

RELOC Section needs to be relocated before loading.

READONLY Section cannot be modified by the child process.

CODE Section contains executable code only.

DATA Section contains data only (no executable code).

ROM Section will reside in ROM.

CONSTRUCTOR.

Section contains data for constructor/destructor lists.

HAS_CONTENTS

Section is not empty.

NEVER_LOAD

An instruction to the linker to not output the section.

COFF_SHARED_LIBRARY

A notification to the linker that the section contains COFF shared library information.

IS_COMMON

Section contains common symbols.

set trust-readonly-sections on

Tell gdb that readonly sections in your object file really are read-only (i.e. that their contents will not change). In that case, gdb can fetch values from these sections out of the object file, rather than from the target program. For some targets (notably embedded ones), this can be a significant enhancement to debugging performance.

The default is off.

set trust-readonly-sections off

Tell gdb not to trust readonly sections. This means that the contents of the section might change while the program is running, and must therefore be fetched from the target when needed.

show trust-readonly-sections

Show the current setting of trusting readonly sections.

All file-specifying commands allow both absolute and relative file names as arguments. GDB always converts the file name to an absolute file name and remembers it that way.

GDB supports GNU/Linux, MS-Windows, HP-UX, SunOS, SVr4, Irix, and IBM RS/6000 AIX shared libraries.

On MS-Windows GDB must be linked with the Expat library to support shared libraries. See [Expat], page 391.

GDB automatically loads symbol definitions from shared libraries when you use the run command, or when you examine a core file. (Before you issue the run command, GDB does not understand references to a function in a shared library, however—unless you are debugging a core file).

On HP-UX, if the program loads a library explicitly, GDB automatically loads the symbols at the time of the shl_load call.

There are times, however, when you may wish to not automatically load symbol definitions from shared libraries, such as when they are particularly large or there are many of them.

To control the automatic loading of shared library symbols, use the commands:

set auto-solib-add mode

If mode is on, symbols from all shared object libraries will be loaded automatically when the inferior begins execution, you attach to an independently started inferior, or when the dynamic linker informs GDB that a new library has been loaded. If mode is off, symbols must be loaded manually, using the sharedlibrary command. The default value is on.

If your program uses lots of shared libraries with debug info that takes large amounts of memory, you can decrease the GDB memory footprint by preventing it from automatically loading the symbols from shared libraries. To that end, type <code>set auto-solib-add off</code> before running the inferior, then load each library whose debug symbols you do need with <code>sharedlibrary regexp</code>, where <code>regexp</code> is a regular expression that matches the libraries whose symbols you want to be loaded.

show auto-solib-add

Display the current autoloading mode.

To explicitly load shared library symbols, use the sharedlibrary command:

info share regex

info sharedlibrary regex

Print the names of the shared libraries which are currently loaded that match regex. If regex is omitted then print all shared libraries that are loaded.

sharedlibrary regex

share regex

Load shared object library symbols for files matching a Unix regular expression. As with files loaded automatically, it only loads shared libraries required by your program for a core file or after typing run. If regex is omitted all shared libraries required by your program are loaded.

nosharedlibrary

Unload all shared object library symbols. This discards all symbols that have been loaded from all shared libraries. Symbols from shared libraries that were loaded by explicit user requests are not discarded.

Sometimes you may wish that GDB stops and gives you control when any of shared library events happen. Use the set stop-on-solib-events command for this:

set stop-on-solib-events

This command controls whether GDB should give you control when the dynamic linker notifies it about some shared library event. The most common event of interest is loading or unloading of a new shared library.

show stop-on-solib-events

Show whether gdb stops and gives you control when shared library events happen.

Shared libraries are also supported in many cross or remote debugging configurations. GDB needs to have access to the target's libraries; this can be accomplished either by providing copies of the libraries on the host system, or by asking GDB to automatically retrieve the libraries from the target. If copies of the target libraries are provided, they need to be the same as the target libraries, although the copies on the target can be stripped as long as the copies on the host are not.

For remote debugging, you need to tell GDB where the target libraries are, so that it can load the correct copies—otherwise, it may try to load the host's libraries. GDB has two variables to specify the search directories for target libraries.

set sysroot path

Use path as the system root for the program being debugged. Any absolute shared library paths will be prefixed with path; many runtime loaders store the absolute paths to the shared library in the target program's memory. If you use set sysroot to find shared libraries, they need to be laid out in the same way that they are on the target, with e.g. a '/lib' and '/usr/lib' hierarchy under path.

If path starts with the sequence 'remote:', GDB will retrieve the target libraries from the remote system. This is only supported when using a remote target that supports the remote get command (see Section 20.2 [Sending files to a remote system], page 205). The part of path following the initial 'remote:' (if present) is used as system root prefix on the remote file system.¹

The set solib-absolute-prefix command is an alias for set sysroot.

You can set the default system root by using the configure-time '--with-sysroot' option. If the system root is inside GDB's configured binary prefix (set with '--prefix' or '--exec-prefix'), then the default system root will be updated automatically if the installed GDB is moved to a new location.

show sysroot

Display the current shared library prefix.

set solib-search-path path

If this variable is set, path is a colon-separated list of directories to search for shared libraries. 'solib-search-path' is used after 'sysroot' fails to locate the library, or if the path to the library is relative instead of absolute. If you want to use 'solib-search-path' instead of 'sysroot', be sure to set 'sysroot' to a nonexistent directory to prevent GDB from finding your host's libraries. 'sysroot' is preferred; setting it to a nonexistent directory may interfere with automatic loading of shared library symbols.

show solib-search-path

Display the current shared library search path.

18.2 Debugging Information in Separate Files

GDB allows you to put a program's debugging information in a file separate from the executable itself, in a way that allows GDB to find and load the debugging information automatically. Since debugging information can be very large—sometimes larger than the executable code itself—some systems distribute debugging information for their executables in separate files, which users can install only when they need to debug a problem.

GDB supports two ways of specifying the separate debug info file:

• The executable contains a *debug link* that specifies the name of the separate debug info file. The separate debug file's name is usually 'executable.debug', where executable is the name of the corresponding executable file without leading directories (e.g., 'ls.debug' for '/usr/bin/ls'). In addition, the debug link specifies a 32-bit

¹ If you want to specify a local system root using a directory that happens to be named 'remote:', you need to use some equivalent variant of the name like './remote:'.

Cyclic Redundancy Check (CRC) checksum for the debug file, which GDB uses to validate that the executable and the debug file came from the same build.

• The executable contains a build ID, a unique bit string that is also present in the corresponding debug info file. (This is supported only on some operating systems, notably those which use the ELF format for binary files and the GNU Binutils.) For more details about this feature, see the description of the '--build-id' command-line option in Section "Command Line Options" in The GNU Linker. The debug info file's name is not specified explicitly by the build ID, but can be computed from the build ID, see below.

Depending on the way the debug info file is specified, GDB uses two different methods of looking for the debug file:

- For the "debug link" method, GDB looks up the named file in the directory of the executable file, then in a subdirectory of that directory named '.debug', and finally under the global debug directory, in a subdirectory whose name is identical to the leading directories of the executable's absolute file name.
- For the "build ID" method, GDB looks in the '.build-id' subdirectory of the global debug directory for a file named 'nn/nnnnnnn.debug', where nn are the first 2 hex characters of the build ID bit string, and nnnnnnnn are the rest of the bit string. (Real build ID strings are 32 or more hex characters, not 10.)

So, for example, suppose you ask GDB to debug '/usr/bin/ls', which has a debug link that specifies the file 'ls.debug', and a build ID whose value in hex is abcdef1234. If the global debug directory is '/usr/lib/debug', then GDB will look for the following debug information files, in the indicated order:

- '/usr/lib/debug/.build-id/ab/cdef1234.debug'
- '/usr/bin/ls.debug'
- '/usr/bin/.debug/ls.debug'
- '/usr/lib/debug/usr/bin/ls.debug'.

You can set the global debugging info directory's name, and view the name GDB is currently using.

set debug-file-directory directories

Set the directories which GDB searches for separate debugging information files to directory. Multiple directory components can be set concatenating them by a directory separator.

show debug-file-directory

Show the directories GDB searches for separate debugging information files.

A debug link is a special section of the executable file named .gnu_debuglink. The section must contain:

- A filename, with any leading directory components removed, followed by a zero byte,
- zero to three bytes of padding, as needed to reach the next four-byte boundary within the section, and
- a four-byte CRC checksum, stored in the same endianness used for the executable file itself. The checksum is computed on the debugging information file's full contents by the function given below, passing zero as the *crc* argument.

Any executable file format can carry a debug link, as long as it can contain a section named .gnu_debuglink with the contents described above.

The build ID is a special section in the executable file (and in other ELF binary files that GDB may consider). This section is often named .note.gnu.build-id, but that name is not mandatory. It contains unique identification for the built files—the ID remains the same across multiple builds of the same build tree. The default algorithm SHA1 produces 160 bits (40 hexadecimal characters) of the content for the build ID string. The same section with an identical value is present in the original built binary with symbols, in its stripped variant, and in the separate debugging information file.

The debugging information file itself should be an ordinary executable, containing a full set of linker symbols, sections, and debugging information. The sections of the debugging information file should have the same names, addresses, and sizes as the original file, but they need not contain any data—much like a .bss section in an ordinary executable.

The GNU binary utilities (Binutils) package includes the 'objcopy' utility that can produce the separated executable / debugging information file pairs using the following commands:

```
objcopy --only-keep-debug foo foo.debug strip -g foo
```

These commands remove the debugging information from the executable file 'foo' and place it in the file 'foo.debug'. You can use the first, second or both methods to link the two files:

• The debug link method needs the following additional command to also leave behind a debug link in 'foo':

```
objcopy --add-gnu-debuglink=foo.debug foo
```

Ulrich Drepper's 'elfutils' package, starting with version 0.53, contains a version of the strip command such that the command strip foo -f foo.debug has the same functionality as the two objcopy commands and the ln -s command above, together.

• Build ID gets embedded into the main executable using 1d --build-id or the GCC counterpart gcc -W1,--build-id. Build ID support plus compatibility fixes for debug files separation are present in GNU binary utilities (Binutils) package since version 2.18.

The CRC used in .gnu_debuglink is the CRC-32 defined in IEEE 802.3 using the polynomial:

```
x^{32} + x^{26} + x^{23} + x^{22} + x^{16} + x^{12} + x^{11} 
+ x^{10} + x^{8} + x^{7} + x^{5} + x^{4} + x^{2} + x + 1
```

The function is computed byte at a time, taking the least significant bit of each byte first. The initial pattern <code>Oxffffffff</code> is used, to ensure leading zeros affect the CRC and the final result is inverted to ensure trailing zeros also affect the CRC.

Note: This is the same CRC polynomial as used in handling the Remote Serial Protocol qCRC packet (see Appendix D [GDB Remote Serial Protocol], page 403). However in the case of the Remote Serial Protocol, the CRC is computed most significant bit first, and the result is not inverted, so trailing zeros have no effect on the CRC value.

To complete the description, we show below the code of the function which produces the CRC used in .gnu_debuglink. Inverting the initially supplied crc argument means that an initial call to this function passing in zero will start computing the CRC using Oxffffffff.

```
unsigned long
gnu_debuglink_crc32 (unsigned long crc,
                     unsigned char *buf, size_t len)
 static const unsigned long crc32_table[256] =
     0x00000000, 0x77073096, 0xee0e612c, 0x990951ba, 0x076dc419,
      0x706af48f, 0xe963a535, 0x9e6495a3, 0x0edb8832, 0x79dcb8a4,
      OxeOd5e91e, 0x97d2d988, 0x09b64c2b, 0x7eb17cbd, 0xe7b82d07,
      0x90bf1d91, 0x1db71064, 0x6ab020f2, 0xf3b97148, 0x84be41de,
      0x1adad47d, 0x6ddde4eb, 0xf4d4b551, 0x83d385c7, 0x136c9856,
      0x646ba8c0, 0xfd62f97a, 0x8a65c9ec, 0x14015c4f, 0x63066cd9,
      Oxfa0f3d63, 0x8d080df5, 0x3b6e20c8, 0x4c69105e, 0xd56041e4,
      0xa2677172, 0x3c03e4d1, 0x4b04d447, 0xd20d85fd, 0xa50ab56b,
      0x35b5a8fa, 0x42b2986c, 0xdbbbc9d6, 0xacbcf940, 0x32d86ce3,
      0x45df5c75, 0xdcd60dcf, 0xabd13d59, 0x26d930ac, 0x51de003a,
      0xc8d75180, 0xbfd06116, 0x21b4f4b5, 0x56b3c423, 0xcfba9599,
      Oxb8bda50f. 0x2802b89e. 0x5f058808. 0xc60cd9b2. 0xb10be924.
      0x2f6f7c87, 0x58684c11, 0xc1611dab, 0xb6662d3d, 0x76dc4190,
      0x01db7106, 0x98d220bc, 0xefd5102a, 0x71b18589, 0x06b6b51f,
      0x9fbfe4a5, 0xe8b8d433, 0x7807c9a2, 0x0f00f934, 0x9609a88e,
      Oxe10e9818, 0x7f6a0dbb, 0x086d3d2d, 0x91646c97, 0xe6635c01,
      0x6b6b51f4, 0x1c6c6162, 0x856530d8, 0xf262004e, 0x6c0695ed,
      0x1b01a57b, 0x8208f4c1, 0xf50fc457, 0x65b0d9c6, 0x12b7e950,
      0x8bbeb8ea, 0xfcb9887c, 0x62dd1ddf, 0x15da2d49, 0x8cd37cf3,
      Oxfbd44c65, 0x4db26158, 0x3ab551ce, 0xa3bc0074, 0xd4bb30e2,
      0x4adfa541, 0x3dd895d7, 0xa4d1c46d, 0xd3d6f4fb, 0x4369e96a,
      0x346ed9fc, 0xad678846, 0xda60b8d0, 0x44042d73, 0x33031de5,
      OxaaOa4c5f, OxddOd7cc9, Ox5005713c, Ox270241aa, OxbeOb1010,
      0xc90c2086, 0x5768b525, 0x206f85b3, 0xb966d409, 0xce61e49f,
      0x5edef90e, 0x29d9c998, 0xb0d09822, 0xc7d7a8b4, 0x59b33d17,
      0x2eb40d81, 0xb7bd5c3b, 0xc0ba6cad, 0xedb88320, 0x9abfb3b6,
      0x03b6e20c, 0x74b1d29a, 0xead54739, 0x9dd277af, 0x04db2615,
      0x73dc1683, 0xe3630b12, 0x94643b84, 0x0d6d6a3e, 0x7a6a5aa8,
      Oxe40ecf0b, 0x9309ff9d, 0x0a00ae27, 0x7d079eb1, 0xf00f9344,
      0x8708a3d2, 0x1e01f268, 0x6906c2fe, 0xf762575d, 0x806567cb,
      0x196c3671, 0x6e6b06e7, 0xfed41b76, 0x89d32be0, 0x10da7a5a,
      0x67dd4acc, 0xf9b9df6f, 0x8ebeeff9, 0x17b7be43, 0x60b08ed5,
      Oxd6d6a3e8, Oxa1d1937e, Ox38d8c2c4, Ox4fdff252, Oxd1bb67f1,
      Oxa6bc5767, Ox3fb506dd, Ox48b2364b, Oxd80d2bda, Oxaf0a1b4c,
      0x36034af6, 0x41047a60, 0xdf60efc3, 0xa867df55, 0x316e8eef,
      0x4669be79, 0xcb61b38c, 0xbc66831a, 0x256fd2a0, 0x5268e236,
      OxccOc7795, OxbbOb4703, Ox220216b9, Ox5505262f, Oxc5ba3bbe,
      0xb2bd0b28, 0x2bb45a92, 0x5cb36a04, 0xc2d7ffa7, 0xb5d0cf31,
      0x2cd99e8b, 0x5bdeae1d, 0x9b64c2b0, 0xec63f226, 0x756aa39c,
      0x026d930a, 0x9c0906a9, 0xeb0e363f, 0x72076785, 0x05005713,
      0x95bf4a82, 0xe2b87a14, 0x7bb12bae, 0x0cb61b38, 0x92d28e9b,
      0xe5d5be0d, 0x7cdcefb7, 0x0bdbdf21, 0x86d3d2d4, 0xf1d4e242,
      0x68ddb3f8, 0x1fda836e, 0x81be16cd, 0xf6b9265b, 0x6fb077e1,
      0x18b74777, 0x88085ae6, 0xff0f6a70, 0x66063bca, 0x11010b5c,
      0x8f659eff, 0xf862ae69, 0x616bffd3, 0x166ccf45, 0xa00ae278,
      Oxd70dd2ee, 0x4e048354, 0x3903b3c2, 0xa7672661, 0xd06016f7,
      0x4969474d, 0x3e6e77db, 0xaed16a4a, 0xd9d65adc, 0x40df0b66,
      0x37d83bf0, 0xa9bcae53, 0xdebb9ec5, 0x47b2cf7f, 0x30b5ffe9,
      Oxbdbdf21c, Oxcabac28a, Ox53b39330, Ox24b4a3a6, Oxbad03605,
      Oxcdd70693, 0x54de5729, 0x23d967bf, 0xb3667a2e, 0xc4614ab8,
      0x5d681b02, 0x2a6f2b94, 0xb40bbe37, 0xc30c8ea1, 0x5a05df1b,
      0x2d02ef8d
```

```
};
unsigned char *end;

crc = ~crc & Oxffffffff;
for (end = buf + len; buf < end; ++buf)
   crc = crc32_table[(crc ^ *buf) & Oxff] ^ (crc >> 8);
return ~crc & Oxffffffff;
}
```

This computation does not apply to the "build ID" method.

18.3 Errors Reading Symbol Files

While reading a symbol file, GDB occasionally encounters problems, such as symbol types it does not recognize, or known bugs in compiler output. By default, GDB does not notify you of such problems, since they are relatively common and primarily of interest to people debugging compilers. If you are interested in seeing information about ill-constructed symbol tables, you can either ask GDB to print only one message about each such type of problem, no matter how many times the problem occurs; or you can ask GDB to print more messages, to see how many times the problems occur, with the set complaints command (see Section 22.7 [Optional Warnings and Messages], page 249).

The messages currently printed, and their meanings, include:

inner block not inside outer block in symbol

The symbol information shows where symbol scopes begin and end (such as at the start of a function or a block of statements). This error indicates that an inner scope block is not fully contained in its outer scope blocks.

GDB circumvents the problem by treating the inner block as if it had the same scope as the outer block. In the error message, *symbol* may be shown as "(don't know)" if the outer block is not a function.

block at address out of order

The symbol information for symbol scope blocks should occur in order of increasing addresses. This error indicates that it does not do so.

GDB does not circumvent this problem, and has trouble locating symbols in the source file whose symbols it is reading. (You can often determine what source file is affected by specifying set verbose on. See Section 22.7 [Optional Warnings and Messages], page 249.)

bad block start address patched

The symbol information for a symbol scope block has a start address smaller than the address of the preceding source line. This is known to occur in the SunOS 4.1.1 (and earlier) C compiler.

GDB circumvents the problem by treating the symbol scope block as starting on the previous source line.

bad string table offset in symbol n

Symbol number n contains a pointer into the string table which is larger than the size of the string table.

GDB circumvents the problem by considering the symbol to have the name foo, which may cause other problems if many symbols end up with this name.

unknown symbol type 0xnn

The symbol information contains new data types that GDB does not yet know how to read. 0xnn is the symbol type of the uncomprehended information, in hexadecimal.

GDB circumvents the error by ignoring this symbol information. This usually allows you to debug your program, though certain symbols are not accessible. If you encounter such a problem and feel like debugging it, you can debug gdb with itself, breakpoint on complain, then go up to the function read_dbx_symtab and examine *bufp to see the symbol.

stub type has NULL name

GDB could not find the full definition for a struct or class.

const/volatile indicator missing (ok if using g++ v1.x), got...

The symbol information for a C++ member function is missing some information that recent versions of the compiler should have output for it.

info mismatch between compiler and debugger

GDB could not parse a type specification output by the compiler.

18.4 GDB Data Files

GDB will sometimes read an auxiliary data file. These files are kept in a directory known as the data directory.

You can set the data directory's name, and view the name GDB is currently using.

set data-directory directory

Set the directory which GDB searches for auxiliary data files to directory.

show data-directory

Show the directory GDB searches for auxiliary data files.

You can set the default data directory by using the configure-time '--with-gdb-datadir' option. If the data directory is inside GDB's configured binary prefix (set with '--prefix' or '--exec-prefix'), then the default data directory will be updated automatically if the installed GDB is moved to a new location.

19 Specifying a Debugging Target

A target is the execution environment occupied by your program.

Often, GDB runs in the same host environment as your program; in that case, the debugging target is specified as a side effect when you use the file or core commands. When you need more flexibility—for example, running GDB on a physically separate host, or controlling a standalone system over a serial port or a realtime system over a TCP/IP connection—you can use the target command to specify one of the target types configured for GDB (see Section 19.2 [Commands for Managing Targets], page 200).

It is possible to build gdb for several different target architectures. When gdb is built like that, you can choose one of the available architectures with the **set architecture** command.

set architecture arch

This command sets the current target architecture to arch. The value of arch can be "auto", in addition to one of the supported architectures.

show architecture

Show the current target architecture.

set processor processor

These are alias commands for, respectively, set architecture and show architecture.

19.1 Active Targets

There are three classes of targets: processes, core files, and executable files. GDB can work concurrently on up to three active targets, one in each class. This allows you to (for example) start a process and inspect its activity without abandoning your work on a core file

For example, if you execute 'gdb a.out', then the executable file a.out is the only active target. If you designate a core file as well—presumably from a prior run that crashed and coredumped—then GDB has two active targets and uses them in tandem, looking first in the corefile target, then in the executable file, to satisfy requests for memory addresses. (Typically, these two classes of target are complementary, since core files contain only a program's read-write memory—variables and so on—plus machine status, while executable files contain only the program text and initialized data.)

When you type run, your executable file becomes an active process target as well. When a process target is active, all GDB commands requesting memory addresses refer to that target; addresses in an active core file or executable file target are obscured while the process target is active.

Use the core-file and exec-file commands to select a new core file or executable target (see Section 18.1 [Commands to Specify Files], page 187). To specify as a target a process that is already running, use the attach command (see Section 4.7 [Debugging an Already-running Process], page 31).

19.2 Commands for Managing Targets

target type parameters

Connects the GDB host environment to a target machine or process. A target is typically a protocol for talking to debugging facilities. You use the argument type to specify the type or protocol of the target machine.

Further parameters are interpreted by the target protocol, but typically include things like device names or host names to connect with, process numbers, and band rates

The target command does not repeat if you press RET again after executing the command.

help target

Displays the names of all targets available. To display targets currently selected, use either info target or info files (see Section 18.1 [Commands to Specify Files], page 187).

help target name

Describe a particular target, including any parameters necessary to select it.

set gnutarget args

GDB uses its own library BFD to read your files. GDB knows whether it is reading an executable, a core, or a .o file; however, you can specify the file format with the set gnutarget command. Unlike most target commands, with gnutarget the target refers to a program, not a machine.

Warning: To specify a file format with set gnutarget, you must know the actual BFD name.

See Section 18.1 [Commands to Specify Files], page 187.

show gnutarget

Use the show gnutarget command to display what file format gnutarget is set to read. If you have not set gnutarget, GDB will determine the file format for each file automatically, and show gnutarget displays 'The current BDF target is "auto".

Here are some common targets (available, or not, depending on the GDB configuration):

target exec program

An executable file. 'target exec program' is the same as 'exec-file program'.

target core filename

A core dump file. 'target core filename' is the same as 'core-file file-name'.

target remote medium

A remote system connected to GDB via a serial line or network connection. This command tells GDB to use its own remote protocol over *medium* for debugging. See Chapter 20 [Remote Debugging], page 203.

For example, if you have a board connected to '/dev/ttya' on the machine running gdb, you could say:

target remote /dev/ttya

target remote supports the load command. This is only useful if you have some other way of getting the stub to the target system, and you can put it somewhere in memory where it won't get clobbered by the download.

target sim

Builtin CPU simulator. GDB includes simulators for most architectures. In general,

target sim load run

works; however, you cannot assume that a specific memory map, device drivers, or even basic I/O is available, although some simulators do provide these. For info about any processor-specific simulator details, see the appropriate section in Section 21.3 [Embedded Processors], page 229.

Some configurations may include these targets as well:

target nrom dev

NetROM ROM emulator. This target only supports downloading.

Different targets are available on different configurations of GDB; your configuration may have more or fewer targets.

Many remote targets require you to download the executable's code once you've successfully established a connection. You may wish to control various aspects of this process.

set hash This command controls whether a hash mark '#' is displayed while downloading a file to the remote monitor. If on, a hash mark is displayed after each S-record is successfully downloaded to the monitor.

show hash Show the current status of displaying the hash mark.

set debug monitor

Enable or disable display of communications messages between GDB and the remote monitor.

show debug monitor

Show the current status of displaying communications between GDB and the remote monitor.

load filename

Depending on what remote debugging facilities are configured into GDB, the load command may be available. Where it exists, it is meant to make filename (an executable) available for debugging on the remote system—by downloading, or dynamic linking, for example. load also records the filename symbol table in GDB, like the add-symbol-file command.

If your GDB does not have a load command, attempting to execute it gets the error message "You can't do that when your target is ..."

The file is loaded at whatever address is specified in the executable. For some object file formats, you can specify the load address when you link the program; for other formats, like a.out, the object file format specifies a fixed address.

Depending on the remote side capabilities, GDB may be able to load programs into flash memory.

load does not repeat if you press RET again after using it.

19.3 Choosing Target Byte Order

Some types of processors, such as the MIPS, PowerPC, and Renesas SH, offer the ability to run either big-endian or little-endian byte orders. Usually the executable or symbol will include a bit to designate the endian-ness, and you will not need to worry about which to use. However, you may still find it useful to adjust GDB's idea of processor endian-ness manually.

set endian big

Instruct gdb to assume the target is big-endian.

set endian little

Instruct gdb to assume the target is little-endian.

set endian auto

Instruct gdb to use the byte order associated with the executable.

show endian

Display GDB's current idea of the target byte order.

Note that these commands merely adjust interpretation of symbolic data on the host, and that they have absolutely no effect on the target system.

20 Debugging Remote Programs

If you are trying to debug a program running on a machine that cannot run GDB in the usual way, it is often useful to use remote debugging. For example, you might use remote debugging on an operating system kernel, or on a small system which does not have a general purpose operating system powerful enough to run a full-featured debugger.

Some configurations of GDB have special serial or TCP/IP interfaces to make this work with particular debugging targets. In addition, GDB comes with a generic serial protocol (specific to GDB, but not specific to any particular target system) which you can use if you write the remote stubs—the code that runs on the remote system to communicate with GDB.

Other remote targets may be available in your configuration of GDB; use help target to list them.

20.1 Connecting to a Remote Target

On the GDB host machine, you will need an unstripped copy of your program, since GDB needs symbol and debugging information. Start up GDB as usual, using the name of the local copy of your program as the first argument.

GDB can communicate with the target over a serial line, or over an IP network using TCP or UDP. In each case, GDB uses the same protocol for debugging your program; only the medium carrying the debugging packets varies. The target remote command establishes a connection to the target. Its arguments indicate which medium to use:

target remote serial-device

Use *serial-device* to communicate with the target. For example, to use a serial line connected to the device named '/dev/ttyb':

```
target remote /dev/ttyb
```

If you're using a serial line, you may want to give GDB the '--baud' option, or use the set remotebaud command (see Section 20.4 [Remote Configuration], page 208) before the target command.

```
target remote host:port
target remote tcp:host:port
```

Debug using a TCP connection to port on host. The host may be either a host name or a numeric IP address; port must be a decimal number. The host could be the target machine itself, if it is directly connected to the net, or it might be a terminal server which in turn has a serial line to the target.

For example, to connect to port 2828 on a terminal server named manyfarms:

```
target remote manyfarms:2828
```

If your remote target is actually running on the same machine as your debugger session (e.g. a simulator for your target running on the same host), you can omit the hostname. For example, to connect to port 1234 on your local machine:

```
target remote :1234
```

Note that the colon is still required here.

target remote udp:host:port

Debug using UDP packets to *port* on *host*. For example, to connect to UDP port 2828 on a terminal server named manyfarms:

```
target remote udp:manyfarms:2828
```

When using a UDP connection for remote debugging, you should keep in mind that the 'U' stands for "Unreliable". UDP can silently drop packets on busy or unreliable networks, which will cause havoc with your debugging session.

target remote | command

Run command in the background and communicate with it using a pipe. The command is a shell command, to be parsed and expanded by the system's command shell, /bin/sh; it should expect remote protocol packets on its standard input, and send replies on its standard output. You could use this to run a stand-alone simulator that speaks the remote debugging protocol, to make net connections using programs like ssh, or for other similar tricks.

If command closes its standard output (perhaps by exiting), GDB will try to send it a SIGTERM signal. (If the program has already exited, this will have no effect.)

Once the connection has been established, you can use all the usual commands to examine and change data. The remote program is already running; you can use **step** and **continue**, and you do not need to use **run**.

Whenever GDB is waiting for the remote program, if you type the interrupt character (often Ctrl-c), GDB attempts to stop the program. This may or may not succeed, depending in part on the hardware and the serial drivers the remote system uses. If you type the interrupt character once again, GDB displays this prompt:

```
Interrupted while waiting for the program. Give up (and stop debugging it)? (y or n)
```

If you type y, GDB abandons the remote debugging session. (If you decide you want to try again later, you can use 'target remote' again to connect once more.) If you type n, GDB goes back to waiting.

detach When you have finished debugging the remote program, you can use the detach command to release it from GDB control. Detaching from the target normally resumes its execution, but the results will depend on your particular remote stub. After the detach command, GDB is free to connect to another target.

disconnect

The disconnect command behaves like detach, except that the target is generally not resumed. It will wait for GDB (this instance or another one) to connect and continue debugging. After the disconnect command, GDB is again free to connect to another target.

monitor cmd

This command allows you to send arbitrary commands directly to the remote monitor. Since GDB doesn't care about the commands it sends like this, this command is the way to extend GDB—you can add new commands that only the external monitor will understand and implement.

20.2 Sending files to a remote system

Some remote targets offer the ability to transfer files over the same connection used to communicate with GDB. This is convenient for targets accessible through other means, e.g. GNU/Linux systems running gdbserver over a network interface. For other targets, e.g. embedded devices with only a single serial port, this may be the only way to upload or download files.

Not all remote targets support these commands.

remote put hostfile targetfile

Copy file hostfile from the host system (the machine running gdb) to targetfile on the target system.

remote get targetfile hostfile

Copy file targetfile from the target system to hostfile on the host system.

remote delete targetfile

Delete targetfile from the target system.

20.3 Using the gdbserver Program

gdbserver is a control program for Unix-like systems, which allows you to connect your program with a remote GDB via target remote—but without linking in the usual debugging stub.

gdbserver is not a complete replacement for the debugging stubs, because it requires essentially the same operating-system facilities that GDB itself does. In fact, a system that can run gdbserver to connect to a remote GDB could also run GDB locally! gdbserver is sometimes useful nevertheless, because it is a much smaller program than GDB itself. It is also easier to port than all of GDB, so you may be able to get started more quickly on a new system by using gdbserver. Finally, if you develop code for real-time systems, you may find that the tradeoffs involved in real-time operation make it more convenient to do as much development work as possible on another system, for example by cross-compiling. You can use gdbserver to make a similar choice for debugging.

GDB and gdbserver communicate via either a serial line or a TCP connection, using the standard GDB remote serial protocol.

Warning: gdbserver does not have any built-in security. Do not run gdbserver connected to any public network; a GDB connection to gdbserver provides access to the target system with the same privileges as the user running gdbserver.

20.3.1 Running gdbserver

Run gdbserver on the target system. You need a copy of the program you want to debug, including any libraries it requires. gdbserver does not need your program's symbol table, so you can strip the program if necessary to save space. GDB on the host system does all the symbol handling.

To use the server, you must tell it how to communicate with GDB; the name of your program; and the arguments for your program. The usual syntax is:

```
target> gdbserver comm program [ args ... ]
```

comm is either a device name (to use a serial line) or a TCP hostname and portnumber. For example, to debug Emacs with the argument 'foo.txt' and communicate with gdb over the serial port '/dev/com1':

```
target> gdbserver /dev/com1 emacs foo.txt
```

gdbserver waits passively for the host GDB to communicate with it.

To use a TCP connection instead of a serial line:

```
target> gdbserver host:2345 emacs foo.txt
```

The only difference from the previous example is the first argument, specifying that you are communicating with the host GDB via TCP. The 'host:2345' argument means that gdbserver is to expect a TCP connection from machine 'host' to local TCP port 2345. (Currently, the 'host' part is ignored.) You can choose any number you want for the port number as long as it does not conflict with any TCP ports already in use on the target system (for example, 23 is reserved for telnet). You must use the same port number with the host GDB target remote command.

20.3.1.1 Attaching to a Running Program

On some targets, gdbserver can also attach to running programs. This is accomplished via the --attach argument. The syntax is:

```
target> gdbserver --attach comm pid
```

pid is the process ID of a currently running process. It isn't necessary to point gdbserver at a binary for the running process.

You can debug processes by name instead of process ID if your target has the pidof utility:

```
target> gdbserver --attach comm 'pidof program'
```

In case more than one copy of *program* is running, or *program* has multiple threads, most versions of pidof support the -s option to only return the first process ID.

20.3.1.2 Multi-Process Mode for gdbserver

When you connect to gdbserver using target remote, gdbserver debugs the specified program only once. When the program exits, or you detach from it, GDB closes the connection and gdbserver exits.

If you connect using target extended-remote, gdbserver enters multi-process mode. When the debugged program exits, or you detach from it, GDB stays connected to gdbserver even though no program is running. The run and attach commands instruct gdbserver to run or attach to a new program. The run command uses set remote exec-file (see [set remote exec-file], page 209) to select the program to run. Command line arguments are supported, except for wildcard expansion and I/O redirection (see Section 4.3 [Arguments], page 28).

To start gdbserver without supplying an initial command to run or process ID to attach, use the '--multi' command line option. Then you can connect using target extended-remote and start the program you want to debug.

 $^{^{1}}$ If you choose a port number that conflicts with another service, **gdbserver** prints an error message and exits.

gdbserver does not automatically exit in multi-process mode. You can terminate it by using monitor exit (see [Monitor Commands for gdbserver], page 207).

20.3.1.3 Other Command-Line Arguments for gdbserver

The '--debug' option tells gdbserver to display extra status information about the debugging process. The '--remote-debug' option tells gdbserver to display remote protocol debug output. These options are intended for gdbserver development and for bug reports to the developers.

The '--wrapper' option specifies a wrapper to launch programs for debugging. The option should be followed by the name of the wrapper, then any command-line arguments to pass to the wrapper, then -- indicating the end of the wrapper arguments.

gdbserver runs the specified wrapper program with a combined command line including the wrapper arguments, then the name of the program to debug, then any arguments to the program. The wrapper runs until it executes your program, and then GDB gains control.

You can use any program that eventually calls execve with its arguments as a wrapper. Several standard Unix utilities do this, e.g. env and nohup. Any Unix shell script ending with exec "\$0" will also work.

For example, you can use **env** to pass an environment variable to the debugged program, without setting the variable in **gdbserver**'s environment:

\$ gdbserver --wrapper env LD_PRELOAD=libtest.so -- :2222 ./testprog

20.3.2 Connecting to gdbserver

Run gdb on the host system.

First make sure you have the necessary symbol files. Load symbols for your application using the file command before you connect. Use set sysroot to locate target libraries (unless your GDB was compiled with the correct sysroot using --with-sysroot).

The symbol file and target libraries must exactly match the executable and libraries on the target, with one exception: the files on the host system should not be stripped, even if the files on the target system are. Mismatched or missing files will lead to confusing results during debugging. On GNU/Linux targets, mismatched or missing files may also prevent gdbserver from debugging multi-threaded programs.

Connect to your target (see Section 20.1 [Connecting to a Remote Target], page 203). For TCP connections, you must start up gdbserver prior to using the target remote command. Otherwise you may get an error whose text depends on the host system, but which usually looks something like 'Connection refused'. Don't use the load command in GDB when using gdbserver, since the program is already on the target.

20.3.3 Monitor Commands for gdbserver

During a GDB session using gdbserver, you can use the monitor command to send special requests to gdbserver. Here are the available commands.

monitor help

List the available monitor commands.

monitor set debug 0 monitor set debug 1

Disable or enable general debugging messages.

```
monitor set remote-debug 0 monitor set remote-debug 1
```

Disable or enable specific debugging messages associated with the remote protocol (see Appendix D [Remote Protocol], page 403).

monitor set libthread-db-search-path [PATH]

When this command is issued, path is a colon-separated list of directories to search for libthread_db (see Section 4.10 [set libthread-db-search-path], page 35). If you omit path, 'libthread-db-search-path' will be reset to an empty list.

monitor exit

Tell gdbserver to exit immediately. This command should be followed by disconnect to close the debugging session. gdbserver will detach from any attached processes and kill any processes it created. Use monitor exit to terminate gdbserver at the end of a multi-process mode debug session.

20.4 Remote Configuration

This section documents the configuration options available when debugging remote programs. For the options related to the File I/O extensions of the remote protocol, see [system], page 445.

set remoteaddresssize bits

Set the maximum size of address in a memory packet to the specified number of bits. GDB will mask off the address bits above that number, when it passes addresses to the remote target. The default value is the number of bits in the target's address.

show remoteaddresssize

Show the current value of remote address size in bits.

set remotebaud n

Set the baud rate for the remote serial I/O to n baud. The value is used to set the speed of the serial port used for debugging remote targets.

show remotebaud

Show the current speed of the remote connection.

set remotebreak

If set to on, GDB sends a BREAK signal to the remote when you type Ctrl-c to interrupt the program running on the remote. If set to off, GDB sends the 'Ctrl-C' character instead. The default is off, since most remote systems expect to see 'Ctrl-C' as the interrupt signal.

show remotebreak

Show whether GDB sends BREAK or 'Ctrl-C' to interrupt the remote program.

set remoteflow on set remoteflow off

Enable or disable hardware flow control (RTS/CTS) on the serial port used to communicate to the remote target.

show remoteflow

Show the current setting of hardware flow control.

set remotelogbase base

Set the base (a.k.a. radix) of logging serial protocol communications to base. Supported values of base are: ascii, octal, and hex. The default is ascii.

show remotelogbase

Show the current setting of the radix for logging remote serial protocol.

set remotelogfile file

Record remote serial communications on the named file. The default is not to record at all.

show remotelogfile.

Show the current setting of the file name on which to record the serial communications.

set remotetimeout num

Set the timeout limit to wait for the remote target to respond to *num* seconds. The default is 2 seconds.

show remotetimeout

Show the current number of seconds to wait for the remote target responses.

set remote hardware-watchpoint-limit limit set remote hardware-breakpoint-limit limit

Restrict GDB to using *limit* remote hardware breakpoint or watchpoints. A limit of -1, the default, is treated as unlimited.

set remote exec-file filename

show remote exec-file

Select the file used for run with target extended-remote. This should be set to a filename valid on the target system. If it is not set, the target will use a default filename (e.g. the last program run).

set remote username username

show remote username

Set the username to be sent to targets that require a login. The username is an arbitary string.

set remote password password

show remote password

Set the password to be sent to targets that require a login. The password is an arbitary string. At present it is stored as cleartext.

set remote interrupt-sequence

Allow the user to select one of 'Ctrl-C', a BREAK or 'BREAK-g' as the sequence to the remote target in order to interrupt the execution. 'Ctrl-C' is a default. Some system prefers BREAK which is high level of serial line for some certain time. Linux kernel prefers 'BREAK-g', a.k.a Magic SysRq g. It is BREAK signal followed by character g.

show interrupt-sequence

Show which of 'Ctrl-C', BREAK or BREAK-g is sent by GDB to interrupt the remote program. BREAK-g is BREAK signal followed by g and also known as Magic SysRq g.

set remote interrupt-on-connect

Specify whether interrupt-sequence is sent to remote target when GDB connects to it. This is mostly needed when you debug Linux kernel. Linux kernel expects BREAK followed by g which is known as Magic SysRq g in order to connect GDB.

show interrupt-on-connect

Show whether interrupt-sequence is sent to remote target when GDB connects to it.

set tcp auto-retry on

Enable auto-retry for remote TCP connections. This is useful if the remote debugging agent is launched in parallel with GDB; there is a race condition because the agent may not become ready to accept the connection before GDB attempts to connect. When auto-retry is enabled, if the initial attempt to connect fails, GDB reattempts to establish the connection using the timeout specified by set tcp connect-timeout.

set tcp auto-retry off

Do not auto-retry failed TCP connections.

show tcp auto-retry

Show the current auto-retry setting.

set tcp connect-timeout seconds

Set the timeout for establishing a TCP connection to the remote target to seconds. The timeout affects both polling to retry failed connections (enabled by set tcp auto-retry on) and waiting for connections that are merely slow to complete, and represents an approximate cumulative value.

show tcp connect-timeout

Show the current connection timeout setting.

The GDB remote protocol autodetects the packets supported by your debugging stub. If you need to override the autodetection, you can use these commands to enable or disable individual packets. Each packet can be set to 'on' (the remote target supports this packet), 'off' (the remote target does not support this packet), or 'auto' (detect remote target support for this packet). They all default to 'auto'. For more information about each packet, see Appendix D [Remote Protocol], page 403.

During normal use, you should not have to use any of these commands. If you do, that may be a bug in your remote debugging stub, or a bug in GDB. You may want to report the problem to the GDB developers.

For each packet name, the command to enable or disable the packet is set remote name-packet. The available settings are:

Command Name Rem

Remote Packet

Related Features

fetch-register	p	info registers
set-register	P	set
binary-download	X	load, set
read-aux-vector	qXfer:auxv:read	info auxv
symbol-lookup	qSymbol	Detecting multiple threads
attach	vAttach	attach
verbose-resume	vCont	Stepping or resuming multiple threads
run	vRun	run
software-breakpoint	ZO	break
hardware-breakpoint	Z1	hbreak
write-watchpoint	Z2	watch
read-watchpoint	Z3	rwatch
access-watchpoint	Z4	awatch
target-features	qXfer:features:read	set architecture
library-info	qXfer:libraries:read	info sharedlibrary
memory-map	qXfer:memory-map:read	info mem
read-spu-object	qXfer:spu:read	info spu
write-spu-object	qXfer:spu:write	info spu
read-siginfo-object	qXfer:siginfo:read	<pre>print \$_siginfo</pre>
write-siginfo-object	qXfer:siginfo:write	set \$_siginfo
threads	qXfer:threads:read	info threads

get-thread-local- storage-address	qGetTLSAddr	Displayingthread variables
search-memory	qSearch:memory	find
supported-packets	qSupported	Remote communications parameters
pass-signals	QPassSignals	handle signal
hostio-close-packet	vFile:close	remote get, remote put
hostio-open-packet	vFile:open	remote get, remote put
hostio-pread-packet	vFile:pread	remote get, remote put
hostio-pwrite-packet	vFile:pwrite	remote get, remote put
hostio-unlink-packet	vFile:unlink	remote delete
noack-packet	QStartNoAckMode	Packet acknowledgment
osdata	qXfer:osdata:read	info os
query-attached	qAttached	Querying remote attach state.

20.5 Implementing a Remote Stub

The stub files provided with GDB implement the target side of the communication protocol, and the GDB side is implemented in the GDB source file 'remote.c'. Normally, you can simply allow these subroutines to communicate, and ignore the details. (If you're implementing your own stub file, you can still ignore the details: start with one of the existing stub files. 'sparc-stub.c' is the best organized, and therefore the easiest to read.)

To debug a program running on another machine (the debugging *target* machine), you must first arrange for all the usual prerequisites for the program to run by itself. For example, for a C program, you need:

1. A startup routine to set up the C runtime environment; these usually have a name like 'crt0'. The startup routine may be supplied by your hardware supplier, or you may have to write your own.

- 2. A C subroutine library to support your program's subroutine calls, notably managing input and output.
- 3. A way of getting your program to the other machine—for example, a download program. These are often supplied by the hardware manufacturer, but you may have to write your own from hardware documentation.

The next step is to arrange for your program to use a serial port to communicate with the machine where gdb is running (the *host* machine). In general terms, the scheme looks like this:

On the host,

GDB already understands how to use this protocol; when everything else is set up, you can simply use the 'target remote' command (see Chapter 19 [Specifying a Debugging Target], page 199).

On the target,

you must link with your program a few special-purpose subroutines that implement the GDB remote serial protocol. The file containing these subroutines is called a *debugging stub*.

On certain remote targets, you can use an auxiliary program gdbserver instead of linking a stub into your program. See Section 20.3 [Using the gdbserver Program], page 205, for details.

The debugging stub is specific to the architecture of the remote machine; for example, use 'sparc-stub.c' to debug programs on SPARC boards.

These working remote stubs are distributed with GDB:

i386-stub.c

For Intel 386 and compatible architectures.

m68k-stub.c

For Motorola 680x0 architectures.

sh-stub.c

For Renesas SH architectures.

sparc-stub.c

For SPARC architectures.

sparcl-stub.c

For Fujitsu Sparclite architectures.

The 'README' file in the GDB distribution may list other recently added stubs.

20.5.1 What the Stub Can Do for You

The debugging stub for your architecture supplies these three subroutines:

set_debug_traps

This routine arranges for handle_exception to run when your program stops. You must call this subroutine explicitly near the beginning of your program.

handle_exception

This is the central workhorse, but your program never calls it explicitly—the setup code arranges for handle_exception to run when a trap is triggered.

handle_exception takes control when your program stops during execution (for example, on a breakpoint), and mediates communications with GDB on the host machine. This is where the communications protocol is implemented; handle_exception acts as the GDB representative on the target machine. It begins by sending summary information on the state of your program, then continues to execute, retrieving and transmitting any information GDB needs, until you execute a GDB command that makes your program resume; at that point, handle_exception returns control to your own code on the target machine.

breakpoint

Use this auxiliary subroutine to make your program contain a breakpoint. Depending on the particular situation, this may be the only way for GDB to get control. For instance, if your target machine has some sort of interrupt button, you won't need to call this; pressing the interrupt button transfers control to handle_exception—in effect, to GDB. On some machines, simply receiving characters on the serial port may also trigger a trap; again, in that situation, you don't need to call breakpoint from your own program—simply running 'target remote' from the host GDB session gets control.

Call breakpoint if none of these is true, or if you simply want to make certain your program stops at a predetermined point for the start of your debugging session.

20.5.2 What You Must Do for the Stub

The debugging stubs that come with GDB are set up for a particular chip architecture, but they have no information about the rest of your debugging target machine.

First of all you need to tell the stub how to communicate with the serial port.

int getDebugChar()

Write this subroutine to read a single character from the serial port. It may be identical to getchar for your target system; a different name is used to allow you to distinguish the two if you wish.

void putDebugChar(int)

Write this subroutine to write a single character to the serial port. It may be identical to putchar for your target system; a different name is used to allow you to distinguish the two if you wish.

If you want GDB to be able to stop your program while it is running, you need to use an interrupt-driven serial driver, and arrange for it to stop when it receives a ^C ('\003', the control-C character). That is the character which GDB uses to tell the remote system to stop.

Getting the debugging target to return the proper status to GDB probably requires changes to the standard stub; one quick and dirty way is to just execute a breakpoint instruction (the "dirty" part is that GDB reports a SIGTRAP instead of a SIGINT).

Other routines you need to supply are:

void exceptionHandler (int exception_number, void *exception_address)

Write this function to install exception_address in the exception handling tables. You need to do this because the stub does not have any way of knowing what the exception handling tables on your target system are like (for example, the processor's table might be in ROM, containing entries which point to a table in RAM). exception_number is the exception number which should be changed; its meaning is architecture-dependent (for example, different numbers might represent divide by zero, misaligned access, etc). When this exception occurs, control should be transferred directly to exception_address, and the processor state (stack, registers, and so on) should be just as it is when a processor exception occurs. So if you want to use a jump instruction to reach exception_address, it should be a simple jump, not a jump to subroutine.

For the 386, exception_address should be installed as an interrupt gate so that interrupts are masked while the handler runs. The gate should be at privilege level 0 (the most privileged level). The SPARC and 68k stubs are able to mask interrupts themselves without help from exceptionHandler.

void flush_i_cache()

On SPARC and SPARCLITE only, write this subroutine to flush the instruction cache, if any, on your target machine. If there is no instruction cache, this subroutine may be a no-op.

On target machines that have instruction caches, GDB requires this function to make certain that the state of your program is stable.

You must also make sure this library routine is available:

```
void *memset(void *, int, int)
```

This is the standard library function memset that sets an area of memory to a known value. If you have one of the free versions of libc.a, memset can be found there; otherwise, you must either obtain it from your hardware manufacturer, or write your own.

If you do not use the GNU C compiler, you may need other standard library subroutines as well; this varies from one stub to another, but in general the stubs are likely to use any of the common library subroutines which GCC generates as inline code.

20.5.3 Putting it All Together

In summary, when your program is ready to debug, you must follow these steps.

1. Make sure you have defined the supporting low-level routines (see Section 20.5.2 [What You Must Do for the Stub], page 214):

```
getDebugChar, putDebugChar,
flush_i_cache, memset, exceptionHandler.
```

2. Insert these lines near the top of your program:

```
set_debug_traps();
breakpoint();
```

3. For the 680x0 stub only, you need to provide a variable called exceptionHook. Normally you just use:

void (*exceptionHook)() = 0;

but if before calling set_debug_traps, you set it to point to a function in your program, that function is called when GDB continues after stopping on a trap (for example, bus error). The function indicated by exceptionHook is called with one parameter: an int which is the exception number.

- 4. Compile and link together: your program, the GDB debugging stub for your target architecture, and the supporting subroutines.
- 5. Make sure you have a serial connection between your target machine and the GDB host, and identify the serial port on the host.
- 6. Download your program to your target machine (or get it there by whatever means the manufacturer provides), and start it.
- 7. Start GDB on the host, and connect to the target (see Section 20.1 [Connecting to a Remote Target], page 203).

21 Configuration-Specific Information

While nearly all GDB commands are available for all native and cross versions of the debugger, there are some exceptions. This chapter describes things that are only available in certain configurations.

There are three major categories of configurations: native configurations, where the host and target are the same, embedded operating system configurations, which are usually the same for several different processor architectures, and bare embedded processors, which are quite different from each other.

21.1 Native

This section describes details specific to particular native configurations.

21.1.1 HP-UX

On HP-UX systems, if you refer to a function or variable name that begins with a dollar sign, GDB searches for a user or system name first, before it searches for a convenience variable.

21.1.2 BSD libkym Interface

BSD-derived systems (FreeBSD/NetBSD/OpenBSD) have a kernel memory interface that provides a uniform interface for accessing kernel virtual memory images, including live systems and crash dumps. GDB uses this interface to allow you to debug live kernels and kernel crash dumps on many native BSD configurations. This is implemented as a special kvm debugging target. For debugging a live system, load the currently running kernel into GDB and connect to the kvm target:

(gdb) target kvm

For debugging crash dumps, provide the file name of the crash dump as an argument:

(gdb) target kvm /var/crash/bsd.0

Once connected to the kvm target, the following commands are available:

kvm pcb Set current context from the Process Control Block (PCB) address.

kvm proc Set current context from proc address. This command isn't available on modern FreeBSD systems.

21.1.3 SVR4 Process Information

Many versions of SVR4 and compatible systems provide a facility called '/proc' that can be used to examine the image of a running process using file-system subroutines. If GDB is configured for an operating system with this facility, the command info proc is available to report information about the process running your program, or about any process running on your system. info proc works only on SVR4 systems that include the procfs code. This includes, as of this writing, GNU/Linux, OSF/1 (Digital Unix), Solaris, Irix, and Unixware, but not HP-UX, for example.

info proc

info proc process-id

Summarize available information about any running process. If a process ID is specified by *process-id*, display information about that process; otherwise

display information about the program being debugged. The summary includes the debugged process ID, the command line used to invoke it, its current working directory, and its executable file's absolute file name.

On some systems, process-id can be of the form '[pid]/tid' which specifies a certain thread ID within a process. If the optional pid part is missing, it means a thread from the process being debugged (the leading '/' still needs to be present, or else GDB will interpret the number as a process ID rather than a thread ID).

info proc mappings

Report the memory address space ranges accessible in the program, with information on whether the process has read, write, or execute access rights to each range. On GNU/Linux systems, each memory range includes the object file which is mapped to that range, instead of the memory access rights to that range.

info proc stat info proc status

These subcommands are specific to GNU/Linux systems. They show the process-related information, including the user ID and group ID; how many threads are there in the process; its virtual memory usage; the signals that are pending, blocked, and ignored; its TTY; its consumption of system and user time; its stack size; its 'nice' value; etc. For more information, see the 'proc' man page (type man 5 proc from your shell prompt).

info proc all

Show all the information about the process described under all of the above info proc subcommands.

set procfs-trace

This command enables and disables tracing of procfs API calls.

show procfs-trace

Show the current state of procfs API call tracing.

set procfs-file file

Tell gdb to write procfs API trace to the named file. gdb appends the trace info to the previous contents of the file. The default is to display the trace on the standard output.

show procfs-file

Show the file to which procfs API trace is written.

```
proc-trace-entry
proc-trace-exit
proc-untrace-entry
proc-untrace-exit
```

These commands enable and disable tracing of entries into and exits from the syscall interface.

info pidlist

For QNX Neutrino only, this command displays the list of all the processes and all the threads within each process.

info meminfo

For QNX Neutrino only, this command displays the list of all mapinfos.

21.1.4 Features for Debugging DJGPP Programs

DJGPP is a port of the GNU development tools to MS-DOS and MS-Windows. DJGPP programs are 32-bit protected-mode programs that use the *DPMI* (DOS Protected-Mode Interface) API to run on top of real-mode DOS systems and their emulations.

GDB supports native debugging of DJGPP programs, and defines a few commands specific to the DJGPP port. This subsection describes those commands.

info dos This is a prefix of DJGPP-specific commands which print information about the target system and important OS structures.

info dos sysinfo

This command displays assorted information about the underlying platform: the CPU type and features, the OS version and flavor, the DPMI version, and the available conventional and DPMI memory.

info dos gdt
info dos ldt
info dos idt

These 3 commands display entries from, respectively, Global, Local, and Interrupt Descriptor Tables (GDT, LDT, and IDT). The descriptor tables are data structures which store a descriptor for each segment that is currently in use. The segment's selector is an index into a descriptor table; the table entry for that index holds the descriptor's base address and limit, and its attributes and access rights.

A typical DJGPP program uses 3 segments: a code segment, a data segment (used for both data and the stack), and a DOS segment (which allows access to DOS/BIOS data structures and absolute addresses in conventional memory). However, the DPMI host will usually define additional segments in order to support the DPMI environment.

These commands allow to display entries from the descriptor tables. Without an argument, all entries from the specified table are displayed. An argument, which should be an integer expression, means display a single entry whose index is given by the argument. For example, here's a convenient way to display information about the debugged program's data segment:

```
(gdb) info dos ldt $ds
0x13f: base=0x11970000 limit=0x0009ffff 32-Bit Data (Read/Write, Exp-up)
```

This comes in handy when you want to see whether a pointer is outside the data segment's limit (i.e. garbled).

info dos pde info dos pte

These two commands display entries from, respectively, the Page Directory and the Page Tables. Page Directories and Page Tables are data structures which control how virtual memory addresses are mapped into physical addresses. A Page Table includes an entry for every page of memory that is mapped into the

program's address space; there may be several Page Tables, each one holding up to 4096 entries. A Page Directory has up to 4096 entries, one each for every Page Table that is currently in use.

Without an argument, info dos pde displays the entire Page Directory, and info dos pte displays all the entries in all of the Page Tables. An argument, an integer expression, given to the info dos pde command means display only that entry from the Page Directory table. An argument given to the info dos pte command means display entries from a single Page Table, the one pointed to by the specified entry in the Page Directory.

These commands are useful when your program uses DMA (Direct Memory Access), which needs physical addresses to program the DMA controller.

These commands are supported only with some DPMI servers.

info dos address-pte addr

This command displays the Page Table entry for a specified linear address. The argument addr is a linear address which should already have the appropriate segment's base address added to it, because this command accepts addresses which may belong to any segment. For example, here's how to display the Page Table entry for the page where a variable i is stored:

(gdb) info dos address-pte __djgpp_base_address + (char *)&i
Page Table entry for address 0x11a00d30:
Base=0x02698000 Dirty Acc. Not-Cached Write-Back Usr Read-Write +0xd30

This says that i is stored at offset 0xd30 from the page whose physical

This says that \mathtt{i} is stored at offset 0xd30 from the page whose physical base address is 0x02698000, and shows all the attributes of that page.

Note that you must cast the addresses of variables to a char *, since otherwise the value of __djgpp_base_address, the base address of all variables and functions in a DJGPP program, will be added using the rules of C pointer arithmetics: if i is declared an int, GDB will add 4 times the value of __djgpp_base_address to the address of i.

Here's another example, it displays the Page Table entry for the transfer buffer: (gdb) info dos address-pte *((unsigned *)&_go32_info_block + 3)
Page Table entry for address 0x29110:

Base=0x00029000 Dirty Acc. Not-Cached Write-Back Usr Read-Write +0x110

(The + 3 offset is because the transfer buffer's address is the 3rd member of the _go32_info_block structure.) The output clearly shows that this DPMI server maps the addresses in conventional memory 1:1, i.e. the physical (0x00029000 + 0x110) and linear (0x29110) addresses are identical.

This command is supported only with some DPMI servers.

In addition to native debugging, the DJGPP port supports remote debugging via a serial data link. The following commands are specific to remote serial debugging in the DJGPP port of GDB.

set comilhase addr

This command sets the base I/O port address of the 'COM1' serial port.

set comling ira

This command sets the *Interrupt Request* (IRQ) line to use for the 'COM1' serial port.

There are similar commands 'set com2base', 'set com3irq', etc. for setting the port address and the IRQ lines for the other 3 COM ports.

The related commands 'show com1base', 'show com1irq' etc. display the current settings of the base address and the IRQ lines used by the COM ports.

info serial

This command prints the status of the 4 DOS serial ports. For each port, it prints whether it's active or not, its I/O base address and IRQ number, whether it uses a 16550-style FIFO, its baudrate, and the counts of various errors encountered so far.

21.1.5 Features for Debugging MS Windows PE Executables

GDB supports native debugging of MS Windows programs, including DLLs with and without symbolic debugging information.

MS-Windows programs that call SetConsoleMode to switch off the special meaning of the 'Ctrl-C' keystroke cannot be interrupted by typing C-c. For this reason, GDB on MS-Windows supports C-BREAK as an alternative interrupt key sequence, which can be used to interrupt the debuggee even if it ignores C-c.

There are various additional Cygwin-specific commands, described in this section. Working with DLLs that have no debugging symbols is described in Section 21.1.5.1 [Non-debug DLL Symbols], page 222.

info w32 This is a prefix of MS Windows-specific commands which print information about the target system and important OS structures.

info w32 selector

This command displays information returned by the Win32 API GetThreadSelectorEntry function. It takes an optional argument that is evaluated to a long value to give the information about this given selector. Without argument, this command displays information about the six segment registers.

info dll This is a Cygwin-specific alias of info shared.

dll-symbols

This command loads symbols from a dll similarly to add-sym command but without the need to specify a base address.

set cygwin-exceptions mode

If mode is on, GDB will break on exceptions that happen inside the Cygwin DLL. If mode is off, GDB will delay recognition of exceptions, and may ignore some exceptions which seem to be caused by internal Cygwin DLL "bookkeeping". This option is meant primarily for debugging the Cygwin DLL itself; the default value is off to avoid annoying GDB users with false SIGSEGV signals.

show cygwin-exceptions

Displays whether gdb will break on exceptions that happen inside the Cygwin DLL itself.

set new-console mode

If mode is on the debuggee will be started in a new console on next start. If mode is offi, the debuggee will be started in the same console as the debugger.

show new-console

Displays whether a new console is used when the debuggee is started.

set new-group mode

This boolean value controls whether the debuggee should start a new group or stay in the same group as the debugger. This affects the way the Windows OS handles 'Ctrl-C'.

show new-group

Displays current value of new-group boolean.

set debugevents

This boolean value adds debug output concerning kernel events related to the debuggee seen by the debugger. This includes events that signal thread and process creation and exit, DLL loading and unloading, console interrupts, and debugging messages produced by the Windows OutputDebugString API call.

set debugexec

This boolean value adds debug output concerning execute events (such as resume thread) seen by the debugger.

set debugexceptions

This boolean value adds debug output concerning exceptions in the debuggee seen by the debugger.

set debugmemory

This boolean value adds debug output concerning debuggee memory reads and writes by the debugger.

set shell This boolean values specifies whether the debuggee is called via a shell or directly (default value is on).

show shell

Displays if the debuggee will be started with a shell.

21.1.5.1 Support for DLLs without Debugging Symbols

Very often on windows, some of the DLLs that your program relies on do not include symbolic debugging information (for example, 'kernel32.dll'). When GDB doesn't recognize any debugging symbols in a DLL, it relies on the minimal amount of symbolic information contained in the DLL's export table. This section describes working with such symbols, known internally to GDB as "minimal symbols".

Note that before the debugged program has started execution, no DLLs will have been loaded. The easiest way around this problem is simply to start the program — either by setting a breakpoint or letting the program run once to completion. It is also possible to force gdb to load a particular DLL before starting the executable — see the shared library information in Section 18.1 [Files], page 187, or the dll-symbols command in Section 21.1.5 [Cygwin Native], page 221. Currently, explicitly loading symbols from a DLL with no debugging information will cause the symbol names to be duplicated in gdb's lookup table, which may adversely affect symbol lookup performance.

21.1.5.2 DLL Name Prefixes

In keeping with the naming conventions used by the Microsoft debugging tools, DLL export symbols are made available with a prefix based on the DLL name, for instance KERNEL32!CreateFileA. The plain name is also entered into the symbol table, so CreateFileA is often sufficient. In some cases there will be name clashes within a program (particularly if the executable itself includes full debugging symbols) necessitating the use of the fully qualified name when referring to the contents of the DLL. Use single-quotes around the name to avoid the exclamation mark ("!") being interpreted as a language operator.

Note that the internal name of the DLL may be all upper-case, even though the file name of the DLL is lower-case, or vice-versa. Since symbols within GDB are *case-sensitive* this may cause some confusion. If in doubt, try the info functions and info variables commands or even maint print msymbols (see Chapter 16 [Symbols], page 175). Here's an example:

```
(gdb) info function CreateFileA
All functions matching regular expression "CreateFileA":
Non-debugging symbols:
0x77e885f4 CreateFileA
0x77e885f4 KERNEL32!CreateFileA
(gdb) info function !
All functions matching regular expression "!":
Non-debugging symbols:
0x6100114c cygwin1!__assert
0x61004034 cygwin1!_dll_crt0@0
0x61004240 cygwin1!dll_crt0(per_process *)
[etc...]
```

21.1.5.3 Working with Minimal Symbols

Symbols extracted from a DLL's export table do not contain very much type information. All that GDB can do is guess whether a symbol refers to a function or variable depending on the linker section that contains the symbol. Also note that the actual contents of the memory contained in a DLL are not available unless the program is running. This means that you cannot examine the contents of a variable or disassemble a function within a DLL without a running program.

Variables are generally treated as pointers and dereferenced automatically. For this reason, it is often necessary to prefix a variable name with the address-of operator ("&") and provide explicit type information in the command. Here's an example of the type of problem:

```
(gdb) print 'cygwin1!__argv'
$1 = 268572168
(gdb) x 'cygwin1!__argv'
0x10021610: "\230y\""

And two possible solutions:
(gdb) print ((char **)'cygwin1!__argv')[0]
$2 = 0x22fd98 "/cygdrive/c/mydirectory/myprogram"
(gdb) x/2x &'cygwin1!__argv'
0x610c0aa8 <cygwin1!__argv>: 0x10021608 0x00000000
```

```
(gdb) x/x 0x10021608

0x10021608: 0x0022fd98

(gdb) x/s 0x0022fd98

0x22fd98: "/cygdrive/c/mydirectory/myprogram"
```

Setting a break point within a DLL is possible even before the program starts execution. However, under these circumstances, GDB can't examine the initial instructions of the function in order to skip the function's frame set-up code. You can work around this by using "*&" to set the breakpoint at a raw memory address:

```
(gdb) break *&'python22!PyOS_Readline'
Breakpoint 1 at 0x1e04eff0
```

The author of these extensions is not entirely convinced that setting a break point within a shared DLL like 'kernel32.dll' is completely safe.

21.1.6 Commands Specific to GNU Hurd Systems

This subsection describes GDB commands specific to the GNU Hurd native debugging.

```
set signals
```

This command toggles the state of inferior signal interception by GDB. Mach exceptions, such as breakpoint traps, are not affected by this command. sigs is a shorthand alias for signals.

show signals

show sigs Show the current state of intercepting inferior's signals.

```
set signal-thread set sigthread
```

This command tells GDB which thread is the libc signal thread. That thread is run when a signal is delivered to a running process. set sigthread is the shorthand alias of set signal-thread.

```
show signal-thread show sigthread
```

These two commands show which thread will run when the inferior is delivered a signal.

set stopped

This commands tells GDB that the inferior process is stopped, as with the SIGSTOP signal. The stopped process can be continued by delivering a signal to it.

show stopped

This command shows whether GDB thinks the debuggee is stopped.

set exceptions

Use this command to turn off trapping of exceptions in the inferior. When exception trapping is off, neither breakpoints nor single-stepping will work. To restore the default, set exception trapping on.

show exceptions

Show the current state of trapping exceptions in the inferior.

set task pause

This command toggles task suspension when GDB has control. Setting it to on takes effect immediately, and the task is suspended whenever GDB gets control. Setting it to off will take effect the next time the inferior is continued. If this option is set to off, you can use set thread default pause on or set thread pause on (see below) to pause individual threads.

show task pause

Show the current state of task suspension.

set task detach-suspend-count

This command sets the suspend count the task will be left with when GDB detaches from it.

show task detach-suspend-count

Show the suspend count the task will be left with when detaching.

set task exception-port set task excp

This command sets the task exception port to which GDB will forward exceptions. The argument should be the value of the send rights of the task. set task excp is a shorthand alias.

set noninvasive

This command switches GDB to a mode that is the least invasive as far as interfering with the inferior is concerned. This is the same as using set task pause, set exceptions, and set signals to values opposite to the defaults.

```
info send-rights
info receive-rights
info port-rights
info port-sets
info dead-names
info ports
info psets
```

These commands display information about, respectively, send rights, receive rights, port rights, port sets, and dead names of a task. There are also shorthand aliases: info ports for info port-rights and info psets for info port-sets.

set thread pause

This command toggles current thread suspension when GDB has control. Setting it to on takes effect immediately, and the current thread is suspended whenever GDB gets control. Setting it to off will take effect the next time the inferior is continued. Normally, this command has no effect, since when GDB has control, the whole task is suspended. However, if you used set task pause off (see above), this command comes in handy to suspend only the current thread.

show thread pause

This command shows the state of current thread suspension.

set thread run

This command sets whether the current thread is allowed to run.

show thread run

Show whether the current thread is allowed to run.

set thread detach-suspend-count

This command sets the suspend count GDB will leave on a thread when detaching. This number is relative to the suspend count found by GDB when it notices the thread; use set thread takeover-suspend-count to force it to an absolute value.

show thread detach-suspend-count

Show the suspend count GDB will leave on the thread when detaching.

$\verb|set thread exception-port|\\$

set thread excp

Set the thread exception port to which to forward exceptions. This overrides the port set by set task exception-port (see above). set thread excp is the shorthand alias.

set thread takeover-suspend-count

Normally, GDB's thread suspend counts are relative to the value GDB finds when it notices each thread. This command changes the suspend counts to be absolute instead.

set thread default

show thread default

Each of the above set thread commands has a set thread default counterpart (e.g., set thread default pause, set thread default exception-port, etc.). The thread default variety of commands sets the default thread properties for all threads; you can then change the properties of individual threads with the non-default commands.

21.1.7 QNX Neutrino

GDB provides the following commands specific to the QNX Neutrino target:

set debug nto-debug

When set to on, enables debugging messages specific to the QNX Neutrino support.

show debug nto-debug

Show the current state of QNX Neutrino messages.

21.1.8 Darwin

GDB provides the following commands specific to the Darwin target:

set debug darwin num

When set to a non zero value, enables debugging messages specific to the Darwin support. Higher values produce more verbose output.

show debug darwin

Show the current state of Darwin messages.

set debug mach-o num

When set to a non zero value, enables debugging messages while GDB is reading Darwin object files. (*Mach-O* is the file format used on Darwin for object and executable files.) Higher values produce more verbose output. This is a command to diagnose problems internal to GDB and should not be needed in normal usage.

show debug mach-o

Show the current state of Mach-O file messages.

set mach-exceptions on set mach-exceptions off

On Darwin, faults are first reported as a Mach exception and are then mapped to a Posix signal. Use this command to turn on trapping of Mach exceptions in the inferior. This might be sometimes useful to better understand the cause of a fault. The default is off.

show mach-exceptions

Show the current state of exceptions trapping.

21.2 Embedded Operating Systems

This section describes configurations involving the debugging of embedded operating systems that are available for several different architectures.

GDB includes the ability to debug programs running on various real-time operating systems.

21.2.1 Using gdb with VxWorks

target vxworks machinename

A VxWorks system, attached via TCP/IP. The argument *machinename* is the target system's machine name or IP address.

On VxWorks, load links filename dynamically on the current target system as well as adding its symbols in GDB.

GDB enables developers to spawn and debug tasks running on networked VxWorks targets from a Unix host. Already-running tasks spawned from the VxWorks shell can also be debugged. GDB uses code that runs on both the Unix host and on the VxWorks target. The program gdb is installed and executed on the Unix host. (It may be installed with the name vxgdb, to distinguish it from a GDB for debugging programs on the host itself.)

VxWorks-timeout args

All VxWorks-based targets now support the option vxworks-timeout. This option is set by the user, and args represents the number of seconds GDB waits for responses to rpc's. You might use this if your VxWorks target is a slow software simulator or is on the far side of a thin network line.

The following information on connecting to VxWorks was current when this manual was produced; newer releases of VxWorks may use revised procedures.

To use GDB with VxWorks, you must rebuild your VxWorks kernel to include the remote debugging interface routines in the VxWorks library 'rdb.a'. To do this, define INCLUDE_RDB in the VxWorks configuration file 'configAll.h' and rebuild your VxWorks kernel. The

resulting kernel contains 'rdb.a', and spawns the source debugging task tRdbTask when VxWorks is booted. For more information on configuring and remaking VxWorks, see the manufacturer's manual.

Once you have included 'rdb.a' in your VxWorks system image and set your Unix execution search path to find GDB, you are ready to run GDB. From your Unix host, run gdb (or vxgdb, depending on your installation).

```
GDB comes up showing the prompt: (vxgdb)
```

21.2.1.1 Connecting to VxWorks

The GDB command target lets you connect to a VxWorks target on the network. To connect to a target whose host name is "tt", type:

```
(vxgdb) target vxworks tt
GDB displays messages like these:
   Attaching remote machine across net...
Connected to tt.
```

GDB then attempts to read the symbol tables of any object modules loaded into the VxWorks target since it was last booted. GDB locates these files by searching the directories listed in the command search path (see Section 4.4 [Your Program's Environment], page 29); if it fails to find an object file, it displays a message such as:

```
prog.o: No such file or directory.
```

When this happens, add the appropriate directory to the search path with the GDB command path, and execute the target command again.

21.2.1.2 VxWorks Download

If you have connected to the VxWorks target and you want to debug an object that has not yet been loaded, you can use the GDB load command to download a file from Unix to VxWorks incrementally. The object file given as an argument to the load command is actually opened twice: first by the VxWorks target in order to download the code, then by GDB in order to read the symbol table. This can lead to problems if the current working directories on the two systems differ. If both systems have NFS mounted the same filesystems, you can avoid these problems by using absolute paths. Otherwise, it is simplest to set the working directory on both systems to the directory in which the object file resides, and then to reference the file by its name, without any path. For instance, a program 'prog.o' may reside in 'vxpath/vw/demo/rdb' in VxWorks and in 'hostpath/vw/demo/rdb' on the host. To load this program, type this on VxWorks:

```
-> cd "vxpath/vw/demo/rdb"

Then, in GDB, type:
    (vxgdb) cd hostpath/vw/demo/rdb
    (vxgdb) load prog.o

GDB displays a response similar to this:
    Reading symbol data from wherever/vw/demo/rdb/prog.o... done.
```

You can also use the load command to reload an object module after editing and recompiling the corresponding source file. Note that this makes GDB delete all currently-defined breakpoints, auto-displays, and convenience variables, and to clear the value history. (This

is necessary in order to preserve the integrity of debugger's data structures that reference the target system's symbol table.)

21.2.1.3 Running Tasks

You can also attach to an existing task using the attach command as follows:

(vxgdb) attach task

where task is the VxWorks hexadecimal task ID. The task can be running or suspended when you attach to it. Running tasks are suspended at the time of attachment.

21.3 Embedded Processors

This section goes into details specific to particular embedded configurations.

Whenever a specific embedded processor has a simulator, GDB allows to send an arbitrary command to the simulator.

sim command

Send an arbitrary *command* string to the simulator. Consult the documentation for the specific simulator in use for information about acceptable commands.

21.3.1 ARM

target rdi dev

ARM Angel monitor, via RDI library interface to ADP protocol. You may use this target to communicate with both boards running the Angel monitor, or with the EmbeddedICE JTAG debug device.

target rdp dev

ARM Demon monitor.

GDB provides the following ARM-specific commands:

set arm disassembler

This commands selects from a list of disassembly styles. The "std" style is the standard style.

show arm disassembler

Show the current disassembly style.

set arm apcs32

This command toggles ARM operation mode between 32-bit and 26-bit.

show arm apcs32

Display the current usage of the ARM 32-bit mode.

set arm fpu fputype

This command sets the ARM floating-point unit (FPU) type. The argument fputype can be one of these:

auto Determine the FPU type by querying the OS ABI.

softfpa Software FPU, with mixed-endian doubles on little-endian ARM

processors.

fpa GCC-compiled FPA co-processor.

softvfp Software FPU with pure-endian doubles.

vfp VFP co-processor.

show arm fpu

Show the current type of the FPU.

set arm abi

This command forces GDB to use the specified ABI.

show arm abi

Show the currently used ABI.

set arm fallback-mode (arm|thumb|auto)

GDB uses the symbol table, when available, to determine whether instructions are ARM or Thumb. This command controls GDB's default behavior when the symbol table is not available. The default is 'auto', which causes GDB to use the current execution mode (from the T bit in the CPSR register).

show arm fallback-mode

Show the current fallback instruction mode.

set arm force-mode (arm|thumb|auto)

This command overrides use of the symbol table to determine whether instructions are ARM or Thumb. The default is 'auto', which causes GDB to use the symbol table and then the setting of 'set arm fallback-mode'.

show arm force-mode

Show the current forced instruction mode.

set debug arm

Toggle whether to display ARM-specific debugging messages from the ARM target support subsystem.

show debug arm

Show whether ARM-specific debugging messages are enabled.

The following commands are available when an ARM target is debugged using the RDI interface:

rdilogfile [file]

Set the filename for the ADP (Angel Debugger Protocol) packet log. With an argument, sets the log file to the specified *file*. With no argument, show the current log file name. The default log file is 'rdi.log'.

rdilogenable [arg]

Control logging of ADP packets. With an argument of 1 or "yes" enables logging, with an argument 0 or "no" disables it. With no arguments displays the current setting. When logging is enabled, ADP packets exchanged between GDB and the RDI target device are logged to a file.

set rdiromatzero

Tell gdb whether the target has ROM at address 0. If on, vector catching is disabled, so that zero address can be used. If off (the default), vector catching is enabled. For this command to take effect, it needs to be invoked prior to the target rdi command.

show rdiromatzero

Show the current setting of ROM at zero address.

set rdiheartbeat

Enable or disable RDI heartbeat packets. It is not recommended to turn on this option, since it confuses ARM and EPI JTAG interface, as well as the Angel monitor.

show rdiheartbeat

Show the setting of RDI heartbeat packets.

21.3.2 Renesas M32R/D and M32R/SDI

target m32r dev

Renesas M32R/D ROM monitor.

target m32rsdi dev

Renesas M32R SDI server, connected via parallel port to the board.

The following gdb commands are specific to the M32R monitor:

set download-path path

Set the default path for finding downloadable SREC files.

show download-path

Show the default path for downloadable SREC files.

set board-address addr

Set the IP address for the M32R-EVA target board.

show board-address

Show the current IP address of the target board.

set server-address addr

Set the IP address for the download server, which is the GDB's host machine.

show server-address

Display the IP address of the download server.

upload [file]

Upload the specified SREC file via the monitor's Ethernet upload capability. If no file argument is given, the current executable file is uploaded.

tload [file]

Test the upload command.

The following commands are available for M32R/SDI:

sdireset This command resets the SDI connection.

sdistatus

This command shows the SDI connection status.

debug_chaos

Instructs the remote that M32R/Chaos debugging is to be used.

use_debug_dma

Instructs the remote to use the DEBUG_DMA method of accessing memory.

use_mon_code

Instructs the remote to use the MON_CODE method of accessing memory.

use_ib_break

Instructs the remote to set breakpoints by IB break.

use_dbt_break

Instructs the remote to set breakpoints by DBT.

21.3.3 M68k

The Motorola m68k configuration includes ColdFire support, and a target command for the following ROM monitor.

target dbug dev

dBUG ROM monitor for Motorola ColdFire.

21.3.4 MicroBlaze

The MicroBlaze is a soft-core processor supported on various Xilinx FPGAs, such as Spartan or Virtex series. Boards with these processors usually have JTAG ports which connect to a host system running the Xilinx Embedded Development Kit (EDK) or Software Development Kit (SDK). This host system is used to download the configuration bitstream to the target FPGA. The Xilinx Microprocessor Debugger (XMD) program communicates with the target board using the JTAG interface and presents a gdbserver interface to the board. By default xmd uses port 1234. (While it is possible to change this default port, it requires the use of undocumented xmd commands. Contact Xilinx support if you need to do this.)

Use these GDB commands to connect to the MicroBlaze target processor.

target remote :1234

Use this command to connect to the target if you are running GDB on the same system as xmd.

target remote xmd-host:1234

Use this command to connect to the target if it is connected to xmd running on a different system named xmd-host.

load Use this command to download a program to the MicroBlaze target.

set debug microblaze n

Enable MicroBlaze-specific debugging messages if non-zero.

show debug microblaze n

Show MicroBlaze-specific debugging level.

21.3.5 MIPS Embedded

GDB can use the MIPS remote debugging protocol to talk to a MIPS board attached to a serial line. This is available when you configure GDB with '--target=mips-idt-ecoff'.

Use these gdb commands to specify the connection to your target board:

target mips port

To run a program on the board, start up gdb with the name of your program as the argument. To connect to the board, use the command 'target mips port', where port is the name of the serial port connected to the board. If the program has not already been downloaded to the board, you may use the load command to download it. You can then use all the usual GDB commands.

For example, this sequence connects to the target board through a serial port, and loads and runs a program called *prog* through the debugger:

```
host$ gdb prog
GDB is free software and ...
(gdb) target mips /dev/ttyb
(gdb) load prog
(gdb) run
```

target mips hostname:portnumber

On some GDB host configurations, you can specify a TCP connection (for instance, to a serial line managed by a terminal concentrator) instead of a serial port, using the syntax 'hostname:portnumber'.

target pmon port

PMON ROM monitor.

target ddb port

NEC's DDB variant of PMON for Vr4300.

target lsi port

LSI variant of PMON.

target r3900 dev

Densan DVE-R3900 ROM monitor for Toshiba R3900 Mips.

target array dev

Array Tech LSI33K RAID controller board.

GDB also supports these special commands for MIPS targets:

```
set mipsfpu double
set mipsfpu single
set mipsfpu none
set mipsfpu auto
show mipsfpu
```

If your target board does not support the MIPS floating point coprocessor, you should use the command 'set mipsfpu none' (if you need this, you may wish to put the command in your GDB init file). This tells GDB how to find the return value of functions which return floating point values. It also allows GDB to avoid saving the floating point registers when calling functions on the board. If you are using a floating point coprocessor with only single precision floating point support, as on the R4650 processor, use the command 'set mipsfpu single'. The default double precision floating point coprocessor may be selected using 'set mipsfpu double'.

In previous versions the only choices were double precision or no floating point, so 'set mipsfpu on' will select double precision and 'set mipsfpu off' will select no floating point.

As usual, you can inquire about the mipsfpu variable with 'show mipsfpu'.

set timeout seconds set retransmit-timeout seconds show timeout show retransmit-timeout

You can control the timeout used while waiting for a packet, in the MIPS remote protocol, with the set timeout seconds command. The default is 5 seconds. Similarly, you can control the timeout used while waiting for an acknowledgment of a packet with the set retransmit-timeout seconds command. The default is 3 seconds. You can inspect both values with show timeout and show retransmit-timeout. (These commands are only available when GDB is configured for '--target=mips-idt-ecoff'.)

The timeout set by set timeout does not apply when GDB is waiting for your program to stop. In that case, GDB waits forever because it has no way of knowing how long the program is going to run before stopping.

set syn-garbage-limit num

Limit the maximum number of characters GDB should ignore when it tries to synchronize with the remote target. The default is 10 characters. Setting the limit to -1 means there's no limit.

show syn-garbage-limit

Show the current limit on the number of characters to ignore when trying to synchronize with the remote system.

set monitor-prompt prompt

Tell GDB to expect the specified *prompt* string from the remote monitor. The default depends on the target:

```
pmon target 'PMON' ddb target 'NEC010' lsi target 'PMON'
```

show monitor-prompt

Show the current strings GDB expects as the prompt from the remote monitor.

set monitor-warnings

Enable or disable monitor warnings about hardware breakpoints. This has effect only for the lsi target. When on, GDB will display warning messages whose codes are returned by the lsi PMON monitor for breakpoint commands.

show monitor-warnings

Show the current setting of printing monitor warnings.

pmon command

This command allows sending an arbitrary *command* string to the monitor. The monitor must be in debug mode for this to work.

21.3.6 OpenRISC 1000

See OR1k Architecture document (www.opencores.org) for more information about platform and commands.

target jtag jtag://host:port

Connects to remote JTAG server. JTAG remote server can be either an or1ksim or JTAG server, connected via parallel port to the board.

Example: target jtag jtag://localhost:9999

or1ksim command

If connected to orlksim OpenRISC 1000 Architectural Simulator, proprietary commands can be executed.

info or1k spr

Displays spr groups.

info or1k spr group

info or1k spr groupno

Displays register names in selected group.

info or1k spr group register

info or1k spr register

info or1k spr groupno registerno

info or1k spr registerno

Shows information about specified spr register.

spr group register value

spr register value

spr groupno registerno value

spr registerno value

Writes value to specified spr register.

Some implementations of OpenRISC 1000 Architecture also have hardware trace. It is very similar to GDB trace, except it does not interfere with normal program execution and is thus much faster. Hardware breakpoints/watchpoint triggers can be set using:

\$LEA/\$LDATA

Load effective address/data

\$SEA/\$SDATA

Store effective address/data

\$AEA/\$ADATA

Access effective address (\$SEA or \$LEA) or data (\$SDATA/\$LDATA)

\$FETCH Fetch data

When triggered, it can capture low level data, like: PC, LSEA, LDATA, SDATA, READSPR, WRITESPR, INSTR.

htrace commands:

hwatch conditional

Set hardware watchpoint on combination of Load/Store Effective Address(es) or Data. For example:

```
hwatch ($LEA == my_var) && ($LDATA < 50) || ($SEA == my_var) && ($SDATA >= 50) hwatch ($LEA == my_var) && ($LDATA < 50) || ($SEA == my_var) && ($SDATA >= 50)
```

htrace info

Display information about current HW trace configuration.

htrace trigger conditional

Set starting criteria for HW trace.

htrace qualifier conditional

Set acquisition qualifier for HW trace.

htrace stop conditional

Set HW trace stopping criteria.

htrace record [data] *

Selects the data to be recorded, when qualifier is met and HW trace was triggered.

htrace enable

htrace disable

Enables/disables the HW trace.

htrace rewind [filename]

Clears currently recorded trace data.

If filename is specified, new trace file is made and any newly collected data will be written there.

htrace print [start [len]]

Prints trace buffer, using current record configuration.

htrace mode continuous

Set continuous trace mode.

htrace mode suspend

Set suspend trace mode.

21.3.7 PowerPC Embedded

GDB provides the following PowerPC-specific commands:

```
set powerpc soft-float
show powerpc soft-float
```

Force GDB to use (or not use) a software floating point calling convention. By default, GDB selects the calling convention based on the selected architecture and the provided executable file.

```
set powerpc vector-abi show powerpc vector-abi
```

Force GDB to use the specified calling convention for vector arguments and return values. The valid options are 'auto'; 'generic', to avoid vector registers even if they are present; 'altivec', to use AltiVec registers; and 'spe' to use

SPE registers. By default, GDB selects the calling convention based on the selected architecture and the provided executable file.

target dink32 dev

DINK32 ROM monitor.

target ppcbug dev

target ppcbug1 dev

PPCBUG ROM monitor for PowerPC.

target sds dev

SDS monitor, running on a PowerPC board (such as Motorola's ADS).

The following commands specific to the SDS protocol are supported by GDB:

set sdstimeout nsec

Set the timeout for SDS protocol reads to be nsec seconds. The default is 2 seconds.

show sdstimeout

Show the current value of the SDS timeout.

sds command

Send the specified *command* string to the SDS monitor.

21.3.8 HP PA Embedded

target op50n dev

OP50N monitor, running on an OKI HPPA board.

target w89k dev

W89K monitor, running on a Winbond HPPA board.

21.3.9 Tsqware Sparclet

GDB enables developers to debug tasks running on Sparclet targets from a Unix host. GDB uses code that runs on both the Unix host and on the Sparclet target. The program gdb is installed and executed on the Unix host.

remotetimeout args

GDB supports the option remotetimeout. This option is set by the user, and args represents the number of seconds GDB waits for responses.

When compiling for debugging, include the options '-g' to get debug information and '-Ttext' to relocate the program to where you wish to load it on the target. You may also want to add the options '-n' or '-N' in order to reduce the size of the sections. Example:

```
sparclet-aout-gcc prog.c -Ttext 0x12010000 -g -o prog -N
```

You can use objdump to verify that the addresses are what you intended:

```
sparclet-aout-objdump --headers --syms prog
```

Once you have set your Unix execution search path to find gdb, you are ready to run gdb. From your Unix host, run gdb (or sparclet-aout-gdb, depending on your installation).

GDB comes up showing the prompt:

(gdbslet)

21.3.9.1 Setting File to Debug

The gdb command file lets you choose with program to debug.

```
(gdbslet) file prog
```

GDB then attempts to read the symbol table of 'prog'. GDB locates the file by searching the directories listed in the command search path. If the file was compiled with debug information (option '-g'), source files will be searched as well. GDB locates the source files by searching the directories listed in the directory search path (see Section 4.4 [Your Program's Environment], page 29). If it fails to find a file, it displays a message such as:

```
prog: No such file or directory.
```

When this happens, add the appropriate directories to the search paths with the GDB commands path and dir, and execute the target command again.

21.3.9.2 Connecting to Sparclet

The GDB command target lets you connect to a Sparclet target. To connect to a target on serial port "ttya", type:

```
(gdbslet) target sparclet /dev/ttya
Remote target sparclet connected to /dev/ttya
main () at ../prog.c:3
GDB displays messages like these:
Connected to ttya.
```

21.3.9.3 Sparclet Download

Once connected to the Sparclet target, you can use the GDB load command to download the file from the host to the target. The file name and load offset should be given as arguments to the load command. Since the file format is aout, the program must be loaded to the starting address. You can use objdump to find out what this value is. The load offset is an offset which is added to the VMA (virtual memory address) of each of the file's sections. For instance, if the program 'prog' was linked to text address 0x1201000, with data at 0x12010160 and bss at 0x12010170, in GDB, type:

```
(gdbslet) load prog 0x12010000
Loading section .text, size 0xdb0 vma 0x12010000
```

If the code is loaded at a different address then what the program was linked to, you may need to use the section and add-symbol-file commands to tell GDB where to map the symbol table.

21.3.9.4 Running and Debugging

You can now begin debugging the task using GDB's execution control commands, b, step, run, etc. See the GDB manual for the list of commands.

21.3.10 Fujitsu Sparclite

target sparclite dev

Fujitsu sparclite boards, used only for the purpose of loading. You must use an additional command to debug the program. For example: target remote dev using GDB standard remote protocol.

21.3.11 Zilog Z8000

When configured for debugging Zilog Z8000 targets, GDB includes a Z8000 simulator.

For the Z8000 family, 'target sim' simulates either the Z8002 (the unsegmented variant of the Z8000 architecture) or the Z8001 (the segmented variant). The simulator recognizes which architecture is appropriate by inspecting the object code.

target sim args

Debug programs on a simulated CPU. If the simulator supports setup options, specify them via args.

After specifying this target, you can debug programs for the simulated CPU in the same style as programs for your host computer; use the file command to load a new program image, the run command to run your program, and so on.

As well as making available all the usual machine registers (see Section 10.11 [Registers], page 110), the Z8000 simulator provides three additional items of information as specially named registers:

cycles Counts clock-ticks in the simulator.

insts Counts instructions run in the simulator.

time Execution time in 60ths of a second.

You can refer to these values in GDB expressions with the usual conventions; for example, 'b fputc if \$cycles>5000' sets a conditional breakpoint that suspends only after at least 5000 simulated clock ticks.

21.3.12 Atmel AVR

When configured for debugging the Atmel AVR, GDB supports the following AVR-specific commands:

info io_registers

This command displays information about the AVR I/O registers. For each register, GDB prints its number and value.

21.3.13 CRIS

When configured for debugging CRIS, GDB provides the following CRIS-specific commands:

set cris-version ver

Set the current CRIS version to ver, either '10' or '32'. The CRIS version affects register names and sizes. This command is useful in case autodetection of the CRIS version fails.

show cris-version

Show the current CRIS version.

set cris-dwarf2-cfi

Set the usage of DWARF-2 CFI for CRIS debugging. The default is 'on'. Change to 'off' when using gcc-cris whose version is below R59.

show cris-dwarf2-cfi

Show the current state of using DWARF-2 CFI.

set cris-mode mode

Set the current CRIS mode to *mode*. It should only be changed when debugging in guru mode, in which case it should be set to 'guru' (the default is 'normal').

show cris-mode

Show the current CRIS mode.

21.3.14 Renesas Super-H

For the Renesas Super-H processor, GDB provides these commands:

regs Show the values of all Super-H registers.

set sh calling-convention convention

Set the calling-convention used when calling functions from GDB. Allowed values are 'gcc', which is the default setting, and 'renesas'. With the 'gcc' setting, functions are called using the GCC calling convention. If the DWARF-2 information of the called function specifies that the function follows the Renesas calling convention, the function is called using the Renesas calling convention. If the calling convention is set to 'renesas', the Renesas calling convention is always used, regardless of the DWARF-2 information. This can be used to override the default of 'gcc' if debug information is missing, or the compiler does not emit the DWARF-2 calling convention entry for a function.

show sh calling-convention

Show the current calling convention setting.

21.4 Architectures

This section describes characteristics of architectures that affect all uses of gdb with the architecture, both native and cross.

21.4.1 x86 Architecture-specific Issues

set struct-convention mode

Set the convention used by the inferior to return structs and unions from functions to mode. Possible values of mode are "pcc", "reg", and "default" (the default). "default" or "pcc" means that structs are returned on the stack, while "reg" means that a struct or a union whose size is 1, 2, 4, or 8 bytes will be returned in a register.

show struct-convention

Show the current setting of the convention to return structs from functions.

21.4.2 A29K

set rstack_high_address address

On AMD 29000 family processors, registers are saved in a separate register stack. There is no way for GDB to determine the extent of this stack. Normally, GDB just assumes that the stack is "large enough". This may result in GDB referencing memory locations that do not exist. If necessary, you can get around this problem by specifying the ending address of the register stack with the set rstack_high_address command. The argument should be an address, which you probably want to precede with '0x' to specify in hexadecimal.

show rstack_high_address

Display the current limit of the register stack, on AMD 29000 family processors.

21.4.3 Alpha

See the following section.

21.4.4 MIPS

Alpha- and MIPS-based computers use an unusual stack frame, which sometimes requires GDB to search backward in the object code to find the beginning of a function.

To improve response time (especially for embedded applications, where gdb may be restricted to a slow serial line for this search) you may want to limit the size of this search, using one of these commands:

set heuristic-fence-post limit

Restrict GDB to examining at most *limit* bytes in its search for the beginning of a function. A value of θ (the default) means there is no limit. However, except for θ , the larger the limit the more bytes heuristic-fence-post must search and therefore the longer it takes to run. You should only need to use this command when debugging a stripped executable.

show heuristic-fence-post

Display the current limit.

These commands are available *only* when gdb is configured for debugging programs on Alpha or MIPS processors.

Several MIPS-specific commands are available when debugging MIPS programs:

set mips abi arg

Tell gdb which MIPS ABI is used by the inferior. Possible values of arg are:

'auto' The default ABI associated with the current binary (this is the default).

'o32'

'o64'

'n32'

'n64'

'eabi32'

'eabi64'

show mips abi

Show the MIPS ABI used by GDB to debug the inferior.

set mips compression arg

Tell gdb which MIPS compressed ISA encoding is used by the inferior. gdb uses this for code disassembly and other internal interpretation purposes. This setting is only referred to when no executable has been associated with the debugging session or the executable does not provide information about the encoding it uses. Otherwise this setting is automatically updated from information provided by the executable.

Possible values of arg are 'mips16' and 'micromips'. The default compressed ISA encoding is 'mips16', as executables containing MIPS16 code frequently are not identified as such.

This setting is "sticky"; that is, it retains its value across debugging sessions until reset either explicitly with this command or implicitly from an executable.

The compiler and/or assembler typically add symbol table annotations to identify functions compiled for the MIPS16 or microMIPS ISAs. If these functionscope annotations are present, GDB uses them in preference to the global compressed ISA encoding setting.

show mips compression

Show the MIPS compressed ISA encoding used by GDB to debug the inferior.

set mipsfpu show mipsfpu

See Section 21.3.5 [MIPS Embedded], page 232.

set mips mask-address arg

This command determines whether the most-significant 32 bits of 64-bit MIPS addresses are masked off. The argument arg can be 'on', 'off', or 'auto'. The latter is the default setting, which lets GDB determine the correct value.

show mips mask-address

Show whether the upper 32 bits of MIPS addresses are masked off or not.

set remote-mips64-transfers-32bit-regs

This command controls compatibility with 64-bit MIPS targets that transfer data in 32-bit quantities. If you have an old MIPS 64 target that transfers 32 bits for some registers, like SR and FSR, and 64 bits for other registers, set this option to 'on'.

show remote-mips64-transfers-32bit-regs

Show the current setting of compatibility with older MIPS 64 targets.

set debug mips

This command turns on and off debugging messages for the MIPS-specific target code in GDB.

show debug mips

Show the current setting of MIPS debugging messages.

21.4.5 HPPA

When gdb is debugging the HP PA architecture, it provides the following special commands:

set debug hppa

This command determines whether HPPA architecture-specific debugging messages are to be displayed.

show debug hppa

Show whether HPPA debugging messages are displayed.

maint print unwind address

This command displays the contents of the unwind table entry at the given address.

21.4.6 Cell Broadband Engine SPU architecture

When gdb is debugging the Cell Broadband Engine SPU architecture, it provides the following special commands:

info spu event

Display SPU event facility status. Shows current event mask and pending event status.

info spu signal

Display SPU signal notification facility status. Shows pending signal-control word and signal notification mode of both signal notification channels.

info spu mailbox

Display SPU mailbox facility status. Shows all pending entries, in order of processing, in each of the SPU Write Outbound, SPU Write Outbound Interrupt, and SPU Read Inbound mailboxes.

info spu dma

Display MFC DMA status. Shows all pending commands in the MFC DMA queue. For each entry, opcode, tag, class IDs, effective and local store addresses and transfer size are shown.

info spu proxydma

Display MFC Proxy-DMA status. Shows all pending commands in the MFC Proxy-DMA queue. For each entry, opcode, tag, class IDs, effective and local store addresses and transfer size are shown.

When GDB is debugging a combined PowerPC/SPU application on the Cell Broadband Engine, it provides in addition the following special commands:

set spu stop-on-load arg

Set whether to stop for new SPE threads. When set to on, GDB will give control to the user when a new SPE thread enters its main function. The default is off.

show spu stop-on-load

Show whether to stop for new SPE threads.

set spu auto-flush-cache arg

Set whether to automatically flush the software-managed cache. When set to on, GDB will automatically cause the SPE software-managed cache to be flushed whenever SPE execution stops. This provides a consistent view of PowerPC memory that is accessed via the cache. If an application does not use the software-managed cache, this option has no effect.

show spu auto-flush-cache

Show whether to automatically flush the software-managed cache.

21.4.7 PowerPC

When GDB is debugging the PowerPC architecture, it provides a set of pseudo-registers to enable inspection of 128-bit wide Decimal Floating Point numbers stored in the floating point registers. These values must be stored in two consecutive registers, always starting at an even register like f0 or f2.

The pseudo-registers go from \$d10 through \$d115, and are formed by joining the even/odd register pairs f0 and f1 for \$d10, f2 and f3 for \$d11 and so on.

For POWER7 processors, GDB provides a set of pseudo-registers, the 64-bit wide Extended Floating Point Registers ('f32' through 'f63').

22 Controlling GDB

You can alter the way GDB interacts with you by using the set command. For commands controlling how GDB displays data, see Section 10.8 [Print Settings], page 102. Other settings are described here.

22.1 Prompt

GDB indicates its readiness to read a command by printing a string called the *prompt*. This string is normally '(gdb)'. You can change the prompt string with the set prompt command. For instance, when debugging GDB with GDB, it is useful to change the prompt in one of the GDB sessions so that you can always tell which one you are talking to.

Note: set prompt does not add a space for you after the prompt you set. This allows you to set a prompt which ends in a space or a prompt that does not.

```
set prompt newprompt
```

Directs gdb to use newprompt as its prompt string henceforth.

show prompt

Prints a line of the form: 'Gdb's prompt is: your-prompt'

22.2 Command Editing

GDB reads its input commands via the *Readline* interface. This GNU library provides consistent behavior for programs which provide a command line interface to the user. Advantages are GNU Emacs-style or *vi*-style inline editing of commands, **csh**-like history substitution, and a storage and recall of command history across debugging sessions.

You may control the behavior of command line editing in GDB with the command set.

```
set editing set editing on
```

Enable command line editing (enabled by default).

set editing off

Disable command line editing.

show editing

Show whether command line editing is enabled.

See Chapter 31 [Command Line Editing], page 365, for more details about the Readline interface. Users unfamiliar with GNU Emacs or vi are encouraged to read that chapter.

22.3 Command History

GDB can keep track of the commands you type during your debugging sessions, so that you can be certain of precisely what happened. Use these commands to manage the GDB command history facility.

GDB uses the GNU History library, a part of the Readline package, to provide the history facility. See Chapter 32 [Using History Interactively], page 385, for the detailed description of the History library.

To issue a command to GDB without affecting certain aspects of the state which is seen by users, prefix it with 'server' (see Section 28.2 [Server Prefix], page 356). This means that this command will not affect the command history, nor will it affect GDB's notion of which command to repeat if RET is pressed on a line by itself.

The server prefix does not affect the recording of values into the value history; to print a value without recording it into the value history, use the output command instead of the print command.

Here is the description of gdb commands related to command history.

set history filename fname

Set the name of the GDB command history file to fname. This is the file where GDB reads an initial command history list, and where it writes the command history from this session when it exits. You can access this list through history expansion or through the history command editing characters listed below. This file defaults to the value of the environment variable GDBHISTFILE, or to './.gdb_history' ('./_gdb_history' on MS-DOS) if this variable is not set.

set history save on

Record command history in a file, whose name may be specified with the set history filename command. By default, this option is disabled.

set history save off

Stop recording command history in a file.

set history size size

Set the number of commands which GDB keeps in its history list. This defaults to the value of the environment variable HISTSIZE, or to 256 if this variable is not set.

History expansion assigns special meaning to the character !. See Section 32.1.1 [Event Designators], page 385, for more details.

Since ! is also the logical not operator in C, history expansion is off by default. If you decide to enable history expansion with the set history expansion on command, you may sometimes need to follow ! (when it is used as logical not, in an expression) with a space or a tab to prevent it from being expanded. The readline history facilities do not attempt substitution on the strings != and ! (, even when history expansion is enabled.

The commands to control history expansion are:

set history expansion on set history expansion

Enable history expansion. History expansion is off by default.

set history expansion off

Disable history expansion.

show history filename show history save show history size show history expansion

These commands display the state of the GDB history parameters. show history by itself displays all four states.

show commands

Display the last ten commands in the command history.

show commands n

Print ten commands centered on command number n.

show commands +

Print ten commands just after the commands last printed.

22.4 Screen Size

Certain commands to GDB may produce large amounts of information output to the screen. To help you read all of it, GDB pauses and asks you for input at the end of each page of output. Type RET when you want to continue the output, or q to discard the remaining output. Also, the screen width setting determines when to wrap lines of output. Depending on what is being printed, GDB tries to break the line at a readable place, rather than simply letting it overflow onto the following line.

Normally GDB knows the size of the screen from the terminal driver software. For example, on Unix GDB uses the termcap data base together with the value of the TERM environment variable and the stty rows and stty cols settings. If this is not correct, you can override it with the set height and set width commands:

set height *lpp* show height set width *cpl* show width

These **set** commands specify a screen height of *lpp* lines and a screen width of *cpl* characters. The associated **show** commands display the current settings.

If you specify a height of zero lines, GDB does not pause during output no matter how long the output is. This is useful if output is to a file or to an editor buffer.

Likewise, you can specify 'set width 0' to prevent GDB from wrapping its output.

set pagination on set pagination off

Turn the output pagination on or off; the default is on. Turning pagination off is the alternative to set height 0.

show pagination

Show the current pagination mode.

22.5 Numbers

You can always enter numbers in octal, decimal, or hexadecimal in GDB by the usual conventions: octal numbers begin with '0', decimal numbers end with '.', and hexadecimal numbers begin with '0x'. Numbers that neither begin with '0' or '0x', nor end with a '.' are, by default, entered in base 10; likewise, the default display for numbers—when no particular format is specified—is base 10. You can change the default base for both input and output with the commands described below.

set input-radix base

Set the default base for numeric input. Supported choices for base are decimal 8, 10, or 16. base must itself be specified either unambiguously or using the current input radix; for example, any of

```
set input-radix 012
set input-radix 10.
set input-radix 0xa
```

sets the input base to decimal. On the other hand, 'set input-radix 10' leaves the input radix unchanged, no matter what it was, since '10', being without any leading or trailing signs of its base, is interpreted in the current radix. Thus, if the current radix is 16, '10' is interpreted in hex, i.e. as 16 decimal, which doesn't change the radix.

set output-radix base

Set the default base for numeric display. Supported choices for *base* are decimal 8, 10, or 16. *base* must itself be specified either unambiguously or using the current input radix.

show input-radix

Display the current default base for numeric input.

show output-radix

Display the current default base for numeric display.

```
set radix [base] show radix
```

These commands set and show the default base for both input and output of numbers. set radix sets the radix of input and output to the same base; without an argument, it resets the radix back to its default value of 10.

22.6 Configuring the Current ABI

GDB can determine the ABI (Application Binary Interface) of your application automatically. However, sometimes you need to override its conclusions. Use these commands to manage GDB's view of the current ABI.

One GDB configuration can debug binaries for multiple operating system targets, either via remote debugging or native emulation. GDB will autodetect the *OS ABI* (Operating System ABI) in use, but you can override its conclusion using the set osabi command. One example where this is useful is in debugging of binaries which use an alternate C library (e.g. uClibc for gnu/Linux) which does not have the same identifying marks that the standard C library for your platform provides.

show osabi

Show the OS ABI currently in use.

set osabi With no argument, show the list of registered available OS ABI's.

set osabi abi

Set the current OS ABI to abi.

Generally, the way that an argument of type float is passed to a function depends on whether the function is prototyped. For a prototyped (i.e. ANSI/ISO style) function, float arguments are passed unchanged, according to the architecture's convention for float. For unprototyped (i.e. K&R style) functions, float arguments are first promoted to type double and then passed.

Unfortunately, some forms of debug information do not reliably indicate whether a function is prototyped. If GDB calls a function that is not marked as prototyped, it consults set coerce-float-to-double.

```
set coerce-float-to-double
set coerce-float-to-double on
```

Arguments of type float will be promoted to double when passed to an unprototyped function. This is the default setting.

set coerce-float-to-double off

Arguments of type float will be passed directly to unprototyped functions.

show coerce-float-to-double

Show the current setting of promoting float to double.

GDB needs to know the ABI used for your program's C++ objects. The correct C++ ABI depends on which C++ compiler was used to build your application. GDB only fully supports programs with a single C++ ABI; if your program contains code using multiple C++ ABI's or if GDB can not identify your program's ABI correctly, you can tell GDB which ABI to use. Currently supported ABI's include "gnu-v2", for g++ versions before 3.0, "gnu-v3", for g++ versions 3.0 and later, and "hpaCC" for the HP ANSI C++ compiler. Other C++ compilers may use the "gnu-v2" or "gnu-v3" ABI's as well. The default setting is "auto".

show cp-abi

Show the C++ ABI currently in use.

set cp-abi

With no argument, show the list of supported C++ ABI's.

set cp-abi *abi* set cp-abi auto

Set the current C++ ABI to abi, or return to automatic detection.

22.7 Optional Warnings and Messages

By default, GDB is silent about its inner workings. If you are running on a slow machine, you may want to use the **set verbose** command. This makes GDB tell you when it does a lengthy internal operation, so you will not think it has crashed.

Currently, the messages controlled by set verbose are those which announce that the symbol table for a source file is being read; see symbol-file in Section 18.1 [Commands to Specify Files], page 187.

set verbose on

Enables gdb output of certain informational messages.

set verbose off

Disables GDB output of certain informational messages.

show verbose

Displays whether set verbose is on or off.

By default, if GDB encounters bugs in the symbol table of an object file, it is silent; but if you are debugging a compiler, you may find this information useful (see Section 18.3 [Errors Reading Symbol Files], page 197).

set complaints limit

Permits gdb to output *limit* complaints about each type of unusual symbols before becoming silent about the problem. Set *limit* to zero to suppress all complaints; set it to a large number to prevent complaints from being suppressed.

show complaints

Displays how many symbol complaints gdb is permitted to produce.

By default, GDB is cautious, and asks what sometimes seems to be a lot of stupid questions to confirm certain commands. For example, if you try to run a program which is already running:

```
(gdb) run
The program being debugged has been started already.
Start it from the beginning? (y or n)
```

If you are willing to unflinchingly face the consequences of your own commands, you can disable this "feature":

set confirm off

Disables confirmation requests.

set confirm on

Enables confirmation requests (the default).

show confirm

Displays state of confirmation requests.

If you need to debug user-defined commands or sourced files you may find it useful to enable *command tracing*. In this mode each command will be printed as it is executed, prefixed with one or more '+' symbols, the quantity denoting the call depth of each command.

set trace-commands on

Enable command tracing.

set trace-commands off

Disable command tracing.

show trace-commands

Display the current state of command tracing.

22.8 Optional Messages about Internal Happenings

GDB has commands that enable optional debugging messages from various GDB subsystems; normally these commands are of interest to GDB maintainers, or when reporting a bug. This section documents those commands.

set exec-done-display

Turns on or off the notification of asynchronous commands' completion. When on, GDB will print a message when an asynchronous command finishes its execution. The default is off.

show exec-done-display

Displays the current setting of asynchronous command completion notification.

set debug arch

Turns on or off display of gdbarch debugging info. The default is off

show debug arch

Displays the current state of displaying gdbarch debugging info.

set debug aix-thread

Display debugging messages about inner workings of the AIX thread module.

show debug aix-thread

Show the current state of AIX thread debugging info display.

set debug dwarf2-die

Dump DWARF2 DIEs after they are read in. The value is the number of nesting levels to print. A value of zero turns off the display.

show debug dwarf2-die

Show the current state of DWARF2 DIE debugging.

set debug displaced

Turns on or off display of gdb debugging info for the displaced stepping support. The default is off.

show debug displaced

Displays the current state of displaying gdb debugging info related to displaced stepping.

set debug event

Turns on or off display of GDB event debugging info. The default is off.

show debug event

Displays the current state of displaying gdb event debugging info.

set debug expression

Turns on or off display of debugging info about GDB expression parsing. The default is off.

show debug expression

Displays the current state of displaying debugging info about GDB expression parsing.

set debug frame

Turns on or off display of GDB frame debugging info. The default is off.

show debug frame

Displays the current state of displaying GDB frame debugging info.

set debug gnu-nat

Turns on or off debugging messages from the GNU/Hurd debug support.

show debug gnu-nat

Show the current state of GNU/Hurd debugging messages.

set debug infrun

Turns on or off display of GDB debugging info for running the inferior. The default is off. 'infrun.c' contains GDB's runtime state machine used for implementing operations such as single-stepping the inferior.

show debug infrun

Displays the current state of GDB inferior debugging.

set debug lin-lwp

Turns on or off debugging messages from the Linux LWP debug support.

show debug lin-lwp

Show the current state of Linux LWP debugging messages.

set debug lin-lwp-async

Turns on or off debugging messages from the Linux LWP async debug support.

show debug lin-lwp-async

Show the current state of Linux LWP async debugging messages.

set debug observer

Turns on or off display of GDB observer debugging. This includes info such as the notification of observable events.

show debug observer

Displays the current state of observer debugging.

set debug overload

Turns on or off display of gdb C++ overload debugging info. This includes info such as ranking of functions, etc. The default is off.

show debug overload

Displays the current state of displaying gdb C++ overload debugging info.

set debug parser

Turns on or off the display of expression parser debugging output. Internally, this sets the yydebug variable in the expression parser. See Section "Tracing Your Parser" in *Bison*, for details. The default is off.

show debug parser

Show the current state of expression parser debugging.

set debug remote

Turns on or off display of reports on all packets sent back and forth across the serial line to the remote machine. The info is printed on the gdb standard output stream. The default is off.

show debug remote

Displays the state of display of remote packets.

set debug serial

Turns on or off display of GDB serial debugging info. The default is off.

show debug serial

Displays the current state of displaying gdb serial debugging info.

set debug solib-frv

Turns on or off debugging messages for FR-V shared-library code.

show debug solib-frv

Display the current state of FR-V shared-library code debugging messages.

set debug target

Turns on or off display of GDB target debugging info. This info includes what is going on at the target level of GDB, as it happens. The default is 0. Set it to 1 to track events, and to 2 to also track the value of large memory transfers. Changes to this flag do not take effect until the next time you connect to a target or use the run command.

show debug target

Displays the current state of displaying GDB target debugging info.

set debug timestamp

Turns on or off display of timestamps with gdb debugging info. When enabled, seconds and microseconds are displayed before each debugging message.

show debug timestamp

Displays the current state of displaying timestamps with gdb debugging info.

set debugvarobj

Turns on or off display of GDB variable object debugging info. The default is off.

show debugvarobj

Displays the current state of displaying gdb variable object debugging info.

set debug xml

Turns on or off debugging messages for built-in XML parsers.

show debug xml

Displays the current state of XML debugging messages.

22.9 Other Miscellaneous Settings

set interactive-mode

If on, forces GDB to operate interactively. If off, forces GDB to operate non-interactively, If auto (the default), GDB guesses which mode to use, based on whether the debugger was started in a terminal or not.

In the vast majority of cases, the debugger should be able to guess correctly which mode should be used. But this setting can be useful in certain specific cases, such as running a MinGW gdb inside a cygwin window.

show interactive-mode

Displays whether the debugger is operating in interactive mode or not.

23 Extending GDB

GDB provides two mechanisms for extension. The first is based on composition of GDB commands, and the second is based on the Python scripting language.

To facilitate the use of these extensions, GDB is capable of evaluating the contents of a file. When doing so, GDB can recognize which scripting language is being used by looking at the filename extension. Files with an unrecognized filename extension are always treated as a GDB Command Files. See Section 23.1.3 [Command files], page 258.

You can control how GDB evaluates these files with the following setting:

set script-extension off

All scripts are always evaluated as GDB Command Files.

set script-extension soft

The debugger determines the scripting language based on filename extension. If this scripting language is supported, GDB evaluates the script using that language. Otherwise, it evaluates the file as a GDB Command File.

set script-extension strict

The debugger determines the scripting language based on filename extension, and evaluates the script using that language. If the language is not supported, then the evaluation fails.

show script-extension

Display the current value of the script-extension option.

23.1 Canned Sequences of Commands

Aside from breakpoint commands (see Section 5.1.7 [Breakpoint Command Lists], page 58), GDB provides two ways to store sequences of commands for execution as a unit: user-defined commands and command files.

23.1.1 User-defined Commands

A user-defined command is a sequence of GDB commands to which you assign a new name as a command. This is done with the define command. User commands may accept up to 10 arguments separated by whitespace. Arguments are accessed within the user command via \$arg0...\$arg9. A trivial example:

```
define adder
  print $arg0 + $arg1 + $arg2
end
```

To execute the command use:

```
adder 1 2 3
```

This defines the command adder, which prints the sum of its three arguments. Note the arguments are text substitutions, so they may reference variables, use complex expressions, or even perform inferior functions calls.

In addition, \$argc may be used to find out how many arguments have been passed. This expands to a number in the range 0...10.

```
define adder
  if $argc == 2
    print $arg0 + $arg1
  end
  if $argc == 3
    print $arg0 + $arg1 + $arg2
  end
end
```

define commandname

Define a command named *commandname*. If there is already a command by that name, you are asked to confirm that you want to redefine it. *commandname* may be a bare command name consisting of letters, numbers, dashes, and underscores. It may also start with any predefined prefix command. For example, 'define target my-target' creates a user-defined 'target my-target' command.

The definition of the command is made up of other GDB command lines, which are given following the define command. The end of these commands is marked by a line containing end.

document commandname

Document the user-defined command commandname, so that it can be accessed by help. The command commandname must already be defined. This command reads lines of documentation just as define reads the lines of the command definition, ending with end. After the document command is finished, help on command commandname displays the documentation you have written.

You may use the document command again to change the documentation of a command. Redefining the command with define does not change the documentation.

dont-repeat

Used inside a user-defined command, this tells GDB that this command should not be repeated when the user hits RET (see Section 3.1 [Command Syntax], page 19).

help user-defined

List all user-defined commands, with the first line of the documentation (if any) for each.

show user

show user commandname

Display the GDB commands used to define *commandname* (but not its documentation). If no *commandname* is given, display the definitions for all user-defined commands.

```
show max-user-call-depth set max-user-call-depth
```

The value of max-user-call-depth controls how many recursion levels are allowed in user-defined commands before GDB suspects an infinite recursion and aborts the command.

In addition to the above commands, user-defined commands frequently use control flow commands, described in Section 23.1.3 [Command Files], page 258.

When user-defined commands are executed, the commands of the definition are not printed. An error in any command stops execution of the user-defined command.

If used interactively, commands that would ask for confirmation proceed without asking when used inside a user-defined command. Many GDB commands that normally print messages to say what they are doing omit the messages when used in a user-defined command.

23.1.2 User-defined Command Hooks

You may define *hooks*, which are a special kind of user-defined command. Whenever you run the command 'foo', if the user-defined command 'hook-foo' exists, it is executed (with no arguments) before that command.

A hook may also be defined which is run after the command you executed. Whenever you run the command 'foo', if the user-defined command 'hookpost-foo' exists, it is executed (with no arguments) after that command. Post-execution hooks may exist simultaneously with pre-execution hooks, for the same command.

It is valid for a hook to call the command which it hooks. If this occurs, the hook is not re-executed, thereby avoiding infinite recursion.

In addition, a pseudo-command, 'stop' exists. Defining ('hook-stop') makes the associated commands execute every time execution stops in your program: before breakpoint commands are run, displays are printed, or the stack frame is printed.

For example, to ignore SIGALRM signals while single-stepping, but treat them normally during normal execution, you could define:

```
define hook-stop
handle SIGALRM nopass
end

define hook-run
handle SIGALRM pass
end

define hook-continue
handle SIGALRM pass
end
```

As a further example, to hook at the beginning and end of the echo command, and to add extra text to the beginning and end of the message, you could define:

```
define hook-echo
echo <<<---
end

define hookpost-echo
echo --->>>\n
end

(gdb) echo Hello World
<<<---Hello World--->>>
(gdb)
```

You can define a hook for any single-word command in GDB, but not for command aliases; you should define a hook for the basic command name, e.g. backtrace rather than

bt. You can hook a multi-word command by adding hook- or hookpost- to the last word of the command, e.g. 'define target hook-remote' to add a hook to 'target remote'.

If an error occurs during the execution of your hook, execution of gdb commands stops and gdb issues a prompt (before the command that you actually typed had a chance to run).

If you try to define a hook which does not match any known command, you get a warning from the define command.

23.1.3 Command Files

A command file for GDB is a text file made of lines that are GDB commands. Comments (lines starting with #) may also be included. An empty line in a command file does nothing; it does not mean to repeat the last command, as it would from the terminal.

You can request the execution of a command file with the source command. Note that the source command is also used to evaluate scripts that are not Command Files. The exact behavior can be configured using the script-extension setting. See Chapter 23 [Extending GDB], page 255.

source [-v] filename

Execute the command file filename.

The lines in a command file are generally executed sequentially, unless the order of execution is changed by one of the *flow-control commands* described below. The commands are not printed as they are executed. An error in any command terminates execution of the command file and control is returned to the console.

GDB searches for *filename* in the current directory and then on the search path (specified with the 'directory' command).

If -v, for verbose mode, is given then GDB displays each command as it is executed. The option must be given before *filename*, and is interpreted as part of the filename anywhere else.

Commands that would ask for confirmation if used interactively proceed without asking when used in a command file. Many GDB commands that normally print messages to say what they are doing omit the messages when called from command files.

GDB also accepts command input from standard input. In this mode, normal output goes to standard output and error output goes to standard error. Errors in a command file supplied on standard input do not terminate execution of the command file—execution continues with the next command.

```
gdb < cmds > log 2>&1
```

(The syntax above will vary depending on the shell used.) This example will execute commands from the file 'cmds'. All output and errors would be directed to 'log'.

Since commands stored on command files tend to be more general than commands typed interactively, they frequently need to deal with complicated situations, such as different or unexpected values of variables and symbols, changes in how the program being debugged is built, etc. GDB provides a set of flow-control commands to deal with these complexities. Using these commands, you can write complex scripts that loop over data structures, execute commands conditionally, etc.

if

This command allows to include in your script conditionally executed commands. The if command takes a single argument, which is an expression to evaluate. It is followed by a series of commands that are executed only if the expression is true (its value is nonzero). There can then optionally be an else line, followed by a series of commands that are only executed if the expression was false. The end of the list is marked by a line containing end.

while This command allows to write loops. Its syntax is similar to if: the command takes a single argument, which is an expression to evaluate, and must be followed by the commands to execute, one per line, terminated by an end. These commands are called the *body* of the loop. The commands in the body of while are executed repeatedly as long as the expression evaluates to true.

loop_break

This command exits the while loop in whose body it is included. Execution of the script continues after that whiles end line.

loop_continue

This command skips the execution of the rest of the body of commands in the while loop in whose body it is included. Execution branches to the beginning of the while loop, where it evaluates the controlling expression.

end Terminate the block of commands that are the body of if, else, or while flow-control commands.

23.1.4 Commands for Controlled Output

During the execution of a command file or a user-defined command, normal GDB output is suppressed; the only output that appears is what is explicitly printed by the commands in the definition. This section describes three commands useful for generating exactly the output you want.

echo text

Print text. Nonprinting characters can be included in text using C escape sequences, such as '\n' to print a newline. No newline is printed unless you specify one. In addition to the standard C escape sequences, a backslash followed by a space stands for a space. This is useful for displaying a string with spaces at the beginning or the end, since leading and trailing spaces are otherwise trimmed from all arguments. To print ' and foo = ', use the command 'echo \ and foo = \'.

A backslash at the end of text can be used, as in C, to continue the command onto subsequent lines. For example,

echo This is some text\n\
which is continued\n\
onto several lines.\n

produces the same output as

echo This is some text\n echo which is continued\n echo onto several lines.\n

output expression

Print the value of expression and nothing but that value: no newlines, no '\$nn = '. The value is not entered in the value history either. See Section 10.1 [Expressions], page 93, for more information on expressions.

output/fmt expression

Print the value of expression in format fmt. You can use the same formats as for print. See Section 10.5 [Output Formats], page 97, for more information.

printf template, expressions...

Print the values of one or more expressions under the control of the string template. To print several values, make expressions be a comma-separated list of individual expressions, which may be either numbers or pointers. Their values are printed as specified by template, exactly as a C program would do by executing the code below:

```
printf (template, expressions...);
```

As in C printf, ordinary characters in *template* are printed verbatim, while conversion specification introduced by the '%' character cause subsequent expressions to be evaluated, their values converted and formatted according to type and style information encoded in the conversion specifications, and then printed.

For example, you can print two values in hex like this:

```
printf "foo, bar-foo = 0x\%x, 0x\%x\n", foo, bar-foo
```

printf supports all the standard C conversion specifications, including the flags and modifiers between the '%' character and the conversion letter, with the following exceptions:

- The argument-ordering modifiers, such as '2\$', are not supported.
- The modifier '*' is not supported for specifying precision or width.
- The '' flag (for separation of digits into groups according to LC_NUMERIC') is not supported.
- The type modifiers 'hh', 'j', 't', and 'z' are not supported.
- The conversion letter 'n' (as in '%n') is not supported.
- The conversion letters 'a' and 'A' are not supported.

Note that the '11' type modifier is supported only if the underlying C implementation used to build GDB supports the long long int type, and the 'L' type modifier is supported only if long double type is available.

As in C, printf supports simple backslash-escape sequences, such as \n, '\t', '\\', '\"', '\a', and '\f', that consist of backslash followed by a single character. Octal and hexadecimal escape sequences are not supported.

Additionally, **printf** supports conversion specifications for DFP (*Decimal Floating Point*) types using the following length modifiers together with a floating point specifier. letters:

- 'H' for printing Decimal32 types.
- 'D' for printing Decimal64 types.
- 'DD' for printing Decimal128 types.

If the underlying C implementation used to build GDB has support for the three length modifiers for DFP types, other modifiers such as width and precision will also be available for GDB to use.

In case there is no such C support, no additional modifiers will be available and the value will be printed in the standard way.

Here's an example of printing DFP types using the above conversion letters: printf "D32: %Hf - D64: %Df - D128: %DDf\n",1.2345df,1.2E10dd,1.2E1dl

23.2 Scripting gdb using Python

You can script GDB using the Python programming language. This feature is available only if GDB was configured using '--with-python'.

23.2.1 Python Commands

GDB provides one command for accessing the Python interpreter, and one related setting:

python [code]

The python command can be used to evaluate Python code.

If given an argument, the python command will evaluate the argument as a Python command. For example:

```
(gdb) python print 23
23
```

If you do not provide an argument to python, it will act as a multi-line command, like define. In this case, the Python script is made up of subsequent command lines, given after the python command. This command list is terminated using a line containing end. For example:

```
(gdb) python
Type python script
End with a line saying just "end".
>print 23
>end
23
```

maint set python print-stack

By default, GDB will print a stack trace when an error occurs in a Python script. This can be controlled using maint set python print-stack: if on, the default, then Python stack printing is enabled; if off, then Python stack printing is disabled.

It is also possible to execute a Python script from the GDB interpreter:

source 'script-name'

The script name must end with '.py' and GDB must be configured to recognize the script language based on filename extension using the script-extension setting. See Chapter 23 [Extending GDB], page 255.

```
python execfile ("script-name")
```

This method is based on the execfile Python built-in function, and thus is always available.

23.2.2 Python API

At startup, GDB overrides Python's sys.stdout and sys.stderr to print using GDB's output-paging streams. A Python program which outputs to one of these streams may have its output interrupted by the user (see Section 22.4 [Screen Size], page 247). In this situation, a Python KeyboardInterrupt exception is thrown.

23.2.2.1 Basic Python

GDB introduces a new Python module, named gdb. All methods and classes added by GDB are placed in this module. GDB automatically imports the gdb module for use in all scripts evaluated by the python command.

execute command [from_tty]

[Function]

Evaluate *command*, a string, as a GDB CLI command. If a GDB exception happens while *command* runs, it is translated as described in Section 23.2.2.2 [Exception Handling], page 263. If no exceptions occur, this function returns None.

from_tty specifies whether GDB ought to consider this command as having originated from the user invoking it interactively. It must be a boolean value. If omitted, it defaults to False.

parameter parameter

[Function]

Return the value of a GDB parameter. parameter is a string naming the parameter to look up; parameter may contain spaces if the parameter has a multi-part name. For example, 'print object' is a valid parameter name.

If the named parameter does not exist, this function throws a RuntimeError. Otherwise, the parameter's value is converted to a Python value of the appropriate type, and returned.

history number

[Function]

Return a value from GDB's value history (see Section 10.9 [Value History], page 108). number indicates which history element to return. If number is negative, then GDB will take its absolute value and count backward from the last element (i.e., the most recent element) to find the value to return. If number is zero, then GDB will return the most recent element. If the element specified by number doesn't exist in the value history, a RuntimeError exception will be raised.

If no exception is raised, the return value is always an instance of gdb.Value (see Section 23.2.2.4 [Values From Inferior], page 264).

parse_and_eval expression

[Function]

Parse expression as an expression in the current language, evaluate it, and return the result as a gdb. Value. expression must be a string.

This function can be useful when implementing a new command (see Section 23.2.2.8 [Commands In Python], page 272), as it provides a way to parse the command's argument as an expression. It is also useful simply to compute values, for example, it is the only way to get the value of a convenience variable (see Section 10.10 [Convenience Vars], page 109) as a gdb.Value.

write string [Function]

Print a string to GDB's paginated standard output stream. Writing to sys.stdout or sys.stderr will automatically call this function.

flush [Function]

Flush GDB's paginated standard output stream. Flushing sys.stdout or sys.stderr will automatically call this function.

23.2.2.2 Exception Handling

When executing the python command, Python exceptions uncaught within the Python code are translated to calls to GDB error-reporting mechanism. If the command that called python does not handle the error, GDB will terminate it and print an error message containing the Python exception name, the associated value, and the Python call stack backtrace at the point where the exception was raised. Example:

```
(gdb) python print foo
Traceback (most recent call last):
   File "<string>", line 1, in <module>
NameError: name 'foo' is not defined
```

GDB errors that happen in GDB commands invoked by Python code are converted to Python RuntimeError exceptions. User interrupt (via C-c or by typing q at a pagination prompt) is translated to a Python KeyboardInterrupt exception. If you catch these exceptions in your Python code, your exception handler will see RuntimeError or KeyboardInterrupt as the exception type, the GDB error message as its value, and the Python call stack backtrace at the Python statement closest to where the GDB error occurred as the traceback.

23.2.2.3 Auto-loading

When a new object file is read (for example, due to the file command, or because the inferior has loaded a shared library), GDB will look for a file named 'objfile-gdb.py', where objfile is the object file's real name, formed by ensuring that the file name is absolute, following all symlinks, and resolving . and . . components. If this file exists and is readable, GDB will evaluate it as a Python script.

If this file does not exist, and if the parameter debug-file-directory is set (see Section 18.2 [Separate Debug Files], page 193), then GDB will use for its each separated directory component component the file named 'component/real-name', where real-name is the object file's real name, as described above.

Finally, if this file does not exist, then GDB will look for a file named 'data-directory/python/auto-load/real-name', where data-directory is GDB's data directory (available via show data-directory, see Section 18.4 [Data Files], page 198), and real-name is the object file's real name, as described above.

When reading an auto-loaded file, GDB sets the "current objfile". This is available via the gdb.current_objfile function (see Section 23.2.2.10 [Objfiles In Python], page 276). This can be useful for registering objfile-specific pretty-printers.

The auto-loading feature is useful for supplying application-specific debugging commands and scripts. You can enable or disable this feature, and view its current state.

```
maint set python auto-load [yes|no]
```

Enable or disable the Python auto-loading feature.

show python auto-load

Show whether Python auto-loading is enabled or disabled.

GDB does not track which files it has already auto-loaded. So, your '-gdb.py' file should take care to ensure that it may be evaluated multiple times without error.

23.2.2.4 Values From Inferior

GDB provides values it obtains from the inferior program in an object of type gdb.Value. GDB uses this object for its internal bookkeeping of the inferior's values, and for fetching values when necessary.

Inferior values that are simple scalars can be used directly in Python expressions that are valid for the value's data type. Here's an example for an integer or floating-point value some_val:

```
bar = some val + 2
```

As result of this, bar will also be a gdb. Value object whose values are of the same type as those of some_val.

Inferior values that are structures or instances of some class can be accessed using the Python dictionary syntax. For example, if some_val is a gdb.Value instance holding a structure, you can access its foo element with:

```
bar = some_val['foo']
```

Again, bar will also be a gdb. Value object.

The following attributes are provided:

address

[Instance Variable of Value]

If this object is addressable, this read-only attribute holds a gdb.Value object representing the address. Otherwise, this attribute holds None.

is_optimized_out

[Instance Variable of Value]

This read-only boolean attribute is true if the compiler optimized out this value, thus it is not available for fetching from the inferior.

type

[Instance Variable of Value]

The type of this gdb.Value. The value of this attribute is a gdb.Type object.

The following methods are provided:

cast type

[Method on Value]

Return a new instance of gdb.Value that is the result of casting this instance to the type described by type, which must be a gdb.Type object. If the cast cannot be performed for some reason, this method throws an exception.

dereference

[Method on Value]

For pointer data types, this method returns a new gdb. Value object whose contents is the object pointed to by the pointer. For example, if foo is a C pointer to an int, declared in your C program as

int *foo;

then you can use the corresponding gdb. Value to access what foo points to like this:

bar = foo.dereference ()

The result bar will be a gdb. Value object holding the value pointed to by foo.

string [encoding] [errors] [length]

[Method on Value]

If this gdb.Value represents a string, then this method converts the contents to a Python string. Otherwise, this method will throw an exception. Strings are recognized in a language-specific way; whether a given gdb.Value represents a string is determined by the current language.

For C-like languages, a value is a string if it is a pointer to or an array of characters or ints. The string is assumed to be terminated by a zero of the appropriate width. However if the optional length argument is given, the string will be converted to that given length, ignoring any embedded zeros that the string may contain.

If the optional encoding argument is given, it must be a string naming the encoding of the string in the gdb.Value, such as "ascii", "iso-8859-6" or "utf-8". It accepts the same encodings as the corresponding argument to Python's string.decode method, and the Python codec machinery will be used to convert the string. If encoding is not given, or if encoding is the empty string, then either the target-charset (see Section 10.18 [Character Sets], page 116) will be used, or a language-specific encoding will be used, if the current language is able to supply one.

The optional *errors* argument is the same as the corresponding argument to Python's string.decode method.

If the optional *length* argument is given, the string will be fetched and converted to the given length.

lazy_string [encoding] [length]

[Method on Value]

If this gdb.Value represents a string, then this method converts the contents to a gdb.LazyString (see Section 23.2.2.12 [Lazy Strings In Python], page 277). Otherwise, this method will throw an exception.

If the optional *encoding* argument is given, it must be a string naming the encoding of the gdb.LazyString. Some examples are: 'ascii', 'iso-8859-6' or 'utf-8'. If the *encoding* argument is an encoding that GDB does recognize, GDB will raise an error.

When a lazy string is printed, the GDB encoding machinery is used to convert the string during printing. If the optional *encoding* argument is not provided, or is an empty string, GDB will automatically select the encoding most suitable for the string type. For further information on encoding in GDB please see Section 10.18 [Character Sets], page 116.

If the optional *length* argument is given, the string will be fetched and encoded to the length of characters specified. If the *length* argument

is not provided, the string will be fetched and encoded until a null of appropriate width is found.

23.2.2.5 Types In Python

GDB represents types from the inferior using the class gdb. Type.

The following type-related functions are available in the gdb module:

lookup_type name [block]

[Function]

This function looks up a type by name. *name* is the name of the type to look up. It must be a string.

Ordinarily, this function will return an instance of gdb.Type. If the named type cannot be found, it will throw an exception.

An instance of Type has the following attributes:

code

[Instance Variable of Type]

The type code for this type. The type code will be one of the TYPE_CODE_constants defined below.

sizeof

[Instance Variable of Type]

The size of this type, in target char units. Usually, a target's char type will be an 8-bit byte. However, on some unusual platforms, this type may have a different size.

tag

[Instance Variable of Type]

The tag name for this type. The tag name is the name after struct, union, or enum in C and C++; not all languages have this concept. If this type has no tag name, then None is returned.

The following methods are provided:

fields

[Method on Type]

For structure and union types, this method returns the fields. Range types have two fields, the minimum and maximum values. Enum types have one field per enum constant. Function and method types have one field per parameter. The base types of C++ classes are also represented as fields. If the type has no fields, or does not fit into one of these categories, an empty sequence will be returned.

Each field is an object, with some pre-defined attributes:

bitpos

This attribute is not available for static fields (as in C++ or Java). For non-static fields, the value is the bit position of the field.

name The name of the field, or None for anonymous fields.

artificial

This is True if the field is artificial, usually meaning that it was provided by the compiler and not the user. This attribute is always provided, and is False if the field is not artificial.

is_base_class

This is True if the field represents a base class of a C++ structure. This attribute is always provided, and is False if the field is not a base class of the type that is the argument of fields, or if that type was not a C++ class.

bitsize

If the field is packed, or is a bitfield, then this will have a non-zero value, which is the size of the field in bits. Otherwise, this will be zero; in this case the field's size is given by its type.

type The type of the field. This is usually an instance of Type, but it can be None in some situations.

const [Method on Type]

Return a new gdb. Type object which represents a const-qualified variant of this type.

volatile [Method on Type]

Return a new gdb.Type object which represents a volatile-qualified variant of this type.

unqualified [Method on Type]

Return a new gdb. Type object which represents an unqualified variant of this type. That is, the result is neither const nor volatile.

range [Method on Type]

Return a Python Tuple object that contains two elements: the low bound of the argument type and the high bound of that type. If the type does not have a range, GDB will raise a RuntimeError exception.

reference [Method on Type]

Return a new gdb. Type object which represents a reference to this type.

pointer [Method on Type]

Return a new gdb. Type object which represents a pointer to this type.

strip_typedefs

[Method on Type]

Return a new gdb. Type that represents the real type, after removing all layers of typedefs.

target [Method on Type]

Return a new gdb. Type object which represents the target type of this type.

For a pointer type, the target type is the type of the pointed-to object. For an array type (meaning C-like arrays), the target type is the type of the elements of the array. For a function or method type, the target type is the type of the return value. For a complex type, the target type is the type of the elements. For a typedef, the target type is the aliased type.

If the type does not have a target, this method will throw an exception.

$template_argument n$

[Method on Type]

If this gdb.Type is an instantiation of a template, this will return a new gdb.Type which represents the type of the *n*th template argument.

If this gdb.Type is not a template type, this will throw an exception. Ordinarily, only C++ code will have template types.

name is searched for globally.

Each type has a code, which indicates what category this type falls into. The available type categories are represented by constants defined in the gdb module:

TYPE_CODE_PTR

The type is a pointer.

TYPE_CODE_ARRAY

The type is an array.

TYPE_CODE_STRUCT

The type is a structure.

TYPE_CODE_UNION

The type is a union.

TYPE_CODE_ENUM

The type is an enum.

TYPE_CODE_FLAGS

A bit flags type, used for things such as status registers.

TYPE_CODE_FUNC

The type is a function.

TYPE_CODE_INT

The type is an integer type.

TYPE_CODE_FLT

A floating point type.

TYPE_CODE_VOID

The special type void.

TYPE_CODE_SET

A Pascal set type.

TYPE CODE RANGE

A range type, that is, an integer type with bounds.

TYPE_CODE_STRING

A string type. Note that this is only used for certain languages with languagedefined string types; C strings are not represented this way.

TYPE_CODE_BITSTRING

A string of bits.

TYPE_CODE_ERROR

An unknown or erroneous type.

```
TYPE_CODE_METHOD
```

A method type, as found in C++ or Java.

TYPE_CODE_METHODPTR

A pointer-to-member-function.

TYPE_CODE_MEMBERPTR

A pointer-to-member.

TYPE_CODE_REF

A reference type.

TYPE_CODE_CHAR

A character type.

TYPE_CODE_BOOL

A boolean type.

TYPE_CODE_COMPLEX

A complex float type.

TYPE_CODE_TYPEDEF

A typedef to some other type.

TYPE_CODE_NAMESPACE

A C++ namespace.

TYPE_CODE_DECFLOAT

A decimal floating point type.

TYPE_CODE_INTERNAL_FUNCTION

A function internal to GDB. This is the type used to represent convenience functions.

23.2.2.6 Pretty Printing

GDB provides a mechanism to allow pretty-printing of values using Python code. The pretty-printer API allows application-specific code to greatly simplify the display of complex objects. This mechanism works for both MI and the CLI.

For example, here is how a C++ std::string looks without a pretty-printer:

```
(gdb) print s
$1 = {
    static npos = 4294967295,
    _M_dataplus = {
        <std::allocator<char>> = {
            <__gnu_cxx::new_allocator<char>> = {<No data fields>}, <No data fields>},
            members of std::basic_string<char, std::char_traits<char>, std::allocator<char> >::_Alloc_hider:
            _M_p = 0x804a014 "abcd"
      }
}
```

After a pretty-printer for std::string has been installed, only the contents are printed:

```
(gdb) print s
$2 = "abcd"
```

A pretty-printer is just an object that holds a value and implements a specific interface, defined here.

children (self)

[Operation on pretty printer]

GDB will call this method on a pretty-printer to compute the children of the pretty-printer's value.

This method must return an object conforming to the Python iterator protocol. Each item returned by the iterator must be a tuple holding two elements. The first element is the "name" of the child; the second element is the child's value. The value can be any Python object which is convertible to a GDB value.

This method is optional. If it does not exist, GDB will act as though the value has no children.

display_hint (self)

[Operation on pretty printer]

The CLI may call this method and use its result to change the formatting of a value. The result will also be supplied to an MI consumer as a 'displayhint' attribute of the variable being printed.

This method is optional. If it does exist, this method must return a string.

Some display hints are predefined by GDB:

'array' Indicate that the object being printed is "array-like". The CLI uses this to respect parameters such as set print elements and set print array.

'map' Indicate that the object being printed is "map-like", and that the children of this value can be assumed to alternate between keys and values.

'string' Indicate that the object being printed is "string-like". If the printer's to_string method returns a Python string of some kind, then GDB will call its internal language-specific string-printing function to format the string. For the CLI this means adding quotation marks, possibly escaping some characters, respecting set print elements, and the like.

to_string (self)

[Operation on pretty printer]

GDB will call this method to display the string representation of the value passed to the object's constructor.

When printing from the CLI, if the to_string method exists, then GDB will prepend its result to the values returned by children. Exactly how this formatting is done is dependent on the display hint, and may change as more hints are added. Also, depending on the print settings (see Section 10.8 [Print Settings], page 102), the CLI may print just the result of to_string in a stack trace, omitting the result of children.

If this method returns a string, it is printed verbatim.

Otherwise, if this method returns an instance of gdb.Value, then GDB prints this value. This may result in a call to another pretty-printer.

If instead the method returns a Python value which is convertible to a gdb.Value, then GDB performs the conversion and prints the resulting value. Again, this may result in a call to another pretty-printer. Python scalars (integers, floats, and booleans) and strings are convertible to gdb.Value; other types are not.

If the result is not one of these types, an exception is raised.

23.2.2.7 Selecting Pretty-Printers

The Python list gdb.pretty_printers contains an array of functions that have been registered via addition as a pretty-printer. Each gdb.Objfile also contains a pretty_printers attribute.

A function on one of these lists is passed a single gdb. Value argument and should return a pretty-printer object conforming to the interface definition above (see Section 23.2.2.6 [Pretty Printing], page 269). If a function cannot create a pretty-printer for the value, it should return None.

GDB first checks the pretty_printers attribute of each gdb.Objfile and iteratively calls each function in the list for that gdb.Objfile until it receives a pretty-printer object. After these lists have been exhausted, it tries the global gdb.pretty-printers list, again calling each function until an object is returned.

The order in which the objfiles are searched is not specified. For a given list, functions are always invoked from the head of the list, and iterated over sequentially until the end of the list, or a printer object is returned.

Here is an example showing how a std::string printer might be written:

```
class StdStringPrinter:
    "Print a std::string"

def __init__ (self, val):
    self.val = val

def to_string (self):
    return self.val['_M_dataplus']['_M_p']

def display_hint (self):
    return 'string'
```

And here is an example showing how a lookup function for the printer example above might be written.

```
def str_lookup_function (val):
    lookup_tag = val.type.tag
    regex = re.compile ("^std::basic_string<char,.*>$")
    if lookup_tag == None:
        return None
    if regex.match (lookup_tag):
        return StdStringPrinter (val)
    return None
```

The example lookup function extracts the value's type, and attempts to match it to a type that it can pretty-print. If it is a type the printer can pretty-print, it will return a printer object. If not, it returns None.

We recommend that you put your core pretty-printers into a Python package. If your pretty-printers are for use with a library, we further recommend embedding a version number into the package name. This practice will enable GDB to load multiple versions of your pretty-printers at the same time, because they will have different names.

You should write auto-loaded code (see Section 23.2.2.3 [Auto-loading], page 263) such that it can be evaluated multiple times without changing its meaning. An ideal auto-load

file will consist solely of imports of your printer modules, followed by a call to a register pretty-printers with the current objfile.

Taken as a whole, this approach will scale nicely to multiple inferiors, each potentially using a different library version. Embedding a version number in the Python package name will ensure that gdb is able to load both sets of printers simultaneously. Then, because the search for pretty-printers is done by objfile, and because your auto-loaded code took care to register your library's printers with a specific objfile, gdb will find the correct printers for the specific version of the library used by each inferior.

To continue the std::string example (see Section 23.2.2.6 [Pretty Printing], page 269), this code might appear in gdb.libstdcxx.v6:

```
def register_printers (objfile):
   objfile.pretty_printers.add (str_lookup_function)
```

And then the corresponding contents of the auto-load file would be:

```
import gdb.libstdcxx.v6
gdb.libstdcxx.v6.register_printers (gdb.current_objfile ())
```

23.2.2.8 Commands In Python

You can implement new GDB CLI commands in Python. A CLI command is implemented using an instance of the gdb.Command class, most commonly using a subclass.

```
__init__ name command_class [completer_class] [Method on Command] [prefix]
```

The object initializer for Command registers the new command with GDB. This initializer is normally invoked from the subclass' own __init__ method.

name is the name of the command. If name consists of multiple words, then the initial words are looked for as prefix commands. In this case, if one of the prefix commands does not exist, an exception is raised.

There is no support for multi-line commands.

command_class should be one of the 'COMMAND_' constants defined below. This argument tells gdb how to categorize the new command in the help system.

completer_class is an optional argument. If given, it should be one of the 'COMPLETE_' constants defined below. This argument tells GDB how to perform completion for this command. If not given, GDB will attempt to complete using the object's complete method (see below); if no such method is found, an error will occur when completion is attempted.

prefix is an optional argument. If True, then the new command is a prefix command; sub-commands of this command may be registered.

The help text for the new command is taken from the Python documentation string for the command's class, if there is one. If no documentation string is provided, the default value "This command is not documented." is used.

dont_repeat [Method on Command]

By default, a GDB command is repeated when the user enters a blank line at the command prompt. A command can suppress this behavior by invoking the dont_repeat method. This is similar to the user command dont-repeat, see Section 23.1.1 [Define], page 255.

invoke argument from_tty

[Method on Command]

This method is called by GDB when this command is invoked.

argument is a string. It is the argument to the command, after leading and trailing whitespace has been stripped.

from_tty is a boolean argument. When true, this means that the command was entered by the user at the terminal; when false it means that the command came from elsewhere.

If this method throws an exception, it is turned into a GDB error call. Otherwise, the return value is ignored.

complete text word

[Method on Command]

This method is called by GDB when the user attempts completion on this command. All forms of completion are handled by this method, that is, the TAB and M-? key bindings (see Section 3.2 [Completion], page 19), and the complete command (see Section 3.3 [Help], page 21).

The arguments text and word are both strings. text holds the complete command line up to the cursor's location. word holds the last word of the command line; this is computed using a word-breaking heuristic.

The complete method can return several values:

- If the return value is a sequence, the contents of the sequence are used as the completions. It is up to complete to ensure that the contents actually do complete the word. A zero-length sequence is allowed, it means that there were no completions available. Only string elements of the sequence are used; other elements in the sequence are ignored.
- If the return value is one of the 'COMPLETE_' constants defined below, then the corresponding gdb-internal completion function is invoked, and its result is used.
- All other results are treated as though there were no available completions.

When a new command is registered, it must be declared as a member of some general class of commands. This is used to classify top-level commands in the on-line help system; note that prefix commands are not listed under their own category but rather that of their top-level command. The available classifications are represented by constants defined in the gdb module:

COMMAND_NONE

The command does not belong to any particular class. A command in this category will not be displayed in any of the help categories.

COMMAND_RUNNING

The command is related to running the inferior. For example, start, step, and continue are in this category. Type help running at the GDB prompt to see a list of commands in this category.

COMMAND_DATA

The command is related to data or variables. For example, call, find, and print are in this category. Type help data at the GDB prompt to see a list of commands in this category.

COMMAND_STACK

The command has to do with manipulation of the stack. For example, backtrace, frame, and return are in this category. Type help stack at the GDB prompt to see a list of commands in this category.

COMMAND_FILES

This class is used for file-related commands. For example, file, list and section are in this category. Type help files at the GDB prompt to see a list of commands in this category.

COMMAND_SUPPORT

This should be used for "support facilities", generally meaning things that are useful to the user when interacting with GDB, but not related to the state of the inferior. For example, help, make, and shell are in this category. Type help support at the GDB prompt to see a list of commands in this category.

COMMAND_STATUS

The command is an 'info'-related command, that is, related to the state of GDB itself. For example, info, macro, and show are in this category. Type help status at the GDB prompt to see a list of commands in this category.

COMMAND_BREAKPOINTS

The command has to do with breakpoints. For example, break, clear, and delete are in this category. Type help breakpoints at the GDB prompt to see a list of commands in this category.

COMMAND_TRACEPOINTS

The command has to do with tracepoints. For example, trace, actions, and tfind are in this category. Type help tracepoints at the GDB prompt to see a list of commands in this category.

COMMAND_OBSCURE

The command is only used in unusual circumstances, or is not of general interest to users. For example, checkpoint, fork, and stop are in this category. Type help obscure at the GDB prompt to see a list of commands in this category.

COMMAND_MAINTENANCE

The command is only useful to GDB maintainers. The maintenance and flushregs commands are in this category. Type help internals at the GDB prompt to see a list of commands in this category.

A new command can use a predefined completion function, either by specifying it via an argument at initialization, or by returning it from the complete method. These predefined completion constants are all defined in the gdb module:

COMPLETE_NONE

This constant means that no completion should be done.

COMPLETE_FILENAME

This constant means that filename completion should be performed.

COMPLETE_LOCATION

This constant means that location completion should be done. See Section 9.2 [Specify Location], page 84.

COMPLETE_COMMAND

This constant means that completion should examine GDB command names.

COMPLETE_SYMBOL

This constant means that completion should be done using symbol names as the source

The following code snippet shows how a trivial CLI command can be implemented in Python:

```
class HelloWorld (gdb.Command):
    """Greet the whole world."""

def __init__ (self):
    super (HelloWorld, self).__init__ ("hello-world", gdb.COMMAND_OBSCURE)

def invoke (self, arg, from_tty):
    print "Hello, World!"

HelloWorld ()
```

The last line instantiates the class, and is necessary to trigger the registration of the command with gdb. Depending on how the Python code is read into gdb, you may need to import the gdb module explicitly.

23.2.2.9 Writing new convenience functions

You can implement new convenience functions (see Section 10.10 [Convenience Vars], page 109) in Python. A convenience function is an instance of a subclass of the class gdb.Function.

```
__init__ name
```

[Method on Function]

The initializer for Function registers the new function with GDB. The argument name is the name of the function, a string. The function will be visible to the user as a convenience variable of type internal function, whose name is the same as the given name.

The documentation for the new function is taken from the documentation string for the new class.

invoke *args

[Method on Function]

When a convenience function is evaluated, its arguments are converted to instances of gdb.Value, and then the function's invoke method is called. Note that GDB does not predetermine the arity of convenience functions. Instead, all available arguments are passed to invoke, following the standard Python calling convention. In particular, a convenience function can have default values for parameters without ill effect.

The return value of this method is used as its value in the enclosing expression. If an ordinary Python value is returned, it is converted to a gdb. Value following the usual rules.

The following code snippet shows how a trivial convenience function can be implemented in Python:

```
class Greet (gdb.Function):
    """Return string to greet someone.
```

```
Takes a name as argument."""

def __init__ (self):
    super (Greet, self).__init__ ("greet")

def invoke (self, name):
    return "Hello, %s!" % name.string ()

Greet ()
```

The last line instantiates the class, and is necessary to trigger the registration of the function with GDB. Depending on how the Python code is read into GDB, you may need to import the gdb module explicitly.

23.2.2.10 Objfiles In Python

GDB loads symbols for an inferior from various symbol-containing files (see Section 18.1 [Files], page 187). These include the primary executable file, any shared libraries used by the inferior, and any separate debug info files (see Section 18.2 [Separate Debug Files], page 193). GDB calls these symbol-containing files objfiles.

The following objfile-related functions are available in the gdb module:

current_objfile

[Function]

When auto-loading a Python script (see Section 23.2.2.3 [Auto-loading], page 263), GDB sets the "current objfile" to the corresponding objfile. This function returns the current objfile. If there is no current objfile, this function returns None.

objfiles

[Function]

Return a sequence of all the objfiles current known to GDB. See Section 23.2.2.10 [Objfiles In Python], page 276.

Each objfile is represented by an instance of the gdb.Objfile class.

filename

[Instance Variable of Objfile]

The file name of the objfile as a string.

pretty_printers

[Instance Variable of Objfile]

The pretty_printers attribute is a list of functions. It is used to look up pretty-printers. A Value is passed to each function in order; if the function returns None, then the search continues. Otherwise, the return value should be an object which is used to format the value. See Section 23.2.2.6 [Pretty Printing], page 269, for more information.

23.2.2.11 Acessing inferior stack frames from Python.

When the debugged program stops, GDB is able to analyze its call stack (see Section 8.1 [Stack frames], page 77). The gdb.Frame class represents a frame in the stack. A gdb.Frame object is only valid while its corresponding frame exists in the inferior's stack. If you try to use an invalid frame object, GDB will throw a RuntimeError exception.

```
Two gdb.Frame objects can be compared for equality with the == operator, like:
    (gdb) python print gdb.newest_frame() == gdb.selected_frame()

True
```

The following frame-related functions are available in the gdb module:

selected_frame [Function]

Return the selected frame object. (see Section 8.3 [Selecting a Frame], page 80).

frame_stop_reason_string reason

[Function]

Return a string explaining the reason why GDB stopped unwinding frames, as expressed by the given reason code (an integer, see the unwind_stop_reason method further down in this section).

A gdb.Frame object has the following methods:

is_valid [Method on Frame]

Returns true if the gdb.Frame object is valid, false if not. A frame object can become invalid if the frame it refers to doesn't exist anymore in the inferior. All gdb.Frame methods will throw an exception if it is invalid at the time the method is called.

name [Method on Frame]

Returns the function name of the frame, or None if it can't be obtained.

type [Method on Frame]

Returns the type of the frame. The value can be one of gdb.NORMAL_FRAME, gdb.DUMMY_FRAME, gdb.SIGTRAMP_FRAME or gdb.SENTINEL_FRAME.

unwind_stop_reason

[Method on Frame]

Return an integer representing the reason why it's not possible to find more frames toward the outermost frame. Use gdb.frame_stop_reason_string to convert the value returned by this function to a string.

pc [Method on Frame]

Returns the frame's resume address.

older [Method on Frame]

Return the frame that called this frame.

newer [Method on Frame]

Return the frame called by this frame.

read_var variable [Method on Frame]

Return the value of the given variable in this frame. *variable* must be a string.

23.2.2.12 Python representation of lazy strings.

A lazy string is a string whose contents is not retrieved or encoded until it is needed.

A gdb.LazyString is represented in GDB as an address that points to a region of memory, an encoding that will be used to encode that region of memory, and a length to delimit the region of memory that represents the string. The difference between a gdb.LazyString and a string wrapped within a gdb.Value is that a gdb.LazyString will be treated differently by GDB when printing. A gdb.LazyString is retrieved and encoded during printing, while a gdb.Value wrapping a string is immediately retrieved and encoded on creation.

A gdb.LazyString object has the following functions:

value

[Method on LazyString]

Convert the gdb.LazyString to a gdb.Value. This value will point to the string in memory, but will lose all the delayed retrieval, encoding and handling that GDB applies to a gdb.LazyString.

address

[Instance Variable of LazyString]

This attribute holds the address of the string. This attribute is not writable.

length

[Instance Variable of LazyString]

This attribute holds the length of the string in characters. If the length is -1, then the string will be fetched and encoded up to the first null of appropriate width. This attribute is not writable.

encoding

[Instance Variable of LazyString]

This attribute holds the encoding that will be applied to the string when the string is printed by GDB. If the encoding is not set, or contains an empty string, then GDB will select the most appropriate encoding when the string is printed. This attribute is not writable.

type

[Instance Variable of LazyString]

This attribute holds the type that is represented by the lazy string's type. For a lazy string this will always be a pointer type. To resolve this to the lazy string's character type, use the type's target method. See Section 23.2.2.5 [Types In Python], page 266. This attribute is not writable.

24 Command Interpreters

GDB supports multiple command interpreters, and some command infrastructure to allow users or user interface writers to switch between interpreters or run commands in other interpreters.

GDB currently supports two command interpreters, the console interpreter (sometimes called the command-line interpreter or CLI) and the machine interface interpreter (or GDB/MI). This manual describes both of these interfaces in great detail.

By default, GDB will start with the console interpreter. However, the user may choose to start GDB with another interpreter by specifying the '-i' or '--interpreter' startup options. Defined interpreters include:

console The traditional console or command-line interpreter. This is the most often used interpreter with gdb. With no interpreter specified at runtime, gdb will use this interpreter.

mi The newest GDB/MI interface (currently mi2). Used primarily by programs wishing to use GDB as a backend for a debugger GUI or an IDE. For more information, see Chapter 27 [The GDB/MI Interface], page 289.

mi2 The current GDB/MI interface.

mi1 The GDB/MI interface included in GDB 5.1, 5.2, and 5.3.

The interpreter being used by GDB may not be dynamically switched at runtime. Although possible, this could lead to a very precarious situation. Consider an IDE using GDB/MI. If a user enters the command "interpreter-set console" in a console view, GDB would switch to using the console interpreter, rendering the IDE inoperable!

Although you may only choose a single interpreter at startup, you may execute commands in any interpreter from the current interpreter using the appropriate command. If you are running the console interpreter, simply use the interpreter-exec command:

interpreter-exec mi "-data-list-register-names"

GDB/MI has a similar command, although it is only available in versions of GDB which support GDB/MI version 2 (or greater).

25 GDB Text User Interface

The GDB Text User Interface (TUI) is a terminal interface which uses the curses library to show the source file, the assembly output, the program registers and GDB commands in separate text windows. The TUI mode is supported only on platforms where a suitable version of the curses library is available.

The TUI mode is enabled by default when you invoke GDB as either 'gdbtui' or 'gdb-tui'. You can also switch in and out of TUI mode while GDB runs by using various TUI commands and key bindings, such as C-x C-a. See Section 25.2 [TUI Key Bindings], page 282.

25.1 TUI Overview

In TUI mode, GDB can display several text windows:

command This window is the GDB command window with the GDB prompt and the GDB output. The GDB input is still managed using readline.

The source window shows the source file of the program. The current line and active breakpoints are displayed in this window.

assembly The assembly window shows the disassembly output of the program.

register This window shows the processor registers. Registers are highlighted when their values change.

The source and assembly windows show the current program position by highlighting the current line and marking it with a '>' marker. Breakpoints are indicated with two markers. The first marker indicates the breakpoint type:

- B Breakpoint which was hit at least once.
- b Breakpoint which was never hit.
- H Hardware breakpoint which was hit at least once.
- h Hardware breakpoint which was never hit.

The second marker indicates whether the breakpoint is enabled or not:

- + Breakpoint is enabled.
- Breakpoint is disabled.

The source, assembly and register windows are updated when the current thread changes, when the frame changes, or when the program counter changes.

These windows are not all visible at the same time. The command window is always visible. The others can be arranged in several layouts:

- source only.
- assembly only,
- source and assembly,
- source and registers, or
- assembly and registers.

A status line above the command window shows the following information:

Indicates the current GDB target. (see Chapter 19 [Specifying a Debugging Target], page 199).

process Gives the current process or thread number. When no process is being debugged, this field is set to No process.

function Gives the current function name for the selected frame. The name is demangled if demangling is turned on (see Section 10.8 [Print Settings], page 102). When there is no symbol corresponding to the current program counter, the string ?? is displayed.

line Indicates the current line number for the selected frame. When the current line number is not known, the string ?? is displayed.

pc Indicates the current program counter address.

25.2 TUI Key Bindings

The TUI installs several key bindings in the readline keymaps (see Chapter 31 [Command Line Editing], page 365). The following key bindings are installed for both TUI mode and the GDB standard mode.

С-х С-а

C-x a

- C-x A Enter or leave the TUI mode. When leaving the TUI mode, the curses window management stops and GDB operates using its standard mode, writing on the terminal directly. When reentering the TUI mode, control is given back to the curses windows. The screen is then refreshed.
- C-x 1 Use a TUI layout with only one window. The layout will either be 'source' or 'assembly'. When the TUI mode is not active, it will switch to the TUI mode. Think of this key binding as the Emacs C-x 1 binding.
- C-x 2 Use a TUI layout with at least two windows. When the current layout already has two windows, the next layout with two windows is used. When a new layout is chosen, one window will always be common to the previous layout and the new one.

Think of it as the Emacs C-x 2 binding.

C-x o Change the active window. The TUI associates several key bindings (like scrolling and arrow keys) with the active window. This command gives the focus to the next TUI window.

Think of it as the Emacs $C-x \circ binding$.

C-x s Switch in and out of the TUI SingleKey mode that binds single keys to GDB commands (see Section 25.3 [TUI Single Key Mode], page 283).

The following key bindings only work in the TUI mode:

PGUP Scroll the active window one page up.

PGDN Scroll the active window one page down.

UP Scroll the active window one line up.

DOWN Scroll the active window one line down.

LEFT Scroll the active window one column left.

RIGHT Scroll the active window one column right.

C-L Refresh the screen.

Because the arrow keys scroll the active window in the TUI mode, they are not available for their normal use by readline unless the command window has the focus. When another window is active, you must use other readline key bindings such as C-p, C-n, C-b and C-f to control the command window.

25.3 TUI Single Key Mode

The TUI also provides a SingleKey mode, which binds several frequently used GDB commands to single keys. Type C-x s to switch into this mode, where the following key bindings are used:

continue С down d finish f next n exit the SingleKey mode. q run rstep S u up info locals V where

Other keys temporarily switch to the GDB command prompt. The key that was pressed is inserted in the editing buffer so that it is possible to type most GDB commands without interaction with the TUI SingleKey mode. Once the command is entered the TUI SingleKey mode is restored. The only way to permanently leave this mode is by typing q or C-x s.

25.4 TUI-specific Commands

The TUI has specific commands to control the text windows. These commands are always available, even when gdb is not in the TUI mode. When gdb is in the standard mode, most of these commands will automatically switch to the TUI mode.

info win List and give the size of all displayed windows.

layout next

Display the next layout.

layout prev

Display the previous layout.

layout src

Display the source window only.

layout asm

Display the assembly window only.

layout split

Display the source and assembly window.

layout regs

Display the register window together with the source or assembly window.

focus next

Make the next window active for scrolling.

focus prev

Make the previous window active for scrolling.

focus src Make the source window active for scrolling.

focus asm Make the assembly window active for scrolling.

focus regs

Make the register window active for scrolling.

focus cmd Make the command window active for scrolling.

refresh Refresh the screen. This is similar to typing C-L.

tui reg float

Show the floating point registers in the register window.

tui reg general

Show the general registers in the register window.

tui reg next

Show the next register group. The list of register groups as well as their order is target specific. The predefined register groups are the following: general, float, system, vector, all, save, restore.

tui reg system

Show the system registers in the register window.

update Update the source window and the current execution point.

winheight name +count winheight name -count

Change the height of the window *name* by *count* lines. Positive counts increase the height, while negative counts decrease it.

tabset nchars

Set the width of tab stops to be nchars characters.

25.5 TUI Configuration Variables

Several configuration variables control the appearance of TUI windows.

set tui border-kind kind

Select the border appearance for the source, assembly and register windows. The possible values are the following:

space Use a space character to draw the border.

ascii Use ASCII characters '+', '-' and '|' to draw the border.

acs Use the Alternate Character Set to draw the border. The border is drawn using character line graphics if the terminal supports them.

set tui border-mode mode

set tui active-border-mode mode

Select the display attributes for the borders of the inactive windows or the active window. The *mode* can be one of the following:

normal Use normal attributes to display the border.

standout Use standout mode.

reverse Use reverse video mode.

half Use half bright mode.

half-standout

Use half bright and standout mode.

bold Use extra bright or bold mode.

bold-standout

Use extra bright or bold and standout mode.

26 Using GDB under GNU Emacs

A special interface allows you to use GNU Emacs to view (and edit) the source files for the program you are debugging with GDB.

To use this interface, use the command $M-x\ gdb$ in Emacs. Give the executable file you want to debug as an argument. This command starts GDB as a subprocess of Emacs, with input and output through a newly created Emacs buffer.

Running GDB under Emacs can be just like running GDB normally except for two things:

• All "terminal" input and output goes through an Emacs buffer, called the GUD buffer. This applies both to GDB commands and their output, and to the input and output done by the program you are debugging.

This is useful because it means that you can copy the text of previous commands and input them again; you can even use parts of the output in this way.

All the facilities of Emacs' Shell mode are available for interacting with your program. In particular, you can send signals the usual way—for example, C-c for an interrupt, C-c for a stop.

• GDB displays source code through Emacs.

Each time GDB displays a stack frame, Emacs automatically finds the source file for that frame and puts an arrow ('=>') at the left margin of the current line. Emacs uses a separate buffer for source display, and splits the screen to show both your GDB session and the source.

Explicit GDB list or search commands still produce output as usual, but you probably have no reason to use them from Emacs.

We call this text command mode. Emacs 22.1, and later, also uses a graphical mode, enabled by default, which provides further buffers that can control the execution and describe the state of your program. See Section "GDB Graphical Interface" in *The GNU Emacs Manual*.

If you specify an absolute file name when prompted for the *M-x gdb* argument, then Emacs sets your current working directory to where your program resides. If you only specify the file name, then Emacs sets your current working directory to to the directory associated with the previous buffer. In this case, GDB may find your program by searching your environment's PATH variable, but on some operating systems it might not find the source. So, although the GDB input and output session proceeds normally, the auxiliary buffer does not display the current source and line of execution.

The initial working directory of GDB is printed on the top line of the GUD buffer and this serves as a default for the commands that specify files for GDB to operate on. See Section 18.1 [Commands to Specify Files], page 187.

By default, M-x gdb calls the program called 'gdb'. If you need to call GDB by a different name (for example, if you keep several configurations around, with different names) you can customize the Emacs variable gud-gdb-command-name to run the one you want.

In the GUD buffer, you can use these special Emacs commands in addition to the standard Shell mode commands:

C-h m Describe the features of Emacs' GUD Mode.

C-c C-s Execute to another source line, like the GDB step command; also update the display window to show the current file and location.

- C-c C-n Execute to next source line in this function, skipping all function calls, like the GDB next command. Then update the display window to show the current file and location.
- C-c C-i Execute one instruction, like the GDB stepi command; update display window accordingly.
- C-c C-f Execute until exit from the selected stack frame, like the GDB finish command.
- C-c C-r Continue execution of your program, like the GDB continue command.
- C-c < Go up the number of frames indicated by the numeric argument (see Section "Numeric Arguments" in *The GNU Emacs Manual*), like the GDB up command.
- C-c > Go down the number of frames indicated by the numeric argument, like the GDB down command.

In any source file, the Emacs command C-x SPC (gud-break) tells GDB to set a breakpoint on the source line point is on.

In text command mode, if you type *M-x speedbar*, Emacs displays a separate frame which shows a backtrace when the GUD buffer is current. Move point to any frame in the stack and type RET to make it become the current frame and display the associated source in the source buffer. Alternatively, click *Mouse-2* to make the selected frame become the current one. In graphical mode, the speedbar displays watch expressions.

If you accidentally delete the source-display buffer, an easy way to get it back is to type the command f in the gdb buffer, to request a frame display; when you run under Emacs, this recreates the source buffer if necessary to show you the context of the current frame.

The source files displayed in Emacs are in ordinary Emacs buffers which are visiting the source files in the usual way. You can edit the files with these buffers if you wish; but keep in mind that GDB communicates with Emacs in terms of line numbers. If you add or delete lines from the text, the line numbers that GDB knows cease to correspond properly with the code.

A more detailed description of Emacs' interaction with gdb is given in the Emacs manual (see Section "Debuggers" in *The GNU Emacs Manual*).

27 The GDB/MI Interface

Function and Purpose

GDB/MI is a line based machine oriented text interface to GDB and is activated by specifying using the '--interpreter' command line option (see Section 2.1.2 [Mode Options], page 13). It is specifically intended to support the development of systems which use the debugger as just one small component of a larger system.

This chapter is a specification of the GDB/MI interface. It is written in the form of a reference manual.

Note that GDB/MI is still under construction, so some of the features described below are incomplete and subject to change (see Section 27.6 [GDB/MI Development and Front Ends], page 294).

Notation and Terminology

This chapter uses the following notation:

- | separates two alternatives.
- [something] indicates that something is optional: it may or may not be given.
- (group)* means that group inside the parentheses may repeat zero or more times.
- (group) + means that group inside the parentheses may repeat one or more times.
- "string" means a literal string.

27.3 GDB/MI General Design

Interaction of a GDB/MI frontend with gdb involves three parts—commands sent to gdb, responses to those commands and notifications. Each command results in exactly one response, indicating either successful completion of the command, or an error. For the commands that do not resume the target, the response contains the requested information. For the commands that resume the target, the response only indicates whether the target was successfully resumed. Notifications is the mechanism for reporting changes in the state of the target, or in gdb state, that cannot conveniently be associated with a command and reported as part of that command response.

The important examples of notifications are:

- Exec notifications. These are used to report changes in target state—when a target is resumed, or stopped. It would not be feasible to include this information in response of resuming commands, because one resume commands can result in multiple events in different threads. Also, quite some time may pass before any event happens in the target, while a frontend needs to know whether the resuming command itself was successfully executed.
- Console output, and status notifications. Console output notifications are used to report output of CLI commands, as well as diagnostics for other commands. Status notifications are used to report the progress of a long-running operation. Naturally, including this information in command response would mean no output is produced until the command is finished, which is undesirable.

• General notifications. Commands may have various side effects on the GDB or target state beyond their official purpose. For example, a command may change the selected thread. Although such changes can be included in command response, using notification allows for more orthogonal frontend design.

There's no guarantee that whenever an MI command reports an error, GDB or the target are in any specific state, and especially, the state is not reverted to the state before the MI command was processed. Therefore, whenever an MI command results in an error, we recommend that the frontend refreshes all the information shown in the user interface.

27.3.1 Context management

In most cases when GDB accesses the target, this access is done in context of a specific thread and frame (see Section 8.1 [Frames], page 77). Often, even when accessing global data, the target requires that a thread be specified. The CLI interface maintains the selected thread and frame, and supplies them to target on each command. This is convenient, because a command line user would not want to specify that information explicitly on each command, and because user interacts with GDB via a single terminal, so no confusion is possible as to what thread and frame are the current ones.

In the case of MI, the concept of selected thread and frame is less useful. First, a frontend can easily remember this information itself. Second, a graphical frontend can have more than one window, each one used for debugging a different thread, and the frontend might want to access additional threads for internal purposes. This increases the risk that by relying on implicitly selected thread, the frontend may be operating on a wrong one. Therefore, each MI command should explicitly specify which thread and frame to operate on. To make it possible, each MI command accepts the '--thread' and '--frame' options, the value to each is GDB identifier for thread and frame to operate on.

Usually, each top-level window in a frontend allows the user to select a thread and a frame, and remembers the user selection for further operations. However, in some cases GDB may suggest that the current thread be changed. For example, when stopping on a breakpoint it is reasonable to switch to the thread where breakpoint is hit. For another example, if the user issues the CLI 'thread' command via the frontend, it is desirable to change the frontend's selected thread to the one specified by user. GDB communicates the suggestion to change current thread using the '=thread-selected' notification. No such notification is available for the selected frame at the moment.

Note that historically, MI shares the selected thread with CLI, so frontends used the -thread-select to execute commands in the right context. However, getting this to work right is cumbersome. The simplest way is for frontend to emit -thread-select command before every command. This doubles the number of commands that need to be sent. The alternative approach is to suppress -thread-select if the selected thread in GDB is supposed to be identical to the thread the frontend wants to operate on. However, getting this optimization right can be tricky. In particular, if the frontend sends several commands to GDB, and one of the commands changes the selected thread, then the behaviour of subsequent commands will change. So, a frontend should either wait for response from such problematic commands, or explicitly add -thread-select for all subsequent commands. No frontend is known to do this exactly right, so it is suggested to just always pass the '--thread' and '--frame' options.

27.3.2 Asynchronous command execution and non-stop mode

On some targets, GDB is capable of processing MI commands even while the target is running. This is called asynchronous command execution (see Section 5.4.3 [Background Execution], page 68). The frontend may specify a preferrence for asynchronous execution using the -gdb-set target-async 1 command, which should be emitted before either running the executable or attaching to the target. After the frontend has started the executable or attached to the target, it can find if asynchronous execution is enabled using the -list-target-features command.

Even if GDB can accept a command while target is running, many commands that access the target do not work when the target is running. Therefore, asynchronous command execution is most useful when combined with non-stop mode (see Section 5.4.2 [Non-Stop Mode], page 67). Then, it is possible to examine the state of one thread, while other threads are running.

When a given thread is running, MI commands that try to access the target in the context of that thread may not work, or may work only on some targets. In particular, commands that try to operate on thread's stack will not work, on any target. Commands that read memory, or modify breakpoints, may work or not work, depending on the target. Note that even commands that operate on global state, such as print, set, and breakpoint commands, still access the target in the context of a specific thread, so frontend should try to find a stopped thread and perform the operation on that thread (using the '--thread' option).

Which commands will work in the context of a running thread is highly target dependent. However, the two commands <code>-exec-interrupt</code>, to stop a thread, and <code>-thread-info</code>, to find the state of a thread, will always work.

27.3.3 Thread groups

GDB may be used to debug several processes at the same time. On some platfroms, GDB may support debugging of several hardware systems, each one having several cores with several different processes running on each core. This section describes the MI mechanism to support such debugging scenarios.

The key observation is that regardless of the structure of the target, MI can have a global list of threads, because most commands that accept the '--thread' option do not need to know what process that thread belongs to. Therefore, it is not necessary to introduce neither additional '--process' option, nor an notion of the current process in the MI interface. The only strictly new feature that is required is the ability to find how the threads are grouped into processes.

To allow the user to discover such grouping, and to support arbitrary hierarchy of machines/cores/processes, MI introduces the concept of a thread group. Thread group is a collection of threads and other thread groups. A thread group always has a string identifier, a type, and may have additional attributes specific to the type. A new command, <code>-list-thread-groups</code>, returns the list of top-level thread groups, which correspond to processes that GDB is debugging at the moment. By passing an identifier of a thread group to the <code>-list-thread-groups</code> command, it is possible to obtain the members of specific thread group.

To allow the user to easily discover processes, and other objects, he wishes to debug, a concept of available thread group is introduced. Available thread group is an thread group that GDB is not debugging, but that can be attached to, using the -target-attach command. The list of available top-level thread groups can be obtained using '-list-thread-groups --available'. In general, the content of a thread group may be only retrieved only after attaching to that thread group.

27.4 GDB/MI Command Syntax

27.4.1 GDB/MI Input Syntax

```
command \mapsto
            cli-command | mi-command
cli-command \mapsto
            [ token ] cli-command nl, where cli-command is any existing GDB CLI com-
            mand.
mi-command \mapsto
            [ token ] "-" operation ( " " option )* [ " --" ] ( " " parameter )* nl
            "any sequence of digits"
token \mapsto
option \mapsto
            "-" parameter [ " " parameter ]
parameter \mapsto
            non-blank-sequence | c-string
operation \mapsto
            any of the operations described in this chapter
non-blank-sequence \mapsto
            anything, provided it doesn't contain special characters such as "-", nl, """ and
            of course " "
c-string \mapsto
            """ seven-bit-iso-c-string-content """
nl \mapsto
            CR | CR-LF
```

El Ott

Notes:

- The CLI commands are still handled by the MI interpreter; their output is described below.
- The token, when present, is passed back when the command finishes.
- Some MI commands accept optional arguments as part of the parameter list. Each option is identified by a leading '-' (dash) and may be followed by an optional argument parameter. Options occur first in the parameter list and can be delimited from normal parameters using '--' (this is useful when some parameters begin with a dash).

Pragmatics:

- We want easy access to the existing CLI syntax (for debugging).
- We want it to be easy to spot a MI operation.

27.4.2 GDB/MI Output Syntax

The output from GDB/MI consists of zero or more out-of-band records followed, optionally, by a single result record. This result record is for the most recent command. The sequence of output records is terminated by '(gdb)'.

If an input command was prefixed with a token then the corresponding output for that command will also be prefixed by that same token.

```
output \mapsto
            ( out-of-band-record )* [ result-record ] "(gdb)" nl
result-record \mapsto
            [token] "^" result-class ("," result )* nl
out-of-band-record \mapsto
            async-record | stream-record
\texttt{async-record} \, \mapsto \,
            exec-async-output | status-async-output | notify-async-output
exec-async-output \mapsto
            [ token ] "*" async-output
status-async-output \mapsto
            [ token ] "+" async-output
notify-async-output \mapsto
            [ token ] "=" async-output
async-output \mapsto
            async-class ( ", " result )* nl
result-class \mapsto
            "done" | "running" | "connected" | "error" | "exit"
async-class \mapsto
            "stopped" | others (where others will be added depending on the needs—this
            is still in development).
result \mapsto
            variable "=" value
variable \mapsto
            string
           const | tuple | list
value \mapsto
const \mapsto
           c-string
           "{}" | "{" result ( "," result )* "}"
\texttt{tuple} \, \mapsto \,
            "[]" | "[" value ( "," value )* "]" | "[" result ( "," result )* "]"
list \mapsto
stream\text{-}record \mapsto
            console-stream-output | target-stream-output | log-stream-output
console-stream-output \mapsto
            "~" c-string
```

Notes:

- All output sequences end in a single line containing a period.
- The token is from the corresponding request. Note that for all async output, while the token is allowed by the grammar and may be output by future versions of GDB for select async output messages, it is generally omitted. Frontends should treat all async output as reporting general changes in the state of the target and there should be no need to associate async output to any prior command.
- status-async-output contains on-going status information about the progress of a slow operation. It can be discarded. All status output is prefixed by '+'.
- exec-async-output contains asynchronous state change on the target (stopped, started, disappeared). All async output is prefixed by '*'.
- notify-async-output contains supplementary information that the client should handle (e.g., a new breakpoint information). All notify output is prefixed by '='.
- console-stream-output is output that should be displayed as is in the console. It is the textual response to a CLI command. All the console output is prefixed by '~'.
- target-stream-output is the output produced by the target program. All the target output is prefixed by '@'.
- log-stream-output is output text coming from GDB's internals, for instance messages that should be displayed as part of an error log. All the log output is prefixed by '&'.
- New gdb/mi commands should only output lists containing values.

See Section 27.7.2 [GDB/MI Stream Records], page 295, for more details about the various output records.

27.5 GDB/MI Compatibility with CLI

For the developers convenience CLI commands can be entered directly, but there may be some unexpected behaviour. For example, commands that query the user will behave as if the user replied yes, breakpoint command lists are not executed and some CLI commands, such as if, when and define, prompt for further input with '>', which is not valid MI output.

This feature may be removed at some stage in the future and it is recommended that front ends use the -interpreter-exec command (see [-interpreter-exec], page 351).

27.6 GDB/MI Development and Front Ends

The application which takes the MI output and presents the state of the program being debugged to the user is called a *front end*.

Although GDB/MI is still incomplete, it is currently being used by a variety of front ends to GDB. This makes it difficult to introduce new functionality without breaking existing usage. This section tries to minimize the problems by describing how the protocol might change.

Some changes in MI need not break a carefully designed front end, and for these the MI version will remain unchanged. The following is a list of changes that may occur within one level, so front ends should parse MI output in a way that can handle them:

- New MI commands may be added.
- New fields may be added to the output of any MI command.
- The range of values for fields with specified values, e.g., in_scope (see [-var-update], page 331) may be extended.

If the changes are likely to break front ends, the MI version level will be increased by one. This will allow the front end to parse the output according to the MI version. Apart from mi0, new versions of GDB will not support old versions of MI and it will be the responsibility of the front end to work with the new one.

The best way to avoid unexpected changes in MI that might break your front end is to make your project known to GDB developers and follow development on gdb@sourceware.org and gdb-patches@sourceware.org.

27.7 GDB/MI Output Records

27.7.1 GDB/MI Result Records

In addition to a number of out-of-band notifications, the response to a GDB/MI command includes one of the following result indications:

```
"^done" [ ", " results ]
```

The synchronous operation was successful, results are the return values.

"^running"

This result record is equivalent to ''done'. Historically, it was output instead of ''done' if the command has resumed the target. This behaviour is maintained for backward compatibility, but all frontends should treat ''done' and ''running' identically and rely on the '*running' output record to determine which threads are resumed.

"^connected"

GDB has connected to a remote target.

```
"^error" "," c-string
```

The operation failed. The *c-string* contains the corresponding error message.

"^exit" GDB has terminated.

27.7.2 GDB/MI Stream Records

GDB internally maintains a number of output streams: the console, the target, and the log. The output intended for each of these streams is funneled through the GDB/MI interface using stream records.

Each stream record begins with a unique prefix character which identifies its stream (see Section 27.4.2 [GDB/MI Output Syntax], page 293). In addition to the prefix, each stream record contains a string-output. This is either raw text (with an implicit new line) or a quoted C string (which does not contain an implicit newline).

"~" string-output

The console output stream contains text that should be displayed in the CLI console window. It contains the textual responses to CLI commands.

"@" string-output

The target output stream contains any textual output from the running target. This is only present when GDB's event loop is truly asynchronous, which is currently only the case for remote targets.

"&" string-output

The log stream contains debugging messages being produced by GDB's internals.

27.7.3 GDB/MI Async Records

Async records are used to notify the GDB/MI client of additional changes that have occurred. Those changes can either be a consequence of GDB/MI commands (e.g., a breakpoint modified) or a result of target activity (e.g., target stopped).

The following is the list of possible async records:

*running,thread-id="thread"

The target is now running. The thread field tells which specific thread is now running, and can be 'all' if all threads are running. The frontend should assume that no interaction with a running thread is possible after this notification is produced. The frontend should not assume that this notification is output only once for any command. GDB may emit this notification several times, either for different threads, because it cannot resume all threads together, or even for a single thread, if the thread must be stepped though some code before letting it run freely.

*stopped,reason="reason",thread-id="id",stopped-threads="stopped",core="core"

The target has stopped. The reason field can have one of the following values:

breakpoint-hit

A breakpoint was reached.

watchpoint-trigger

A watchpoint was triggered.

read-watchpoint-trigger

A read watchpoint was triggered.

access-watchpoint-trigger

An access watchpoint was triggered.

function-finished

An -exec-finish or similar CLI command was accomplished.

location-reached

An -exec-until or similar CLI command was accomplished.

watchpoint-scope

A watchpoint has gone out of scope.

end-stepping-range

An -exec-next, -exec-next-instruction, -exec-step, -exec-step-instruction or similar CLI command was accomplished.

exited-signalled

The inferior exited because of a signal.

exited The inferior exited.

exited-normally

The inferior exited normally.

signal-received

A signal was received by the inferior.

The *id* field identifies the thread that directly caused the stop – for example by hitting a breakpoint. Depending on whether all-stop mode is in effect (see Section 5.4.1 [All-Stop Mode], page 65), GDB may either stop all threads, or only the thread that directly triggered the stop. If all threads are stopped, the *stopped* field will have the value of "all". Otherwise, the value of the *stopped* field will be a list of thread identifiers. Presently, this list will always include a single thread, but frontend should be prepared to see several threads in the list. The *core* field reports the processor core on which the stop event has happened. This field may be absent if such information is not available.

=thread-group-created,id="id"

=thread-group-exited,id="id"

A thread thread group either was attached to, or has exited/detached from. The *id* field contains the GDB identifier of the thread group.

```
=thread-created,id="id",group-id="gid"
=thread-exited,id="id",group-id="gid"
```

A thread either was created, or has exited. The *id* field contains the gdb identifier of the thread. The *gid* field identifies the thread group this thread belongs to.

=thread-selected,id="id"

Informs that the selected thread was changed as result of the last command. This notification is not emitted as result of -thread-select command but is emitted whenever an MI command that is not documented to change the selected thread actually changes it. In particular, invoking, directly or indirectly (via user-defined command), the CLI thread command, will generate this notification.

We suggest that in response to this notification, front ends highlight the selected thread and cause subsequent commands to apply to that thread.

=library-loaded,...

Reports that a new library file was loaded by the program. This notification has 4 fields—id, target-name, host-name, and symbols-loaded. The id field is

an opaque identifier of the library. For remote debugging case, target-name and host-name fields give the name of the library file on the target, and on the host respectively. For native debugging, both those fields have the same value. The symbols-loaded field reports if the debug symbols for this library are loaded.

=library-unloaded,...

Reports that a library was unloaded by the program. This notification has 3 fields—id, target-name and host-name with the same meaning as for the =library-loaded notification

27.7.4 GDB/MI Frame Information

Response from many MI commands includes an information about stack frame. This information is a tuple that may have the following fields:

The level of the stack frame. The innermost frame has the level of zero. This field is always present.

func The name of the function corresponding to the frame. This field may be absent if GDB is unable to determine the function name.

addr The code address for the frame. This field is always present.

The name of the source files that correspond to the frame's code address. This field may be absent.

The source line corresponding to the frames' code address. This field may be absent.

The name of the binary file (either executable or shared library) the corresponds to the frame's code address. This field may be absent.

27.7.5 GDB/MI Thread Information

Whenever GDB has to report an information about a thread, it uses a tuple with the following fields:

id The numeric id assigned to the thread by GDB. This field is always present.

target-id

Target-specific string identifying the thread. This field is always present.

details Additional information about the thread provided by the target. It is supposed to be human-readable and not interpreted by the frontend. This field is optional.

Either 'stopped' or 'running', depending on whether the thread is presently running. This field is always present.

The value of this field is an integer number of the processor core the thread was last seen on. This field is optional.

27.8 Simple Examples of GDB/MI Interaction

This subsection presents several simple examples of interaction using the GDB/MI interface. In these examples, '->' means that the following line is passed to GDB/MI as input, while '<-' means the output received from GDB/MI.

Note the line breaks shown in the examples are here only for readability, they don't appear in the real output.

Setting a Breakpoint

Setting a breakpoint generates synchronous output which contains detailed information of the breakpoint.

```
-> -break-insert main
<- ^done,bkpt={number="1",type="breakpoint",disp="keep",
        enabled="y",addr="0x08048564",func="main",file="myprog.c",
        fullname="/home/nickrob/myprog.c",line="68",times="0"}
<- (gdb)</pre>
```

Program Execution

Program execution generates asynchronous records and MI gives the reason that execution stopped.

```
-> -exec-run
<- ^running
<- (gdb)
<- *stopped,reason="breakpoint-hit",disp="keep",bkptno="1",thread-id="0",
    frame={addr="0x08048564",func="main",
        args=[{name="argc",value="1"},{name="argv",value="0xbfc4d4d4"}],
    file="myprog.c",fullname="/home/nickrob/myprog.c",line="68"}
<- (gdb)
-> -exec-continue
<- ^running
<- (gdb)
<- *stopped,reason="exited-normally"
<- (gdb)</pre>
```

Quitting GDB

Quitting GDB just prints the result class '^exit'.

```
-> (gdb)
<- -gdb-exit
<- ^exit</pre>
```

Please note that '^exit' is printed immediately, but it might take some time for gdb to actually exit. During that time, gdb performs necessary cleanups, including killing programs being debugged or disconnecting from debug hardware, so the frontend should wait till gdb exits and should only forcibly kill gdb if it fails to exit in reasonable time.

A Bad Command

Here's what happens if you pass a non-existent command:

```
-> -rubbish
<- ^error,msg="Undefined MI command: rubbish"
<- (gdb)</pre>
```

27.9 GDB/MI Command Description Format

The remaining sections describe blocks of commands. Each block of commands is laid out in a fashion similar to this section.

Motivation

The motivation for this collection of commands.

Introduction

A brief introduction to this collection of commands as a whole.

Commands

For each command in the block, the following is described:

Synopsis

```
-command args...
```

Result

GDB Command

The corresponding GDB CLI command(s), if any.

Example

Example(s) formatted for readability. Some of the described commands have not been implemented yet and these are labeled N.A. (not available).

27.10 gdb/mi Breakpoint Commands

This section documents GDB/MI commands for manipulating breakpoints.

The -break-after Command

Synopsis

```
-break-after number count
```

The breakpoint number number is not in effect until it has been hit *count* times. To see how this is reflected in the output of the '-break-list' command, see the description of the '-break-list' command below.

GDB Command

The corresponding GDB command is 'ignore'.

Example

```
done,BreakpointTable={nr_rows="1",nr_cols="6",
hdr=[{width="3",alignment="-1",col_name="number",colhdr="Num"},
{width="14",alignment="-1",col_name="type",colhdr="Type"},
{width="4",alignment="-1",col_name="disp",colhdr="Disp"},
{width="3",alignment="-1",col_name="enabled",colhdr="Enb"},
{width="10",alignment="-1",col_name="addr",colhdr="Address"},
{width="40",alignment="2",col_name="what",colhdr="What"}],
body=[bkpt={number="1",type="breakpoint",disp="keep",enabled="y",
addr="0x000100d0",func="main",file="hello.c",fullname="/home/foo/hello.c",
line="5",times="0",ignore="3"}]}
(gdb)
```

The -break-commands Command

Synopsis

```
-break-commands number [ command1 ... commandN ]
```

Specifies the CLI commands that should be executed when breakpoint *number* is hit. The parameters *command1* to *commandN* are the commands. If no command is specified, any previously-set commands are cleared. See Section 5.1.7 [Break Commands], page 58. Typical use of this functionality is tracing a program, that is, printing of values of some variables whenever breakpoint is hit and then continuing.

GDB Command

The corresponding GDB command is 'commands'.

Example

The -break-condition Command

Synopsis

```
-break-condition number expr
```

Breakpoint *number* will stop the program only if the condition in *expr* is true. The condition becomes part of the '-break-list' output (see the description of the '-break-list' command below).

GDB Command

The corresponding GDB command is 'condition'.

Example

```
(gdb)
-break-condition 1 1
^done
```

```
(gdb)
-break-list
   done,BreakpointTable={nr_rows="1",nr_cols="6",
   hdr=[{width="3",alignment="-1",col_name="number",colhdr="Num"},
   {width="14",alignment="-1",col_name="type",colhdr="Type"},
   {width="4",alignment="-1",col_name="disp",colhdr="Disp"},
   {width="3",alignment="-1",col_name="enabled",colhdr="Enb"},
   {width="10",alignment="-1",col_name="addr",colhdr="Address"},
   {width="40",alignment="2",col_name="what",colhdr="What"}],
   body=[bkpt={number="1",type="breakpoint",disp="keep",enabled="y",
   addr="0x000100d0",func="main",file="hello.c",fullname="/home/foo/hello.c",
   line="5",cond="1",times="0",ignore="3"}]}
   (gdb)
```

The -break-delete Command

Synopsis

```
-break-delete ( breakpoint )+
```

Delete the breakpoint(s) whose number(s) are specified in the argument list. This is obviously reflected in the breakpoint list.

GDB Command

The corresponding GDB command is 'delete'.

Example

The -break-disable Command

Synopsis

```
-break-disable ( breakpoint )+
```

Disable the named breakpoint(s). The field 'enabled' in the break list is now set to 'n' for the named breakpoint(s).

gdb Command

The corresponding GDB command is 'disable'.

Example

(gdb)

```
-break-disable 2

'done
(gdb)
-break-list

'done,BreakpointTable={nr_rows="1",nr_cols="6",
hdr=[{width="3",alignment="-1",col_name="number",colhdr="Num"},
{width="14",alignment="-1",col_name="type",colhdr="Type"},
{width="4",alignment="-1",col_name="disp",colhdr="Disp"},
{width="3",alignment="-1",col_name="enabled",colhdr="Enb"},
{width="10",alignment="-1",col_name="addr",colhdr="Address"},
{width="40",alignment="2",col_name="what",colhdr="What"}],
body=[bkpt={number="2",type="breakpoint",disp="keep",enabled="n",
addr="0x000100d0",func="main",file="hello.c",fullname="/home/foo/hello.c",
line="5",times="0"}]}
(gdb)
```

The -break-enable Command

Synopsis

```
-break-enable (breakpoint)+
Enable (previously disabled) breakpoint(s).
```

GDB Command

The corresponding GDB command is 'enable'.

Example

```
(gdb)
-break-enable 2
^done
(gdb)
-break-list
^done,BreakpointTable={nr_rows="1",nr_cols="6",
hdr=[{width="3",alignment="-1",col_name="number",colhdr="Num"},
{width="14",alignment="-1",col_name="type",colhdr="Type"},
{width="4",alignment="-1",col_name="disp",colhdr="Disp"},
{width="3",alignment="-1",col_name="enabled",colhdr="Enb"},
{width="10",alignment="-1",col_name="addr",colhdr="Address"},
{width="40",alignment="2",col_name="what",colhdr="What"}],
body=[bkpt={number="2",type="breakpoint",disp="keep",enabled="y",
addr="0x000100d0",func="main",file="hello.c",fullname="/home/foo/hello.c",
line="5",times="0"}]}
(gdb)
```

The -break-info Command

Synopsis

```
-break-info breakpoint
```

Get information about a single breakpoint.

gdb Command

The corresponding GDB command is 'info break breakpoint'.

Example

N.A.

The -break-insert Command

Synopsis

```
-break-insert [ -t ] [ -h ] [ -f ] [ -d ] [ -c condition ] [ -i ignore-count ] [ -p thread ] [ location ]
```

If specified, *location*, can be one of:

- function
- filename:linenum
- filename:function
- *address

The possible optional parameters of this command are:

'-t' Insert a temporary breakpoint.

'-h' Insert a hardware breakpoint.

'-c condition'

Make the breakpoint conditional on condition.

'-i ignore-count'

Initialize the ignore-count.

- '-f' If location cannot be parsed (for example if it refers to unknown files or functions), create a pending breakpoint. Without this flag, GDB will report an error, and won't create a breakpoint, if location cannot be parsed.
- '-d' Create a disabled breakpoint.

Result

The result is in the form:

```
^done,bkpt={number="number",type="type",disp="del"|"keep",
enabled="y"|"n",addr="hex",func="funcname",file="filename",
fullname="full_filename",line="lineno",[thread="threadno,]
times="times"}
```

where number is the GDB number for this breakpoint, function is the name of the function where the breakpoint was inserted, filename is the name of the source file which contains this function, lineno is the source line number within that file and times the number of times that the breakpoint has been hit (always 0 for -break-insert but may be greater for -break-info or -break-list which use the same output).

Note: this format is open to change.

GDB Command

The corresponding GDB commands are 'break', 'tbreak', 'hbreak', 'thbreak', and 'rbreak'.

Example

```
(gdb)
-break-insert main
^done,bkpt={number="1",addr="0x0001072c",file="recursive2.c",
fullname="/home/foo/recursive2.c,line="4",times="0"}
-break-insert -t foo
^done,bkpt={number="2",addr="0x00010774",file="recursive2.c",
fullname="/home/foo/recursive2.c,line="11",times="0"}
-break-list
^done,BreakpointTable={nr_rows="2",nr_cols="6",
hdr=[{width="3",alignment="-1",col_name="number",colhdr="Num"},
{width="14",alignment="-1",col_name="type",colhdr="Type"},
{width="4",alignment="-1",col_name="disp",colhdr="Disp"},
{width="3",alignment="-1",col_name="enabled",colhdr="Enb"},
{width="10",alignment="-1",col_name="addr",colhdr="Address"},
{width="40",alignment="2",col_name="what",colhdr="What"}],
body=[bkpt={number="1",type="breakpoint",disp="keep",enabled="y",
addr="0x0001072c", func="main",file="recursive2.c",
fullname="/home/foo/recursive2.c,"line="4",times="0"},
bkpt={number="2",type="breakpoint",disp="del",enabled="y",
addr="0x00010774",func="foo",file="recursive2.c",
fullname="/home/foo/recursive2.c",line="11",times="0"}]}
(gdb)
-break-insert -r foo.*
~int foo(int, int);
^done,bkpt={number="3",addr="0x00010774",file="recursive2.c,
"fullname="/home/foo/recursive2.c",line="11",times="0"}
```

The -break-list Command

Synopsis

```
-break-list
```

Displays the list of inserted breakpoints, showing the following fields:

'Number' number of the breakpoint

'Type' type of the breakpoint: 'breakpoint' or 'watchpoint'

'Disposition'

should the breakpoint be deleted or disabled when it is hit: 'keep' or 'nokeep'

'Enabled' is the breakpoint enabled or no: 'y' or 'n'

'Address' memory location at which the breakpoint is set

'What' logical location of the breakpoint, expressed by function name, file name, line number

'Times' number of times the breakpoint has been hit

If there are no breakpoints or watchpoints, the BreakpointTable body field is an empty list.

gdb Command

The corresponding GDB command is 'info break'.

Example

```
(gdb)
   -break-list
   ^done,BreakpointTable={nr_rows="2",nr_cols="6",
   hdr=[{width="3",alignment="-1",col_name="number",colhdr="Num"},
   {width="14",alignment="-1",col_name="type",colhdr="Type"},
   {width="4",alignment="-1",col_name="disp",colhdr="Disp"},
   {width="3",alignment="-1",col_name="enabled",colhdr="Enb"},
   {width="10",alignment="-1",col_name="addr",colhdr="Address"},
   {width="40",alignment="2",col_name="what",colhdr="What"}],
   body=[bkpt={number="1",type="breakpoint",disp="keep",enabled="y",
   addr="0x000100d0",func="main",file="hello.c",line="5",times="0"},
   bkpt={number="2",type="breakpoint",disp="keep",enabled="y",
   addr="0x00010114",func="foo",file="hello.c",fullname="/home/foo/hello.c",
   line="13",times="0"}]}
Here's an example of the result when there are no breakpoints:
   (gdb)
   -break-list
   ^done,BreakpointTable={nr_rows="0",nr_cols="6",
   hdr=[{width="3",alignment="-1",col_name="number",colhdr="Num"},
   {width="14",alignment="-1",col_name="type",colhdr="Type"},
   {width="4",alignment="-1",col_name="disp",colhdr="Disp"},
   {width="3",alignment="-1",col_name="enabled",colhdr="Enb"}
   {width="10",alignment="-1",col_name="addr",colhdr="Address"},
   {width="40",alignment="2",col_name="what",colhdr="What"}],
   body=[]}
   (gdb)
```

The -break-watch Command

Synopsis

```
-break-watch [ -a | -r ]
```

Create a watchpoint. With the '-a' option it will create an access watchpoint, i.e., a watchpoint that triggers either on a read from or on a write to the memory location. With the '-r' option, the watchpoint created is a read watchpoint, i.e., it will trigger only when the memory location is accessed for reading. Without either of the options, the watchpoint created is a regular watchpoint, i.e., it will trigger when the memory location is accessed for writing. See Section 5.1.2 [Setting Watchpoints], page 49.

Note that '-break-list' will report a single list of watchpoints and breakpoints inserted.

GDB Command

The corresponding GDB commands are 'watch', 'awatch', and 'rwatch'.

Example

Setting a watchpoint on a variable in the main function:

```
(gdb)
*stopped,reason="watchpoint-trigger",wpt={number="2",exp="x"},
value={old="-268439212",new="55"},
frame={func="main",args=[],file="recursive2.c",
fullname="/home/foo/bar/recursive2.c",line="5"}
(gdb)
```

Setting a watchpoint on a variable local to a function. GDB will stop the program execution twice: first for the variable changing value, then for the watchpoint going out of scope.

```
(gdb)
-break-watch C
^done, wpt={number="5",exp="C"}
-exec-continue
^running
(gdb)
*stopped,reason="watchpoint-trigger",
wpt={number="5",exp="C"},value={old="-276895068",new="3"},
frame={func="callee4",args=[],
file="../../devo/gdb/testsuite/gdb.mi/basics.c",
fullname="/home/foo/bar/devo/gdb/testsuite/gdb.mi/basics.c",line="13"}
(gdb)
-exec-continue
^running
(gdb)
*stopped, reason="watchpoint-scope", wpnum="5",
frame={func="callee3", args=[{name="strarg",
value="0x11940 \"A string argument.\""}],
file="../../devo/gdb/testsuite/gdb.mi/basics.c",
fullname="/home/foo/bar/devo/gdb/testsuite/gdb.mi/basics.c",line="18"}
```

Listing breakpoints and watchpoints, at different points in the program execution. Note that once the watchpoint goes out of scope, it is deleted.

```
-break-watch C
^done, wpt={number="2", exp="C"}
(gdb)
-break-list
^done,BreakpointTable={nr_rows="2",nr_cols="6",
hdr=[{width="3",alignment="-1",col_name="number",colhdr="Num"},
{width="14",alignment="-1",col_name="type",colhdr="Type"},
{width="4",alignment="-1",col_name="disp",colhdr="Disp"},
{width="3",alignment="-1",col_name="enabled",colhdr="Enb"},
{width="10",alignment="-1",col_name="addr",colhdr="Address"},
{width="40",alignment="2",col_name="what",colhdr="What"}],
body=[bkpt={number="1",type="breakpoint",disp="keep",enabled="y",
addr="0x00010734",func="callee4",
file="../../devo/gdb/testsuite/gdb.mi/basics.c",
fullname="/home/foo/devo/gdb/testsuite/gdb.mi/basics.c"line="8",times="1"},
bkpt={number="2",type="watchpoint",disp="keep",
enabled="y",addr="",what="C",times="0"}]}
(gdb)
-exec-continue
^running
(gdb)
*stopped,reason="watchpoint-trigger",wpt={number="2",exp="C"},
```

```
value={old="-276895068",new="3"},
frame={func="callee4",args=[],
file="../../devo/gdb/testsuite/gdb.mi/basics.c",
fullname="/home/foo/bar/devo/gdb/testsuite/gdb.mi/basics.c",line="13"}
(gdb)
-break-list
^done,BreakpointTable={nr_rows="2",nr_cols="6",
hdr=[{width="3",alignment="-1",col_name="number",colhdr="Num"},
{width="14",alignment="-1",col_name="type",colhdr="Type"},
{width="4",alignment="-1",col_name="disp",colhdr="Disp"},
{width="3",alignment="-1",col_name="enabled",colhdr="Enb"},
{width="10",alignment="-1",col_name="addr",colhdr="Address"},
{width="40",alignment="2",col_name="what",colhdr="What"}],
body=[bkpt={number="1",type="breakpoint",disp="keep",enabled="y",
addr="0x00010734",func="callee4",
file="../../devo/gdb/testsuite/gdb.mi/basics.c",
fullname="/home/foo/devo/gdb/testsuite/gdb.mi/basics.c",line="8",times="1"},
bkpt={number="2",type="watchpoint",disp="keep",
enabled="y",addr="",what="C",times="-5"}]}
(gdb)
-exec-continue
^running
^done,reason="watchpoint-scope",wpnum="2",
frame={func="callee3",args=[{name="strarg",
value="0x11940 \"A string argument.\""}],
file="../../devo/gdb/testsuite/gdb.mi/basics.c",
fullname="/home/foo/bar/devo/gdb/testsuite/gdb.mi/basics.c",line="18"}
(gdb)
^done,BreakpointTable={nr_rows="1",nr_cols="6",
hdr=[{width="3",alignment="-1",col_name="number",colhdr="Num"},
{width="14",alignment="-1",col_name="type",colhdr="Type"},
{width="4",alignment="-1",col_name="disp",colhdr="Disp"},
{width="3",alignment="-1",col_name="enabled",colhdr="Enb"},
{width="10",alignment="-1",col_name="addr",colhdr="Address"},
{width="40",alignment="2",col_name="what",colhdr="What"}],
body=[bkpt={number="1",type="breakpoint",disp="keep",enabled="y",
addr="0x00010734",func="callee4",
file="../../devo/gdb/testsuite/gdb.mi/basics.c",
fullname="/home/foo/devo/gdb/testsuite/gdb.mi/basics.c",line="8",
times="1"}]}
(gdb)
```

27.11 GDB/MI Program Context

The -exec-arguments Command

Synopsis

```
-exec-arguments args
```

Set the inferior program arguments, to be used in the next '-exec-run'.

gdb Command

The corresponding GDB command is 'set args'.

Example

```
(gdb)
-exec-arguments -v word
^done
(gdb)
```

The -environment-cd Command

Synopsis

```
-environment-cd pathdir
Set GDB's working directory.
```

GDB Command

The corresponding GDB command is 'cd'.

Example

```
(gdb)
-environment-cd /kwikemart/marge/ezannoni/flathead-dev/devo/gdb
^done
(gdb)
```

The -environment-directory Command

Synopsis

```
-environment-directory [ -r ] [ pathdir ]+
```

Add directories pathdir to beginning of search path for source files. If the '-r' option is used, the search path is reset to the default search path. If directories pathdir are supplied in addition to the '-r' option, the search path is first reset and then addition occurs as normal. Multiple directories may be specified, separated by blanks. Specifying multiple directories in a single command results in the directories added to the beginning of the search path in the same order they were presented in the command. If blanks are needed as part of a directory name, double-quotes should be used around the name. In the command output, the path will show up separated by the system directory-separator character. The directory-separator character must not be used in any directory name. If no directories are specified, the current search path is displayed.

GDB Command

The corresponding GDB command is 'dir'.

Example

```
(gdb)
-environment-directory -r
^done,source-path="$cdir:$cwd"
(gdb)
```

The -environment-path Command

Synopsis

```
-environment-path [ -r ] [ pathdir ]+
```

Add directories pathdir to beginning of search path for object files. If the '-r' option is used, the search path is reset to the original search path that existed at gdb start-up. If directories pathdir are supplied in addition to the '-r' option, the search path is first reset and then addition occurs as normal. Multiple directories may be specified, separated by blanks. Specifying multiple directories in a single command results in the directories added to the beginning of the search path in the same order they were presented in the command. If blanks are needed as part of a directory name, double-quotes should be used around the name. In the command output, the path will show up separated by the system directory-separator character. The directory-separator character must not be used in any directory name. If no directories are specified, the current path is displayed.

GDB Command

The corresponding gdb command is 'path'.

Example

The -environment-pwd Command

Synopsis

```
-environment-pwd
```

Show the current working directory.

gdb Command

The corresponding gdb command is 'pwd'.

Example

27.12 GDB/MI Thread Commands

The -thread-info Command

Synopsis

```
-thread-info [ thread-id ]
```

Reports information about either a specific thread, if the *thread-id* parameter is present, or about all threads. When printing information about all threads, also reports the current thread.

gdb Command

The 'info thread' command prints the same information about all threads.

Example

The 'state' field may have the following values:

stopped The thread is stopped. Frame information is available for stopped threads.

running The thread is running. There's no frame information for running threads.

The -thread-list-ids Command

Synopsis

```
-thread-list-ids
```

Produces a list of the currently known GDB thread ids. At the end of the list it also prints the total number of such threads.

This command is retained for historical reasons, the -thread-info command should be used instead.

GDB Command

Part of 'info threads' supplies the same information.

Example

The -thread-select Command

Synopsis

```
-thread-select threadnum
```

Make threadnum the current thread. It prints the number of the new current thread, and the topmost frame for that thread.

This command is deprecated in favor of explicitly using the '--thread' option to each command.

GDB Command

The corresponding GDB command is 'thread'.

Example

```
(gdb)
-exec-next
^running
(gdb)
*stopped, reason="end-stepping-range", thread-id="2", line="187",
file="../../devo/gdb/testsuite/gdb.threads/linux-dp.c"
(gdb)
-thread-list-ids
^done,
thread-ids={thread-id="3",thread-id="2",thread-id="1"},
number-of-threads="3"
(gdb)
-thread-select 3
^done, new-thread-id="3",
frame={level="0",func="vprintf",
args=[{name="format",value="0x8048e9c \"%*s%c %d %c\\n\""},
{name="arg",value="0x2"}],file="vprintf.c",line="31"}
(gdb)
```

27.13 GDB/MI Program Execution

These are the asynchronous commands which generate the out-of-band record '*stopped'. Currently GDB only really executes asynchronously with remote targets and this interaction is mimicked in other cases.

The -exec-continue Command

Synopsis

```
-exec-continue [--reverse] [--all|--thread-group N]
```

Resumes the execution of the inferior program, which will continue to execute until it reaches a debugger stop event. If the '--reverse' option is specified, execution resumes in reverse until it reaches a stop event. Stop events may include

- breakpoints or watchpoints
- signals or exceptions
- the end of the process (or its beginning under '--reverse')
- the end or beginning of a replay log if one is being used.

In all-stop mode (see Section 5.4.1 [All-Stop Mode], page 65), may resume only one thread, or all threads, depending on the value of the 'scheduler-locking' variable. If

'--all' is specified, all threads will be resumed. The '--all' option is ignored in all-stop mode. If the '--thread-group' options is specified, then all threads in that thread group are resumed.

GDB Command

The corresponding GDB corresponding is 'continue'.

Example

```
-exec-continue
^running
(gdb)
@Hello world
*stopped,reason="breakpoint-hit",disp="keep",bkptno="2",frame={
func="foo",args=[],file="hello.c",fullname="/home/foo/bar/hello.c",
line="13"}
(gdb)
```

The -exec-finish Command

Synopsis

```
-exec-finish [--reverse]
```

Resumes the execution of the inferior program until the current function is exited. Displays the results returned by the function. If the '--reverse' option is specified, resumes the reverse execution of the inferior program until the point where current function was called.

GDB Command

The corresponding GDB command is 'finish'.

Example

Function returning void.

Function returning other than void. The name of the internal GDB variable storing the result is printed, together with the value itself.

The -exec-interrupt Command

Synopsis

```
-exec-interrupt [--all|--thread-group N]
```

Interrupts the background execution of the target. Note how the token associated with the stop message is the one for the execution command that has been interrupted. The token for the interrupt itself only appears in the '<code>^done</code>' output. If the user is trying to interrupt a non-running program, an error message will be printed.

Note that when asynchronous execution is enabled, this command is asynchronous just like other execution commands. That is, first the ''done' response will be printed, and the target stop will be reported after that using the '*stopped' notification.

In non-stop mode, only the context thread is interrupted by default. All threads will be interrupted if the '--all' option is specified. If the '--thread-group' option is specified, all threads in that group will be interrupted.

GDB Command

The corresponding gdb command is 'interrupt'.

Example

```
(gdb)
111-exec-continue
111^running
(gdb)
222-exec-interrupt
222^done
(gdb)
111*stopped,signal-name="SIGINT",signal-meaning="Interrupt",frame={addr="0x00010140",func="foo",args=[],file="try.c",fullname="/home/foo/bar/try.c",line="13"}
(gdb)
(gdb)
-exec-interrupt
^error,msg="mi_cmd_exec_interrupt: Inferior not executing."
(gdb)
```

The -exec-jump Command

Synopsis

```
-exec-jump location
```

Resumes execution of the inferior program at the location specified by parameter. See Section 9.2 [Specify Location], page 84, for a description of the different forms of location.

GDB Command

The corresponding GDB command is 'jump'.

Example

```
-exec-jump foo.c:10
*running,thread-id="all"
^running
```

The -exec-next Command

Synopsis

```
-exec-next [--reverse]
```

Resumes execution of the inferior program, stopping when the beginning of the next source line is reached.

If the '--reverse' option is specified, resumes reverse execution of the inferior program, stopping at the beginning of the previous source line. If you issue this command on the first line of a function, it will take you back to the caller of that function, to the source line where the function was called.

GDB Command

The corresponding GDB command is 'next'.

Example

```
-exec-next
^running
(gdb)
*stopped,reason="end-stepping-range",line="8",file="hello.c"
(gdb)
```

The -exec-next-instruction Command

Synopsis

```
-exec-next-instruction [--reverse]
```

Executes one machine instruction. If the instruction is a function call, continues until the function returns. If the program stops at an instruction in the middle of a source line, the address will be printed as well.

If the '--reverse' option is specified, resumes reverse execution of the inferior program, stopping at the previous instruction. If the previously executed instruction was a return from another function, it will continue to execute in reverse until the call to that function (from the current stack frame) is reached.

GDB Command

The corresponding GDB command is 'nexti'.

Example

```
(gdb)
-exec-next-instruction
    running

(gdb)
*stopped,reason="end-stepping-range",
addr="0x000100d4",line="5",file="hello.c"
(gdb)
```

The -exec-return Command

Synopsis

```
-exec-return
```

Makes current function return immediately. Doesn't execute the inferior. Displays the new current frame.

GDB Command

The corresponding gdb command is 'return'.

Example

```
200-break-insert callee4
200^done, bkpt={number="1", addr="0x00010734",
file="../../devo/gdb/testsuite/gdb.mi/basics.c",line="8"}
000-exec-run
000^running
(gdb)
000*stopped,reason="breakpoint-hit",disp="keep",bkptno="1",
frame={func="callee4",args=[],
file="../../devo/gdb/testsuite/gdb.mi/basics.c",
fullname="/home/foo/bar/devo/gdb/testsuite/gdb.mi/basics.c",line="8"}
(gdb)
205-break-delete
205<sup>done</sup>
(gdb)
111-exec-return
111^done,frame={level="0",func="callee3",
args=[{name="strarg",
value="0x11940 \"A string argument.\""}],
file="../../devo/gdb/testsuite/gdb.mi/basics.c",
fullname="/home/foo/bar/devo/gdb/testsuite/gdb.mi/basics.c",line="18"}
(gdb)
```

The -exec-run Command

Synopsis

```
-exec-rur
```

Starts execution of the inferior from the beginning. The inferior executes until either a breakpoint is encountered or the program exits. In the latter case the output will include an exit code, if the program has exited exceptionally.

GDB Command

The corresponding gdb command is 'run'.

Examples

```
(gdb)
-break-insert main
^done,bkpt={number="1",addr="0x0001072c",file="recursive2.c",line="4"}
(gdb)
-exec-run
^running
(gdb)
```

```
*stopped,reason="breakpoint-hit",disp="keep",bkptno="1",
      frame={func="main",args=[],file="recursive2.c";
      fullname="/home/foo/bar/recursive2.c",line="4"}
Program exited normally:
      (gdb)
      -exec-run
      ^running
      (gdb)
      x = 55
      *stopped,reason="exited-normally"
Program exited exceptionally:
      (gdb)
      -exec-run
      ^running
      (gdb)
      x = 55
      *stopped,reason="exited",exit-code="01"
```

Another way the program can terminate is if it receives a signal such as SIGINT. In this case, GDB/MI displays this:

```
(gdb)
*stopped, reason="exited-signalled", signal-name="SIGINT",
signal-meaning="Interrupt"
```

The -exec-step Command

Synopsis

```
-exec-step [--reverse]
```

Resumes execution of the inferior program, stopping when the beginning of the next source line is reached, if the next source line is not a function call. If it is, stop at the first instruction of the called function. If the '--reverse' option is specified, resumes reverse execution of the inferior program, stopping at the beginning of the previously executed source line.

gdb Command

The corresponding GDB command is 'step'.

Example

Stepping into a function:

```
-exec-step
^running
(gdb)
*stopped, reason="end-stepping-range",
frame={func="foo", args=[{name="a", value="10"},
{name="b",value="0"}],file="recursive2.c",
fullname="/home/foo/bar/recursive2.c",line="11"}
(gdb)
```

Regular stepping:

```
-exec-step
^running
(gdb)
*stopped,reason="end-stepping-range",line="14",file="recursive2.c"
(gdb)
```

The -exec-step-instruction Command

Synopsis

```
-exec-step-instruction [--reverse]
```

Resumes the inferior which executes one machine instruction. If the '--reverse' option is specified, resumes reverse execution of the inferior program, stopping at the previously executed instruction. The output, once GDB has stopped, will vary depending on whether we have stopped in the middle of a source line or not. In the former case, the address at which the program stopped will be printed as well.

gdb Command

The corresponding GDB command is 'stepi'.

Example

The -exec-until Command

Synopsis

```
-exec-until [ location ]
```

Executes the inferior until the *location* specified in the argument is reached. If there is no argument, the inferior executes until a source line greater than the current one is reached. The reason for stopping in this case will be 'location-reached'.

GDB Command

The corresponding GDB command is 'until'.

Example

(gdb)

27.14 GDB/MI Stack Manipulation Commands

The -stack-info-frame Command

Synopsis

```
-stack-info-frame
```

Get info on the selected frame.

GDB Command

The corresponding GDB command is 'info frame' or 'frame' (without arguments).

Example

The -stack-info-depth Command

Synopsis

```
-stack-info-depth [ max-depth ]
```

Return the depth of the stack. If the integer argument max-depth is specified, do not count beyond max-depth frames.

GDB Command

There's no equivalent gdb command.

Example

For a stack with frame levels 0 through 11:

```
(gdb)
-stack-info-depth
^done,depth="12"
(gdb)
-stack-info-depth 4
^done,depth="4"
(gdb)
-stack-info-depth 12
^done,depth="12"
(gdb)
-stack-info-depth 11
```

```
^done,depth="11"
(gdb)
-stack-info-depth 13
^done,depth="12"
(gdb)
```

The -stack-list-arguments Command

Synopsis

```
-stack-list-arguments print-values
[ low-frame high-frame ]
```

Display a list of the arguments for the frames between *low-frame* and *high-frame* (inclusive). If *low-frame* and *high-frame* are not provided, list the arguments for the whole call stack. If the two arguments are equal, show the single frame at the corresponding level. It is an error if *low-frame* is larger than the actual number of frames. On the other hand, *high-frame* may be larger than the actual number of frames, in which case only existing frames will be returned.

If print-values is 0 or --no-values, print only the names of the variables; if it is 1 or --all-values, print also their values; and if it is 2 or --simple-values, print the name, type and value for simple data types, and the name and type for arrays, structures and unions.

Use of this command to obtain arguments in a single frame is deprecated in favor of the '-stack-list-variables' command.

GDB Command

GDB does not have an equivalent command. gdbtk has a 'gdb_get_args' command which partially overlaps with the functionality of '-stack-list-arguments'.

Example

```
(gdb)
-stack-list-frames
^done,
stack=[
frame={level="0",addr="0x00010734",func="callee4",
file="../../devo/gdb/testsuite/gdb.mi/basics.c",
fullname="/home/foo/bar/devo/gdb/testsuite/gdb.mi/basics.c",line="8"},
frame={level="1",addr="0x0001076c",func="callee3",
file="../../devo/gdb/testsuite/gdb.mi/basics.c",
fullname="/home/foo/bar/devo/gdb/testsuite/gdb.mi/basics.c",line="17"},
frame={level="2",addr="0x0001078c",func="callee2",
file="../../devo/gdb/testsuite/gdb.mi/basics.c",
fullname="/home/foo/bar/devo/gdb/testsuite/gdb.mi/basics.c",line="22"},
frame={level="3",addr="0x000107b4",func="callee1",
file="../../devo/gdb/testsuite/gdb.mi/basics.c",
fullname="/home/foo/bar/devo/gdb/testsuite/gdb.mi/basics.c",line="27"},
frame={level="4",addr="0x000107e0",func="main",
file="../../devo/gdb/testsuite/gdb.mi/basics.c",
fullname="/home/foo/bar/devo/gdb/testsuite/gdb.mi/basics.c",line="32"}]
-stack-list-arguments 0
^done,
stack-args=[
```

```
frame={level="0",args=[]},
frame={level="1",args=[name="strarg"]},
frame={level="2",args=[name="intarg",name="strarg"]},
frame={level="3",args=[name="intarg",name="strarg",name="fltarg"]},
frame={level="4",args=[]}]
(gdb)
-stack-list-arguments 1
^done,
stack-args=[
frame={level="0",args=[]},
frame={level="1",
args=[{name="strarg",value="0x11940 \"A string argument.\""}]},
frame={level="2",args=[
{name="intarg",value="2"},
{name="strarg",value="0x11940 \"A string argument.\""}]},
{frame={level="3",args=[
{name="intarg",value="2"},
{name="strarg",value="0x11940 \"A string argument.\""},
{name="fltarg",value="3.5"}]}.
frame={level="4",args=[]}]
(gdb)
-stack-list-arguments 0 2 2
^done,stack-args=[frame={level="2",args=[name="intarg",name="strarg"]}]
-stack-list-arguments 1 2 2
^done, stack-args=[frame={level="2",
args=[{name="intarg", value="2"},
{name="strarg", value="0x11940 \"A string argument.\""}]}]
```

The -stack-list-frames Command

Synopsis

```
-stack-list-frames [ low-frame high-frame ]
```

List the frames currently on the stack. For each frame it displays the following info:

'level' The frame number, 0 being the topmost frame, i.e., the innermost function.

'addr' The \$pc value for that frame.

'func' Function name.

'file' File name of the source file where the function lives.

'line' Line number corresponding to the \$pc.

If invoked without arguments, this command prints a backtrace for the whole stack. If given two integer arguments, it shows the frames whose levels are between the two arguments (inclusive). If the two arguments are equal, it shows the single frame at the corresponding level. It is an error if *low-frame* is larger than the actual number of frames. On the other hand, *high-frame* may be larger than the actual number of frames, in which case only existing frames will be returned.

gdb Command

The corresponding GDB commands are 'backtrace' and 'where'.

Example

```
Full stack backtrace:
      (gdb)
      -stack-list-frames
      ^done.stack=
      [frame={level="0",addr="0x0001076c",func="foo",
       file="recursive2.c",fullname="/home/foo/bar/recursive2.c",line="11"},
      frame={level="1",addr="0x000107a4",func="foo",
       file="recursive2.c",fullname="/home/foo/bar/recursive2.c",line="14"},
      frame={level="2",addr="0x000107a4",func="foo",
        file="recursive2.c",fullname="/home/foo/bar/recursive2.c",line="14"},
      frame={level="3",addr="0x000107a4",func="foo",
       file="recursive2.c",fullname="/home/foo/bar/recursive2.c",line="14"},
      frame={level="4",addr="0x000107a4",func="foo",
        file="recursive2.c",fullname="/home/foo/bar/recursive2.c",line="14"},
      frame={level="5",addr="0x000107a4",func="foo",
        file="recursive2.c",fullname="/home/foo/bar/recursive2.c",line="14"},
      frame={level="6",addr="0x000107a4",func="foo",
        file="recursive2.c",fullname="/home/foo/bar/recursive2.c",line="14"},
      frame={level="7",addr="0x000107a4",func="foo",
        file="recursive2.c",fullname="/home/foo/bar/recursive2.c",line="14"},
      frame={level="8",addr="0x000107a4",func="foo",
        file="recursive2.c",fullname="/home/foo/bar/recursive2.c",line="14"},
      frame={level="9",addr="0x000107a4",func="foo",
        file="recursive2.c",fullname="/home/foo/bar/recursive2.c",line="14"},
      frame={level="10",addr="0x000107a4",func="foo",
        file="recursive2.c",fullname="/home/foo/bar/recursive2.c",line="14"},
      frame={level="11",addr="0x00010738",func="main",
        file="recursive2.c",fullname="/home/foo/bar/recursive2.c",line="4"}]
   Show frames between low_frame and high_frame:
      (gdb)
      -stack-list-frames 3 5
      ^done, stack=
      [frame={level="3",addr="0x000107a4",func="foo",
       file="recursive2.c",fullname="/home/foo/bar/recursive2.c",line="14"},
      frame={level="4",addr="0x000107a4",func="foo",
       file="recursive2.c",fullname="/home/foo/bar/recursive2.c",line="14"},
      frame={level="5",addr="0x000107a4",func="foo",
        file="recursive2.c",fullname="/home/foo/bar/recursive2.c",line="14"}]
      (gdb)
   Show a single frame:
      (gdb)
      -stack-list-frames 3 3
      ^done, stack=
      [frame={level="3",addr="0x000107a4",func="foo",
        file="recursive2.c",fullname="/home/foo/bar/recursive2.c",line="14"}]
```

The -stack-list-locals Command

Synopsis

```
-stack-list-locals print-values
```

Display the local variable names for the selected frame. If *print-values* is 0 or --no-values, print only the names of the variables; if it is 1 or --all-values, print also their

values; and if it is 2 or --simple-values, print the name, type and value for simple data types, and the name and type for arrays, structures and unions. In this last case, a frontend can immediately display the value of simple data types and create variable objects for other data types when the user wishes to explore their values in more detail.

This command is deprecated in favor of the '-stack-list-variables' command.

GDB Command

'info locals' in GDB, 'gdb_get_locals' in gdbtk.

Example

The -stack-list-variables Command

Synopsis

```
-stack-list-variables print-values
```

Display the names of local variables and function arguments for the selected frame. If print-values is 0 or --no-values, print only the names of the variables; if it is 1 or --all-values, print also their values; and if it is 2 or --simple-values, print the name, type and value for simple data types, and the name and type for arrays, structures and unions.

Example

```
(gdb)
-stack-list-variables --thread 1 --frame 0 --all-values
^done,variables=[{name="x",value="11"},{name="s",value="{a = 1, b = 2}"}]
(gdb)
```

The -stack-select-frame Command

Synopsis

```
-stack-select-frame framenum
```

Change the selected frame. Select a different frame framenum on the stack.

This command in deprecated in favor of passing the '--frame' option to every command.

GDB Command

The corresponding GDB commands are 'frame', 'up', 'down', 'select-frame', 'up-silent', and 'down-silent'.

Example

```
(gdb)
-stack-select-frame 2
done
(gdb)
```

27.15 GDB/MI Variable Objects

Introduction to Variable Objects

Variable objects are "object-oriented" MI interface for examining and changing values of expressions. Unlike some other MI interfaces that work with expressions, variable objects are specifically designed for simple and efficient presentation in the frontend. A variable object is identified by string name. When a variable object is created, the frontend specifies the expression for that variable object. The expression can be a simple variable, or it can be an arbitrary complex expression, and can even involve CPU registers. After creating a variable object, the frontend can invoke other variable object operations—for example to obtain or change the value of a variable object, or to change display format.

Variable objects have hierarchical tree structure. Any variable object that corresponds to a composite type, such as structure in C, has a number of child variable objects, for example corresponding to each element of a structure. A child variable object can itself have children, recursively. Recursion ends when we reach leaf variable objects, which always have built-in types. Child variable objects are created only by explicit request, so if a frontend is not interested in the children of a particular variable object, no child will be created.

For a leaf variable object it is possible to obtain its value as a string, or set the value from a string. String value can be also obtained for a non-leaf variable object, but it's generally a string that only indicates the type of the object, and does not list its contents. Assignment to a non-leaf variable object is not allowed.

A frontend does not need to read the values of all variable objects each time the program stops. Instead, MI provides an update command that lists all variable objects whose values has changed since the last update operation. This considerably reduces the amount of data that must be transferred to the frontend. As noted above, children variable objects are created on demand, and only leaf variable objects have a real value. As result, gdb will read target memory only for leaf variables that frontend has created.

The automatic update is not always desirable. For example, a frontend might want to keep a value of some expression for future reference, and never update it. For another example, fetching memory is relatively slow for embedded targets, so a frontend might want to disable automatic update for the variables that are either not visible on the screen, or "closed". This is possible using so called "frozen variable objects". Such variable objects are never implicitly updated.

Variable objects can be either fixed or floating. For the fixed variable object, the expression is parsed when the variable object is created, including associating identifiers to specific variables. The meaning of expression never changes. For a floating variable object the values of variables whose names appear in the expressions are re-evaluated every time in the context of the current frame. Consider this example:

```
void do_work(...)
{
```

```
struct work_state state;

if (...)
          do_work(...);
}
```

If a fixed variable object for the state variable is created in this function, and we enter the recursive call, the the variable object will report the value of state in the top-level do_work invocation. On the other hand, a floating variable object will report the value of state in the current frame.

If an expression specified when creating a fixed variable object refers to a local variable, the variable object becomes bound to the thread and frame in which the variable object is created. When such variable object is updated, GDB makes sure that the thread/frame combination the variable object is bound to still exists, and re-evaluates the variable object in context of that thread/frame.

The following is the complete set of GDB/MI operations defined to access this functionality:

Operation

Description

-enable-pretty-printing -var-create -var-delete -var-set-format -var-show-format -var-info-num-children -var-list-children -var-info-type -var-info-expression	enable Python-based pretty-printing create a variable object delete the variable object and/or its children set the display format of this variable show the display format of this variable tells how many children this object has return a list of the object's children show the type of this variable object print parent-relative expression that this variable object represents
-var-info-path-expression	print full expression that this variable object represents
-var-show-attributes	is this variable editable? does it exist here?
-var-evaluate-expression	get the value of this variable
-var-assign	set the value of this variable
-var-update	update the variable and its children
-var-set-frozen	set frozeness attribute
-var-set-update-range	set range of children to display on update

In the next subsection we describe each operation in detail and suggest how it can be used.

Description And Use of Operations on Variable Objects

The -enable-pretty-printing Command

```
-enable-pretty-printing
```

GDB allows Python-based visualizers to affect the output of the MI variable object commands. However, because there was no way to implement this in a fully backward-compatible way, a front end must request that this functionality be enabled.

Once enabled, this feature cannot be disabled.

Note that if Python support has not been compiled into GDB, this command will still succeed (and do nothing).

This feature is currently (as of gdb 7.0) experimental, and may work differently in future versions of gdb.

The -var-create Command

Synopsis

```
-var-create {name | "-"} {frame-addr | "*" | "@"} expression
```

This operation creates a variable object, which allows the monitoring of a variable, the result of an expression, a memory cell or a CPU register.

The name parameter is the string by which the object can be referenced. It must be unique. If '-' is specified, the varobj system will generate a string "varNNNNNN" automatically. It will be unique provided that one does not specify name of that format. The command fails if a duplicate name is found.

The frame under which the expression should be evaluated can be specified by frame-addr. A '*' indicates that the current frame should be used. A '@' indicates that a floating variable object must be created.

expression is any expression valid on the current language set (must not begin with a '*'), or one of the following:

- '*addr', where addr is the address of a memory cell
- '*addr-addr' a memory address range (TBD)
- '\$regname' a CPU register name

A varobj's contents may be provided by a Python-based pretty-printer. In this case the varobj is known as a *dynamic varobj*. Dynamic varobjs have slightly different semantics in some cases. If the -enable-pretty-printing command is not sent, then GDB will never create a dynamic varobj. This ensures backward compatibility for existing clients.

Result

This operation returns attributes of the newly-created varibj. These are:

'name' The name of the varobj.

'numchild'

The number of children of the varobj. This number is not necessarily reliable for a dynamic varobj. Instead, you must examine the 'has_more' attribute.

'value' The varobj's scalar value. For a varobj whose type is some sort of aggregate (e.g., a struct), or for a dynamic varobj, this value will not be interesting.

'type' The varobj's type. This is a string representation of the type, as would be printed by the gdb CLI.

'thread-id'

If a variable object is bound to a specific thread, then this is the thread's identifier.

'has_more'

For a dynamic varobj, this indicates whether there appear to be any children available. For a non-dynamic varobj, this will be 0.

'dynamic' This attribute will be present and have the value '1' if the varobj is a dynamic varobj. If the varobj is not a dynamic varobj, then this attribute will not be present.

'displayhint'

A dynamic varobj can supply a display hint to the front end. The value comes directly from the Python pretty-printer object's display_hint method. See Section 23.2.2.6 [Pretty Printing], page 269.

Typical output will look like this:

```
name="name",numchild="N",type="type",thread-id="M",
has_more="has_more"
```

The -var-delete Command

Synopsis

```
-var-delete [ -c ] name
```

Deletes a previously created variable object and all of its children. With the '-c' option, just deletes the children.

Returns an error if the object name is not found.

The -var-set-format Command

Synopsis

```
-var-set-format name format-spec
```

Sets the output format for the value of the object name to be format-spec.

The syntax for the format-spec is as follows:

```
\begin{array}{lll} \textit{format-spec} & \mapsto \\ \{ \texttt{binary} \mid \texttt{decimal} \mid \texttt{hexadecimal} \mid \texttt{octal} \mid \texttt{natural} \} \end{array}
```

The natural format is the default format choosen automatically based on the variable type (like decimal for an int, hex for pointers, etc.).

For a variable with children, the format is set only on the variable itself, and the children are not affected.

The -var-show-format Command

Synopsis

```
-var-show-format name
```

Returns the format used to display the value of the object name.

```
format \mapsto format-spec
```

The -var-info-num-children Command

328 Debugging with GDB

Synopsis

-var-info-num-children name

Returns the number of children of a variable object name:

numchild=n

Note that this number is not completely reliable for a dynamic varobj. It will return the current number of children, but more children may be available.

The -var-list-children Command

Synopsis

-var-list-children [print-values] name [from to]

Return a list of the children of the specified variable object and create variable objects for them, if they do not already exist. With a single argument or if print-values has a value for of 0 or --no-values, print only the names of the variables; if print-values is 1 or --all-values, also print their values; and if it is 2 or --simple-values print the name and value for simple data types and just the name for arrays, structures and unions.

from and to, if specified, indicate the range of children to report. If from or to is less than zero, the range is reset and all children will be reported. Otherwise, children starting at from (zero-based) and up to and excluding to will be reported.

If a child range is requested, it will only affect the current call to -var-list-children, but not future calls to -var-update. For this, you must instead use -var-set-updaterange. The intent of this approach is to enable a front end to implement any update approach it likes; for example, scrolling a view may cause the front end to request more children with -var-list-children, and then the front end could call -var-set-updaterange with a different range to ensure that future updates are restricted to just the visible items.

For each child the following results are returned:

Name of the variable object created for this child. name

The expression to be shown to the user by the front end to designate this child. expFor example this may be the name of a structure member.

> For a dynamic varobj, this value cannot be used to form an expression. There is no way to do this at all with a dynamic varobj.

> For C/C++ structures there are several pseudo children returned to designate access qualifiers. For these pseudo children exp is 'public', 'private', or 'protected'. In this case the type and value are not present.

> A dynamic varobj will not report the access qualifying pseudo-children, regardless of the language. This information is not available at all with a dynamic varobj.

numchild Number of children this child has. For a dynamic varibj, this will be 0.

typeThe type of the child.

value If values were requested, this is the value.

thread-id If this variable object is associated with a thread, this is the thread id. Other-

wise this result is not present.

frozen If the variable object is frozen, this variable will be present with a value of 1.

The result may have its own attributes:

'displayhint'

A dynamic varobj can supply a display hint to the front end. The value comes directly from the Python pretty-printer object's display_hint method. See Section 23.2.2.6 [Pretty Printing], page 269.

'has_more

This is an integer attribute which is nonzero if there are children remaining after the end of the selected range.

Example

The -var-info-type Command

Synopsis

```
-var-info-type name
```

Returns the type of the specified variable *name*. The type is returned as a string in the same format as it is output by the GDB CLI:

```
type=typename
```

The -var-info-expression Command

Synopsis

```
-var-info-expression name
```

Returns a string that is suitable for presenting this variable object in user interface. The string is generally not valid expression in the current language, and cannot be evaluated.

For example, if a is an array, and variable object A was created for a, then we'll get this output:

```
(gdb) -var-info-expression A.1 ^done,lang="C",exp="1"
```

Here, the values of lang can be $\{"C" \mid "C++" \mid "Java"\}$.

Note that the output of the -var-list-children command also includes those expressions, so the -var-info-expression command is of limited use.

The -var-info-path-expression Command

Synopsis

```
-var-info-path-expression name
```

Returns an expression that can be evaluated in the current context and will yield the same value that a variable object has. Compare this with the -var-info-expression command, which result can be used only for UI presentation. Typical use of the -var-info-path-expression command is creating a watchpoint from a variable object.

This command is currently not valid for children of a dynamic varobj, and will give an error when invoked on one.

For example, suppose C is a C++ class, derived from class Base, and that the Base class has a member called m_size. Assume a variable c is has the type of C and a variable object C was created for variable c. Then, we'll get this output:

```
(gdb) -var-info-path-expression C.Base.public.m_size
^done,path_expr=((Base)c).m_size)
```

The -var-show-attributes Command

Synopsis

```
-var-show-attributes name

List attributes of the specified variable object name:

status=attr [ ( ,attr )* ]

where attr is { { editable | noneditable } | TBD }.
```

The -var-evaluate-expression Command

Synopsis

```
-var-evaluate-expression [-f format-spec] name
```

Evaluates the expression that is represented by the specified variable object and returns its value as a string. The format of the string can be specified with the '-f' option. The possible values of this option are the same as for -var-set-format (see [-var-set-format], page 327). If the '-f' option is not specified, the current display format will be used. The current display format can be changed using the -var-set-format command.

```
value=value
```

Note that one must invoke -var-list-children for a variable before the value of a child variable can be evaluated.

The -var-assign Command

Synopsis

```
-var-assign name expression
```

Assigns the value of expression to the variable object specified by name. The object must be 'editable'. If the variable's value is altered by the assign, the variable will show up in any subsequent -var-update list.

Example

```
(gdb)
-var-assign var1 3
^done,value="3"
(gdb)
-var-update *
^done,changelist=[{name="var1",in_scope="true",type_changed="false"}]
(gdb)
```

The -var-update Command

Synopsis

```
-var-update [print-values] {name | "*"}
```

Reevaluate the expressions corresponding to the variable object name and all its direct and indirect children, and return the list of variable objects whose values have changed; name must be a root variable object. Here, "changed" means that the result of -var-evaluate-expression before and after the -var-update is different. If '*' is used as the variable object names, all existing variable objects are updated, except for frozen ones (see [-var-set-frozen], page 332). The option print-values determines whether both names and values, or just names are printed. The possible values of this option are the same as for -var-list-children (see [-var-list-children], page 328). It is recommended to use the '--all-values' option, to reduce the number of MI commands needed on each program stop.

With the '*' parameter, if a variable object is bound to a currently running thread, it will not be updated, without any diagnostic.

If -var-set-update-range was previously used on a varobj, then only the selected range of children will be reported.

-var-update reports all the changed varobjs in a tuple named 'changelist'.

Each item in the change list is itself a tuple holding:

'name' The name of the varobj.

'value' If values were requested for this update, then this field will be present and will hold the value of the varobj.

'in_scope'

This field is a string which may take one of three values:

"true" The variable object's current value is valid.

"false" The variable object does not currently hold a valid value but it may hold one in the future if its associated expression comes back into scope.

"invalid"

The variable object no longer holds a valid value. This can occur when the executable file being debugged has changed, either through recompilation or by using the GDB file command. The front end should normally choose to delete these variable objects.

In the future new values may be added to this list so the front should be prepared for this possibility. See Section 27.6 [GDB/MI Development and Front Ends], page 294.

'type_changed'

This is only present if the varobj is still valid. If the type changed, then this will be the string 'true'; otherwise it will be 'false'.

'new_type'

If the varobj's type changed, then this field will be present and will hold the new type.

'new_num_children'

For a dynamic varobj, if the number of children changed, or if the type changed, this will be the new number of children.

The 'numchild' field in other varobj responses is generally not valid for a dynamic varobj – it will show the number of children that GDB knows about, but because dynamic varobjs lazily instantiate their children, this will not reflect the number of children which may be available.

The 'new_num_children' attribute only reports changes to the number of children known by GDB. This is the only way to detect whether an update has removed children (which necessarily can only happen at the end of the update range).

'displayhint'

The display hint, if any.

'has_more'

This is an integer value, which will be 1 if there are more children available outside the varobj's update range.

'dynamic' This attribute will be present and have the value '1' if the varobj is a dynamic varobj. If the varobj is not a dynamic varobj, then this attribute will not be present.

'new_children'

If new children were added to a dynamic varobj within the selected update range (as set by -var-set-update-range), then they will be listed in this attribute.

Example

The -var-set-frozen Command

Synopsis

```
-var-set-frozen name flag
```

Set the frozenness flag on the variable object name. The flag parameter should be either '1' to make the variable frozen or '0' to make it unfrozen. If a variable object is frozen, then neither itself, nor any of its children, are implicitly updated by -var-update of a parent variable or by -var-update *. Only -var-update of the variable itself will update its value and values of its children. After a variable object is unfrozen, it is implicitly updated by all subsequent -var-update operations. Unfreezing a variable does not update it, only subsequent -var-update does.

Example

```
(gdb)
-var-set-frozen V 1
^done
(gdb)
```

The -var-set-update-range command

Synopsis

```
-var-set-update-range name from to
```

Set the range of children to be returned by future invocations of -var-update.

from and to indicate the range of children to report. If from or to is less than zero, the range is reset and all children will be reported. Otherwise, children starting at from (zero-based) and up to and excluding to will be reported.

Example

```
(gdb)
-var-set-update-range V 1 2
```

The -var-set-visualizer command

Synopsis

```
-var-set-visualizer name visualizer
```

Set a visualizer for the variable object name.

visualizer is the visualizer to use. The special value 'None' means to disable any visualizer in use.

If not 'None', visualizer must be a Python expression. This expression must evaluate to a callable object which accepts a single argument. GDB will call this object with the value of the varobj name as an argument (this is done so that the same Python pretty-printing code can be used for both the CLI and MI). When called, this object must return an object which conforms to the pretty-printing interface (see Section 23.2.2.6 [Pretty Printing], page 269).

The pre-defined function gdb.default_visualizer may be used to select a visualizer by following the built-in process (see Section 23.2.2.7 [Selecting Pretty-Printers], page 271). This is done automatically when a varobj is created, and so ordinarily is not needed.

This feature is only available if Python support is enabled. The MI command -list-features (see Section 27.22 [GDB/MI Miscellaneous Commands], page 347) can be used to check this.

Example

```
Resetting the visualizer:
    (gdb)
    -var-set-visualizer V None
    ^done

Reselecting the default (type-based) visualizer:
    (gdb)
    -var-set-visualizer V gdb.default_visualizer
```

Suppose SomeClass is a visualizer class. A lambda expression can be used to instantiate this class for a varobj:

```
(gdb)
-var-set-visualizer V "lambda val: SomeClass()"
^done
```

27.16 GDB/MI Data Manipulation

This section describes the GDB/MI commands that manipulate data: examine memory and registers, evaluate expressions, etc.

The -data-disassemble Command

[-s start-addr -e end-addr]

| [-f filename -l linenum [-n lines]]

-data-disassemble

Synopsis

```
-- mode
Where:
'start-addr'
            is the beginning address (or $pc)
'end-addr'
            is the end address
'filename'
            is the name of the file to disassemble
'linenum'
           is the line number to disassemble around
'lines'
            is the number of disassembly lines to be produced. If it is -1, the whole function
            will be disassembled, in case no end-addr is specified. If end-addr is specified
            as a non-zero value, and lines is lower than the number of disassembly lines
            between start-addr and end-addr, only lines lines are displayed; if lines is higher
            than the number of lines between start-addr and end-addr, only the lines up to
            end-addr are displayed.
'mode'
            is either 0 (meaning only disassembly) or 1 (meaning mixed source and disas-
            sembly).
```

Result

The output for each instruction is composed of four fields:

- Address
- Func-name
- Offset
- Instruction

Note that whatever included in the instruction field, is not manipulated directly by GDB/MI, i.e., it is not possible to adjust its format.

GDB Command

There's no direct mapping from this command to the CLI.

Example

```
Disassemble from the current value of $pc to $pc + 20:
```

Disassemble the whole main function. Line 32 is part of main.

```
-data-disassemble -f basics.c -l 32 -- 0

'done,asm_insns=[

{address="0x000107bc",func-name="main",offset="0",
inst="save %sp, -112, %sp"},

{address="0x000107c0",func-name="main",offset="4",
inst="mov 2, %o0"},

{address="0x000107c4",func-name="main",offset="8",
inst="sethi %hi(0x11800), %o2"},

[...]

{address="0x0001081c",func-name="main",offset="96",inst="ret "},

{address="0x00010820",func-name="main",offset="100",inst="restore "}]

(mdh)
```

Disassemble 3 instructions from the start of main:

```
inst="sethi %hi(0x11800), %o2"}]
   (gdb)
Disassemble 3 instructions from the start of main in mixed mode:
   -data-disassemble -f basics.c -l 32 -n 3 -- 1
   ^done,asm_insns=[
   src_and_asm_line={line="31",
   file="/kwikemart/marge/ezannoni/flathead-dev/devo/gdb/ \
     testsuite/gdb.mi/basics.c",line_asm_insn=[
   {address="0x000107bc",func-name="main",offset="0",
   inst="save %sp, -112, %sp"}]},
   src_and_asm_line={line="32",
   file="/kwikemart/marge/ezannoni/flathead-dev/devo/gdb/ \
    testsuite/gdb.mi/basics.c",line_asm_insn=[
   {address="0x000107c0",func-name="main",offset="4",
   inst="mov 2, %o0"},
   {address="0x000107c4",func-name="main",offset="8",
```

The -data-evaluate-expression Command

Synopsis

```
-data-evaluate-expression expr
```

inst="sethi %hi(0x11800), %o2"}]}]

Evaluate expr as an expression. The expression could contain an inferior function call. The function call will execute synchronously. If the expression contains spaces, it must be enclosed in double quotes.

GDB Command

The corresponding GDB commands are 'print', 'output', and 'call'. In gdbtk only, there's a corresponding 'gdb_eval' command.

Example

In the following example, the numbers that precede the commands are the *tokens* described in Section 27.4 [GDB/MI Command Syntax], page 292. Notice how GDB/MI returns the same tokens in its output.

```
211-data-evaluate-expression A
211^done,value="1"
(gdb)
311-data-evaluate-expression &A
311^done,value="0xefffeb7c"
(gdb)
411-data-evaluate-expression A+3
411^done,value="4"
(gdb)
511-data-evaluate-expression "A + 3"
511^done,value="4"
(gdb)
```

The -data-list-changed-registers Command

Synopsis

```
-data-list-changed-registers
```

Display a list of the registers that have changed.

GDB Command

GDB doesn't have a direct analog for this command; gdbtk has the corresponding command 'gdb_changed_register_list'.

Example

On a PPC MBX board:

```
(gdb)
-exec-continue
running

(gdb)
*stopped,reason="breakpoint-hit",disp="keep",bkptno="1",frame={
func="main",args=[],file="try.c",fullname="/home/foo/bar/try.c",
line="5"}
(gdb)
-data-list-changed-registers
done,changed-registers=["0","1","2","4","5","6","7","8","9",
"10","11","13","14","15","16","17","18","19","20","21","22","23",
"24","25","26","27","28","30","31","64","65","66","67","69"]
(gdb)
```

The -data-list-register-names Command

Synopsis

```
-data-list-register-names [ ( regno )+ ]
```

Show a list of register names for the current target. If no arguments are given, it shows a list of the names of all the registers. If integer numbers are given as arguments, it will print a list of the names of the registers corresponding to the arguments. To ensure consistency between a register name and its number, the output list may include empty register names.

GDB Command

GDB does not have a command which corresponds to '-data-list-register-names'. In gdbtk there is a corresponding command 'gdb_regnames'.

Example

For the PPC MBX board:

```
(gdb)
-data-list-register-names

'done,register-names=["r0","r1","r2","r3","r4","r5","r6","r7",
"r8","r9","r10","r11","r12","r13","r14","r15","r16","r17","r18",
"r19","r20","r21","r22","r23","r24","r25","r26","r27","r28","r29",
"r30","r31","f0","f1","f2","f3","f4","f5","f6","f7","f8","f9",
"f10","f11","f12","f13","f14","f15","f16","f17","f18","f19","f20",
"f21","f22","f23","f24","f25","f26","f27","f28","f29","f30","f31",
"", "pc","ps","cr","lr","ctr","xer"]
(gdb)
```

```
-data-list-register-names 1 2 3
^done,register-names=["r1","r2","r3"]
(gdb)
```

The -data-list-register-values Command

Synopsis

```
-data-list-register-values fmt [ ( regno )*]
```

Display the registers' contents. *fmt* is the format according to which the registers' contents are to be returned, followed by an optional list of numbers specifying the registers to display. A missing list of numbers indicates that the contents of all the registers must be returned.

Allowed formats for fmt are:

- x Hexadecimal
- o Octal
- t Binary
- d Decimal
- r Raw
- N Natural

GDB Command

The corresponding GDB commands are 'info reg', 'info all-reg', and (in gdbtk) 'gdb_fetch_registers'.

Example

For a PPC MBX board (note: line breaks are for readability only, they don't appear in the actual output):

```
(gdb)
-data-list-register-values r 64 65
^done,register-values=[{number="64",value="0xfe00a300"},
{number="65", value="0x00029002"}]
-data-list-register-values x
^done,register-values=[{number="0",value="0xfe0043c8"},
{number="1", value="0x3fff88"}, {number="2", value="0xfffffffe"},
{number="3",value="0x0"},{number="4",value="0xa"},
{number="5",value="0x3fff68"}, {number="6",value="0x3fff58"},
{number="7",value="0xfe011e98"},{number="8",value="0x2"},
{number="9",value="0xfa202820"},{number="10",value="0xfa202808"},
{number="11",value="0x1"},{number="12",value="0x0"},
{number="13", value="0x4544"}, {number="14", value="0xffdfffff"},
{number="15",value="0xffffffff"}, {number="16",value="0xfffffffff"},
{number="17",value="0xefffffed"}, {number="18",value="0xfffffffe"},
{number="19",value="0xffffffff"}, {number="20",value="0xfffffffff"},
{number="21", value="0xffffffff"}, {number="22", value="0xffffffff7"},
{number="23",value="0xffffffff"}, {number="24",value="0xfffffffff"},
{number="25",value="0xffffffff"}, {number="26",value="0xfffffffb"},
{number="27",value="0xffffffff"},{number="28",value="0xf7bffffff"},
```

```
{number="29",value="0x0"},{number="30",value="0xfe010000"},
{number="31",value="0x0"},{number="32",value="0x0"},
{number="33",value="0x0"},{number="34",value="0x0"},
{number="35", value="0x0"}, {number="36", value="0x0"},
{number="37", value="0x0"}, {number="38", value="0x0"},
{number="39", value="0x0"}, {number="40", value="0x0"},
{number="41", value="0x0"}, {number="42", value="0x0"},
{number="43", value="0x0"}, {number="44", value="0x0"},
{number="45", value="0x0"}, {number="46", value="0x0"},
{number="47", value="0x0"}, {number="48", value="0x0"},
{number="49",value="0x0"},{number="50",value="0x0"},
{number="51",value="0x0"},{number="52",value="0x0"},
{number="53",value="0x0"},{number="54",value="0x0"},
{number="55",value="0x0"},{number="56",value="0x0"},
{number="57", value="0x0"}, {number="58", value="0x0"},
{number="59", value="0x0"}, {number="60", value="0x0"},
{number="61", value="0x0"}, {number="62", value="0x0"},
{number="63", value="0x0"}, {number="64", value="0xfe00a300"},
{number="65", value="0x29002"}, {number="66", value="0x202f04b5"},
{number="67",value="0xfe0043b0"},{number="68",value="0xfe00b3e4"},
{number="69",value="0x20002b03"}]
(gdb)
```

The -data-read-memory Command

Synopsis

```
-data-read-memory [ -o byte-offset ]
address word-format word-size
nr-rows nr-cols [ aschar ]
```

where:

'address' An expression specifying the address of the first memory word to be read. Complex expressions containing embedded white space should be quoted using the C convention.

'word-format'

The format to be used to print the memory words. The notation is the same as for GDB's print command (see Section 10.5 [Output Formats], page 97).

'word-size'

The size of each memory word in bytes.

'nr-rows' The number of rows in the output table.

'nr-cols' The number of columns in the output table.

'aschar' If present, indicates that each row should include an ASCII dump. The value of aschar is used as a padding character when a byte is not a member of the printable ASCII character set (printable ASCII characters are those whose code is between 32 and 126, inclusively).

'byte-offset'

An offset to add to the address before fetching memory.

This command displays memory contents as a table of nr-rows by nr-cols words, each word being word-size bytes. In total, nr-rows * nr-cols * word-size bytes are read

(returned as 'total-bytes'). Should less than the requested number of bytes be returned by the target, the missing words are identified using 'N/A'. The number of bytes read from the target is returned in 'nr-bytes' and the starting address used to read memory in 'addr'.

The address of the next/previous row or page is available in 'next-row' and 'prev-row', 'next-page' and 'prev-page'.

GDB Command

The corresponding GDB command is 'x'. gdbtk has 'gdb_get_mem' memory read command.

Example

Read six bytes of memory starting at bytes+6 but then offset by -6 bytes. Format as three rows of two columns. One byte per word. Display each word in hex.

```
(gdb)
9-data-read-memory -o -6 -- bytes+6 x 1 3 2
9^done,addr="0x00001390",nr-bytes="6",total-bytes="6",
next-row="0x00001396",prev-row="0x0000138e",next-page="0x00001396",
prev-page="0x0000138a",memory=[
{addr="0x00001390",data=["0x00","0x01"]},
{addr="0x00001392",data=["0x02","0x03"]},
{addr="0x00001394",data=["0x04","0x05"]}]
(gdb)
```

Read two bytes of memory starting at address shorts + 64 and display as a single word formatted in decimal.

```
(gdb)
5-data-read-memory shorts+64 d 2 1 1
5^done,addr="0x00001510",nr-bytes="2",total-bytes="2",
next-row="0x00001512",prev-row="0x0000150e",
next-page="0x00001512",prev-page="0x0000150e",memory=[
{addr="0x00001510",data=["128"]}]
(gdb)
```

Read thirty two bytes of memory starting at bytes+16 and format as eight rows of four columns. Include a string encoding with 'x' used as the non-printable character.

```
(gdb)
4-data-read-memory bytes+16 x 1 8 4 x
4^done,addr="0x000013a0",nr-bytes="32",total-bytes="32",
next-row="0x000013c0",prev-row="0x0000139c",
next-page="0x000013c0",prev-page="0x00001380",memory=[
{addr="0x000013a0",data=["0x10","0x11","0x12","0x13"],ascii="xxxx"},
{addr="0x000013a4",data=["0x14","0x15","0x16","0x17"],ascii="xxxx"},
{addr="0x000013a8",data=["0x18","0x19","0x1a","0x1b"],ascii="xxxx"},
{addr="0x000013ac",data=["0x1c","0x1d","0x1e","0x1f"],ascii="xxxx"},
{addr="0x000013b0",data=["0x20","0x21","0x22","0x23"],ascii="!\"#"},
{addr="0x000013b4",data=["0x24","0x25","0x26","0x27"],ascii="$%&,"},
{addr="0x000013b8",data=["0x28","0x29","0x2a","0x2b"],ascii="()*+"},
{addr="0x000013bc",data=["0x2c","0x2d","0x2e","0x2f"],ascii=",-./"}]
(gdb)
```

27.17 GDB/MI Tracepoint Commands

The tracepoint commands are not yet implemented.

27.18 GDB/MI Symbol Query Commands

The -symbol-list-lines Command

Synopsis

```
-symbol-list-lines filename
```

Print the list of lines that contain code and their associated program addresses for the given source filename. The entries are sorted in ascending PC order.

GDB Command

There is no corresponding gdb command.

Example

```
(gdb)
-symbol-list-lines basics.c

done,lines=[{pc="0x08048554",line="7"},{pc="0x0804855a",line="8"}]
(gdb)
```

27.19 GDB/MI File Commands

This section describes the GDB/MI commands to specify executable file names and to read in and obtain symbol table information.

The -file-exec-and-symbols Command

Synopsis

```
-file-exec-and-symbols file
```

Specify the executable file to be debugged. This file is the one from which the symbol table is also read. If no file is specified, the command clears the executable and symbol information. If breakpoints are set when using this command with no arguments, GDB will produce error messages. Otherwise, no output is produced, except a completion notification.

GDB Command

The corresponding GDB command is 'file'.

Example

```
(gdb)
-file-exec-and-symbols /kwikemart/marge/ezannoni/TRUNK/mbx/hello.mbx
^done
(gdb)
```

The -file-exec-file Command

Synopsis

```
-file-exec-file file
```

Specify the executable file to be debugged. Unlike '-file-exec-and-symbols', the symbol table is *not* read from this file. If used without argument, GDB clears the information about the executable file. No output is produced, except a completion notification.

GDB Command

The corresponding gdb command is 'exec-file'.

Example

```
(gdb)
-file-exec-file /kwikemart/marge/ezannoni/TRUNK/mbx/hello.mbx
^done
(gdb)
```

The -file-list-exec-source-file Command

Synopsis

```
-file-list-exec-source-file
```

List the line number, the current source file, and the absolute path to the current source file for the current executable. The macro information field has a value of '1' or '0' depending on whether or not the file includes preprocessor macro information.

GDB Command

The gdb equivalent is 'info source'

Example

```
(gdb)
123-file-list-exec-source-file
123^done,line="1",file="foo.c",fullname="/home/bar/foo.c,macro-info="1"
(gdb)
```

The -file-list-exec-source-files Command

Synopsis

```
-file-list-exec-source-files
```

List the source files for the current executable.

It will always output the filename, but only when GDB can find the absolute file name of a source file, will it output the fullname.

GDB Command

The GDB equivalent is 'info sources'. gdbtk has an analogous command 'gdb_listfiles'.

Example

```
(gdb)
-file-list-exec-source-files
^done,files=[
{file=foo.c,fullname=/home/foo.c},
{file=/home/bar.c,fullname=/home/bar.c},
{file=gdb_could_not_find_fullpath.c}]
(gdb)
```

The -file-symbol-file Command

Synopsis

```
-file-symbol-file file
```

Read symbol table info from the specified *file* argument. When used without arguments, clears gdb's symbol table info. No output is produced, except for a completion notification.

GDB Command

The corresponding GDB command is 'symbol-file'.

Example

```
(gdb)
-file-symbol-file /kwikemart/marge/ezannoni/TRUNK/mbx/hello.mbx
^done
(gdb)
```

27.20 GDB/MI Target Manipulation Commands

The -target-attach Command

Synopsis

```
-target-attach pid | gid | file
```

Attach to a process *pid* or a file *file* outside of GDB, or a thread group *gid*. If attaching to a thread group, the id previously returned by '-list-thread-groups --available' must be used.

GDB Command

The corresponding GDB command is 'attach'.

Example

```
(gdb)
-target-attach 34
=thread-created,id="1"
*stopped,thread-id="1",frame={addr="0xb7f7e410",func="bar",args=[]}
^done
(gdb)
```

The -target-detach Command

Synopsis

```
-target-detach [ pid | gid ]
```

Detach from the remote target which normally resumes its execution. If either *pid* or *gid* is specified, detaches from either the specified process, or specified thread group. There's no output.

GDB Command

The corresponding GDB command is 'detach'.

Example

```
(gdb)
-target-detach
^done
(gdb)
```

The -target-disconnect Command

Synopsis

```
-target-disconnect
```

Disconnect from the remote target. There's no output and the target is generally not resumed.

GDB Command

The corresponding GDB command is 'disconnect'.

Example

```
(gdb)
-target-disconnect
^done
(gdb)
```

The -target-download Command

Synopsis

```
-target-download
```

Loads the executable onto the remote target. It prints out an update message every half second, which includes the fields:

```
'section' The name of the section.
```

```
'section-sent'
```

The size of what has been sent so far for that section.

```
'section-size'
```

The size of the section.

'total-sent'

The total size of what was sent so far (the current and the previous sections).

'total-size'

The size of the overall executable to download.

Each message is sent as status record (see Section 27.4.2 [GDB/MI Output Syntax], page 293).

In addition, it prints the name and size of the sections, as they are downloaded. These messages include the following fields:

```
'section' The name of the section.
```

```
'section-size'
```

The size of the section.

'total-size'

The size of the overall executable to download.

At the end, a summary is printed.

GDB Command

The corresponding GDB command is 'load'.

Example

Note: each status message appears on a single line. Here the messages have been broken down so that they can fit onto a page.

```
(gdb)
-target-download
+download, {section=".text", section-size="6668", total-size="9880"}
+download, {section=".text", section-sent="512", section-size="6668",
total-sent="512",total-size="9880"}
+download, {section=".text", section-sent="1024", section-size="6668",
total-sent="1024",total-size="9880"}
+download, {section=".text", section-sent="1536", section-size="6668",
total-sent="1536",total-size="9880"}
+download, {section=".text", section-sent="2048", section-size="6668",
total-sent="2048",total-size="9880"}
+download, {section=".text", section-sent="2560", section-size="6668",
total-sent="2560",total-size="9880"}
+download, {section=".text", section-sent="3072", section-size="6668",
total-sent="3072",total-size="9880"}
+download, {section=".text", section-sent="3584", section-size="6668",
total-sent="3584",total-size="9880"}
+download, {section=".text", section-sent="4096", section-size="6668",
total-sent="4096",total-size="9880"}
+download, {section=".text", section-sent="4608", section-size="6668",
total-sent="4608",total-size="9880"}
+download, {section=".text", section-sent="5120", section-size="6668",
total-sent="5120",total-size="9880"}
+download, {section=".text", section-sent="5632", section-size="6668",
total-sent="5632",total-size="9880"}
+download, {section=".text", section-sent="6144", section-size="6668",
total-sent="6144",total-size="9880"}
+download, {section=".text", section-sent="6656", section-size="6668",
total-sent="6656",total-size="9880"}
+download, {section=".init", section-size="28", total-size="9880"}
+download, {section=".fini", section-size="28", total-size="9880"}
+download, {section=".data", section-size="3156", total-size="9880"}
+download, {section=".data", section-sent="512", section-size="3156",
total-sent="7236",total-size="9880"}
+download, {section=".data", section-sent="1024", section-size="3156",
total-sent="7748",total-size="9880"}
+download, {section=".data", section-sent="1536", section-size="3156",
total-sent="8260",total-size="9880"}
+download, {section=".data", section-sent="2048", section-size="3156",
total-sent="8772",total-size="9880"}
+download, {section=".data", section-sent="2560", section-size="3156",
total-sent="9284",total-size="9880"}
+download, {section=".data", section-sent="3072", section-size="3156",
total-sent="9796",total-size="9880"}
^done,address="0x10004",load-size="9880",transfer-rate="6586",
```

```
write-rate="429" (gdb)
```

GDB Command

No equivalent.

Example

N.A.

The -target-select Command

Synopsis

```
-target-select type parameters ...
```

Connect gdb to the remote target. This command takes two args:

'type' The type of target, for instance 'remote', etc.

'parameters'

Device names, host names and the like. See Section 19.2 [Commands for Managing Targets], page 200, for more details.

The output is a connection notification, followed by the address at which the target program is, in the following form:

```
^connected,addr="address",func="function name",
   args=[arg list]
```

GDB Command

The corresponding GDB command is 'target'.

Example

```
(gdb)
-target-select remote /dev/ttya
^connected,addr="0xfe00a300",func="??",args=[]
(gdb)
```

27.21 GDB/MI File Transfer Commands

The -target-file-put Command

Synopsis

```
-target-file-put hostfile targetfile
```

Copy file hostfile from the host system (the machine running GDB) to targetfile on the target system.

GDB Command

The corresponding GDB command is 'remote put'.

Example

```
(gdb)
-target-file-put localfile remotefile
^done
(gdb)
```

The -target-file-get Command

Synopsis

```
-target-file-get targetfile hostfile
```

Copy file targetfile from the target system to hostfile on the host system.

GDB Command

The corresponding GDB command is 'remote get'.

Example

```
(gdb)
-target-file-get remotefile localfile
^done
(gdb)
```

The -target-file-delete Command

Synopsis

```
-target-file-delete targetfile
```

Delete targetfile from the target system.

gdb Command

The corresponding GDB command is 'remote delete'.

Example

```
(gdb)
-target-file-delete remotefile
^done
(gdb)
```

27.22 Miscellaneous GDB/MI Commands

The -gdb-exit Command

Synopsis

```
-gdb-exit
```

Exit GDB immediately.

GDB Command

Approximately corresponds to 'quit'.

Example

```
(gdb)
-gdb-exit
^exit
```

The -gdb-set Command

Synopsis

```
-gdb-set
```

Set an internal gdb variable.

GDB Command

The corresponding GDB command is 'set'.

Example

```
(gdb)
-gdb-set $foo=3
^done
(gdb)
```

The -gdb-show Command

Synopsis

```
-gdb-show
```

Show the current value of a GDB variable.

GDB Command

The corresponding GDB command is 'show'.

Example

```
(gdb)
-gdb-show annotate
^done,value="0"
(gdb)
```

The -gdb-version Command

Synopsis

```
-gdb-version
```

Show version information for GDB. Used mostly in testing.

gdb Command

The GDB equivalent is 'show version'. GDB by default shows this information when you start an interactive session.

Example

```
(gdb)
-gdb-version
```

```
"GNU gdb 5.2.1
"Copyright 2000 Free Software Foundation, Inc.
"GDB is free software, covered by the GNU General Public License, and "you are welcome to change it and/or distribute copies of it under "certain conditions.
"Type "show copying" to see the conditions.
"There is absolutely no warranty for GDB. Type "show warranty" for "details.
"This GDB was configured as "--host=sparc-sun-solaris2.5.1 --target=ppc-eabi".
"done (gdb)
```

The -list-features Command

Returns a list of particular features of the MI protocol that this version of gdb implements. A feature can be a command, or a new field in an output of some command, or even an important bugfix. While a frontend can sometimes detect presence of a feature at runtime, it is easier to perform detection at debugger startup.

The command returns a list of strings, with each string naming an available feature. Each returned string is just a name, it does not have any internal structure. The list of possible feature names is given below.

Example output:

```
(gdb) -list-features
^done,result=["feature1","feature2"]
```

The current list of features is:

'frozen-varobjs'

Indicates presence of the -var-set-frozen command, as well as possible presense of the frozen field in the output of -varobj-create.

'pending-breakpoints'

Indicates presence of the '-f' option to the -break-insert command.

'python' Indicates presence of Python scripting support, Python-based pretty-printing commands, and possible presence of the 'display_hint' field in the output of -var-list-children

'thread-info'

Indicates presence of the -thread-info command.

The -list-target-features Command

Returns a list of particular features that are supported by the target. Those features affect the permitted MI commands, but unlike the features reported by the <code>-list-features</code> command, the features depend on which target GDB is using at the moment. Whenever a target can change, due to commands such as <code>-target-select</code>, <code>-target-attach</code> or <code>-exec-run</code>, the list of target features may change, and the frontend should obtain it again. Example output:

```
(gdb) -list-features
^done,result=["async"]
```

The current list of features is:

'async' Indicates that the target is capable of asynchronous command execution, which means that GDB will accept further commands while the target is running.

The -list-thread-groups Command

Synopsis

```
-list-thread-groups [ --available ] [ --recurse 1 ] [ group ... ]
```

Lists thread groups (see Section 27.3.3 [Thread groups], page 291). When a single thread group is passed as the argument, lists the children of that group. When several thread group are passed, lists information about those thread groups. Without any parameters, lists information about all top-level thread groups.

Normally, thread groups that are being debugged are reported. With the '--available' option, GDB reports thread groups available on the target.

The output of this command may have either a 'threads' result or a 'groups' result. The 'thread' result has a list of tuples as value, with each tuple describing a thread (see Section 27.7.5 [GDB/MI Thread Information], page 298). The 'groups' result has a list of tuples as value, each tuple describing a thread group. If top-level groups are requested (that is, no parameter is passed), or when several groups are passed, the output always has a 'groups' result. The format of the 'group' result is described below.

To reduce the number of roundtrips it's possible to list thread groups together with their children, by passing the '--recurse' option and the recursion depth. Presently, only recursion depth of 1 is permitted. If this option is present, then every reported thread group will also include its children, either as 'group' or 'threads' field.

In general, any combination of option and parameters is permitted, with the following caveats:

- When a single thread group is passed, the output will typically be the 'threads' result. Because threads may not contain anything, the 'recurse' option will be ignored.
- When the '--available' option is passed, limited information may be available. In particular, the list of threads of a process might be inaccessible. Further, specifying specific thread groups might not give any performance advantage over listing all thread groups. The frontend should assume that '-list-thread-groups --available' is always an expensive operation and cache the results.

The 'groups' result is a list of tuples, where each tuple may have the following fields:

id Identifier of the thread group. This field is always present.

The type of the thread group. At present, only 'process' is a valid type.

The target-specific process identifier. This field is only present for thread groups of type 'process'.

num_children

The number of children this thread group has. This field may be absent for an available thread group.

This field has a list of tuples as value, each tuple describing a thread. It may be present if the '--recurse' option is specified, and it's actually possible to obtain the threads.

cores

This field is a list of integers, each identifying a core that one thread of the group is running on. This field may be absent if such information is not available.

Example

```
gdb
-list-thread-groups
^done,groups=[{id="17",type="process",pid="yyy",num_children="2"}]
-list-thread-groups 17
^done,threads=[{id="2",target-id="Thread 0xb7e14b90 (LWP 21257)",
  frame={level="0",addr="0xffffe410",func="__kernel_vsyscall",args=[]},state="running"},
{id="1",target-id="Thread 0xb7e156b0 (LWP 21254)",
  frame={level="0",addr="0x0804891f",func="foo",args=[{name="i",value="10"}],
          file="/tmp/a.c",fullname="/tmp/a.c",line="158"},state="running"}]]
-list-thread-groups --available
^done,groups=[{id="17",type="process",pid="yyy",num_children="2",cores=[1,2]}]
-list-thread-groups --available --recurse 1
^done,groups=[{id="17", types="process",pid="yyy",num_children="2",cores=[1,2],
                threads=[{id="1",target-id="Thread 0xb7e14b90",cores=[1]},
                         {id="2",target-id="Thread 0xb7e14b90",cores=[2]}]},..]
-list-thread-groups --available --recurse 1 17 18
^done,groups=[{id="17", types="process",pid="yyy",num_children="2",cores=[1,2],
               threads=[{id="1",target-id="Thread 0xb7e14b90",cores=[1]},
                        {id="2",target-id="Thread 0xb7e14b90",cores=[2]}}},...]
```

The -interpreter-exec Command

Synopsis

```
-interpreter-exec interpreter command
```

Execute the specified command in the given interpreter.

gdb Command

The corresponding GDB command is 'interpreter-exec'.

Example

```
(gdb)
-interpreter-exec console "break main"
&"During symbol reading, couldn't parse type; debugger out of date?.\n"
&"During symbol reading, bad structure-type format.\n"
~"Breakpoint 1 at 0x8074fc6: file ../../src/gdb/main.c, line 743.\n"
^done
(gdb)
```

The -inferior-tty-set Command

Synopsis

```
-inferior-tty-set /dev/pts/1
```

Set terminal for future runs of the program being debugged.

gdb Command

The corresponding GDB command is 'set inferior-tty' /dev/pts/1.

Example

```
(gdb)
-inferior-tty-set /dev/pts/1
^done
(gdb)
```

The -inferior-tty-show Command

Synopsis

```
-inferior-tty-show
```

Show terminal for future runs of program being debugged.

gdb Command

The corresponding GDB command is 'show inferior-tty'.

Example

```
(gdb)
-inferior-tty-set /dev/pts/1
^done
(gdb)
-inferior-tty-show
^done,inferior_tty_terminal="/dev/pts/1"
(gdb)
```

The -enable-timings Command

Synopsis

```
-enable-timings [yes | no]
```

Toggle the printing of the wallclock, user and system times for an MI command as a field in its output. This command is to help frontend developers optimize the performance of their code. No argument is equivalent to 'yes'.

gdb Command

No equivalent.

Example

```
(gdb)
-enable-timings
^done
(gdb)
-break-insert main
^done,bkpt={number="1",type="breakpoint",disp="keep",enabled="y",addr="0x080484ed",func="main",file="myprog.c",fullname="/home/nickrob/myprog.c",line="73",times="0"},
time={wallclock="0.05185",user="0.00800",system="0.00000"}
(gdb)
-enable-timings no
^done
(gdb)
-exec-run
^running
```

```
(gdb)
*stopped,reason="breakpoint-hit",disp="keep",bkptno="1",thread-id="0",
frame={addr="0x080484ed",func="main",args=[{name="argc",value="1"},
{name="argv",value="0xbfb60364"}],file="myprog.c",
fullname="/home/nickrob/myprog.c",line="73"}
(gdb)
```

28 GDB Annotations

This chapter describes annotations in GDB. Annotations were designed to interface GDB to graphical user interfaces or other similar programs which want to interact with GDB at a relatively high level.

The annotation mechanism has largely been superseded by GDB/MI (see Chapter 27 [GDB/MI], page 289).

28.1 What is an Annotation?

Annotations start with a newline character, two 'control-z' characters, and the name of the annotation. If there is no additional information associated with this annotation, the name of the annotation is followed immediately by a newline. If there is additional information, the name of the annotation is followed by a space, the additional information, and a newline. The additional information cannot contain newline characters.

Any output not beginning with a newline and two 'control-z' characters denotes literal output from gdb. Currently there is no need for gdb to output a newline followed by two 'control-z' characters, but if there was such a need, the annotations could be extended with an 'escape' annotation which means those three characters as output.

The annotation level, which is specified using the '--annotate' command line option (see Section 2.1.2 [Mode Options], page 13), controls how much information GDB prints together with its prompt, values of expressions, source lines, and other types of output. Level 0 is for no annotations, level 1 is for use when GDB is run as a subprocess of GNU Emacs, level 3 is the maximum annotation suitable for programs that control GDB, and level 2 annotations have been made obsolete (see Section "Limitations of the Annotation Interface" in GDB's Obsolete Annotations).

set annotate level

The GDB command set annotate sets the level of annotations to the specified level.

show annotate

Show the current annotation level.

This chapter describes level 3 annotations.

A simple example of starting up GDB with annotations is:

```
$ gdb --annotate=3
GNU gdb 6.0
Copyright 2003 Free Software Foundation, Inc.
GDB is free software, covered by the GNU General Public License,
and you are welcome to change it and/or distribute copies of it
under certain conditions.
Type "show copying" to see the conditions.
There is absolutely no warranty for GDB. Type "show warranty"
for details.
This GDB was configured as "i386-pc-linux-gnu"

2^2pre-prompt
(gdb)
2^2prompt
quit
```

```
^Z^Zpost-prompt
$
```

Here 'quit' is input to GDB; the rest is output from GDB. The three lines beginning '^Z^Z' (where '^Z' denotes a 'control-z' character) are annotations; the rest is output from GDB.

28.2 The Server Prefix

If you prefix a command with 'server' then it will not affect the command history, nor will it affect gdb's notion of which command to repeat if RET is pressed on a line by itself. This means that commands can be run behind a user's back by a front-end in a transparent manner.

The server prefix does not affect the recording of values into the value history; to print a value without recording it into the value history, use the output command instead of the print command.

Using this prefix also disables confirmation requests (see [confirmation requests], page 250).

28.3 Annotation for gdb Input

When GDB prompts for input, it annotates this fact so it is possible to know when to send output, when the output from a given command is over, etc.

Different kinds of input each have a different *input type*. Each input type has three annotations: a pre- annotation, which denotes the beginning of any prompt which is being output, a plain annotation, which denotes the end of the prompt, and then a post- annotation which denotes the end of any echo which may (or may not) be associated with the input. For example, the prompt input type features the following annotations:

```
^Z^Zpre-prompt
^Z^Zprompt
```

^Z^Zpost-prompt

The input types are

when gdb is prompting for a command (the main gdb prompt).

commands When gdb prompts for a set of commands, like in the commands command. The annotations are repeated for each command which is input.

overload-choice

When gdb wants the user to select between various overloaded functions.

query When GDB wants the user to confirm a potentially dangerous operation.

prompt-for-continue

When GDB is asking the user to press return to continue. Note: Don't expect this to work well; instead use set height 0 to disable prompting. This is because the counting of lines is buggy in the presence of annotations.

28.4 Errors

^Z^Zquit

This annotation occurs right before gdb responds to an interrupt.

^Z^Zerror

This annotation occurs right before GDB responds to an error.

Quit and error annotations indicate that any annotations which GDB was in the middle of may end abruptly. For example, if a value-history-begin annotation is followed by a error, one cannot expect to receive the matching value-history-end. One cannot expect not to receive it either, however; an error annotation does not necessarily mean that GDB is immediately returning all the way to the top level.

A quit or error annotation may be preceded by

^Z^Zerror-begin

Any output between that and the quit or error annotation is the error message.

Warning messages are not yet annotated.

28.5 Invalidation Notices

The following annotations say that certain pieces of state may have changed.

^Z^Zframes-invalid

The frames (for example, output from the backtrace command) may have changed.

^Z^Zbreakpoints-invalid

The breakpoints may have changed. For example, the user just added or deleted a breakpoint.

28.6 Running the Program

When the program starts executing due to a GDB command such as step or continue,

^Z^Zstarting

is output. When the program stops,

^Z^Zstopped

is output. Before the **stopped** annotation, a variety of annotations describe how the program stopped.

^Z^Zexited exit-status

The program exited, and exit-status is the exit status (zero for successful exit, otherwise nonzero).

^Z^Zsignalled

The program exited with a signal. After the ^Z^Zsignalled, the annotation continues:

intro-text
^Z^Zsignal-name
name
^Z^Zsignal-name-end
middle-text
^Z^Zsignal-string

```
string
^Z^Zsignal-string-end
end-text
```

where name is the name of the signal, such as SIGILL or SIGSEGV, and string is the explanation of the signal, such as Illegal Instruction or Segmentation fault. intro-text, middle-text, and end-text are for the user's benefit and have no particular format.

^Z^Zsignal

The syntax of this annotation is just like signalled, but GDB is just saying that the program received the signal, not that it was terminated with it.

^Z^Zbreakpoint number

The program hit breakpoint number number.

^Z^Zwatchpoint number

The program hit watchpoint number number.

28.7 Displaying Source

The following annotation is used instead of displaying source code:

```
^Z^Zsource filename:line:character:middle:addr
```

where filename is an absolute file name indicating which source file, line is the line number within that file (where 1 is the first line in the file), character is the character position within the file (where 0 is the first character in the file) (for most debug formats this will necessarily point to the beginning of a line), middle is 'middle' if addr is in the middle of the line, or 'beg' if addr is at the beginning of the line, and addr is the address in the target program associated with the source which is being displayed. addr is in the form '0x' followed by one or more lowercase hex digits (note that this does not depend on the language).

29 JIT Compilation Interface

This chapter documents GDB's just-in-time (JIT) compilation interface. A JIT compiler is a program or library that generates native executable code at runtime and executes it, usually in order to achieve good performance while maintaining platform independence.

Programs that use JIT compilation are normally difficult to debug because portions of their code are generated at runtime, instead of being loaded from object files, which is where GDB normally finds the program's symbols and debug information. In order to debug programs that use JIT compilation, GDB has an interface that allows the program to register in-memory symbol files with GDB at runtime.

If you are using GDB to debug a program that uses this interface, then it should work transparently so long as you have not stripped the binary. If you are developing a JIT compiler, then the interface is documented in the rest of this chapter. At this time, the only known client of this interface is the LLVM JIT.

Broadly speaking, the JIT interface mirrors the dynamic loader interface. The JIT compiler communicates with GDB by writing data into a global variable and calling a fuction at a well-known symbol. When GDB attaches, it reads a linked list of symbol files from the global variable to find existing code, and puts a breakpoint in the function so that it can find out about additional code.

29.1 JIT Declarations

These are the relevant struct declarations that a C program should include to implement the interface:

```
typedef enum
  JIT_NOACTION = O,
  JIT_REGISTER_FN,
  JIT_UNREGISTER_FN
} jit_actions_t;
struct jit_code_entry
  struct jit_code_entry *next_entry;
 struct jit_code_entry *prev_entry;
 const char *symfile_addr;
 uint64_t symfile_size;
struct jit_descriptor
 uint32_t version;
 /* This type should be jit_actions_t, but we use uint32_t
     to be explicit about the bitwidth. */
 uint32_t action_flag;
 struct jit_code_entry *relevant_entry;
 struct jit_code_entry *first_entry;
};
/* GDB puts a breakpoint in this function. */
void __attribute__((noinline)) __jit_debug_register_code() { };
```

```
/* Make sure to specify the version statically, because the
  debugger may check the version before we can set it. */
struct jit_descriptor __jit_debug_descriptor = { 1, 0, 0, 0 };
```

If the JIT is multi-threaded, then it is important that the JIT synchronize any modifications to this global data properly, which can easily be done by putting a global mutex around modifications to these structures.

29.2 Registering Code

To register code with GDB, the JIT should follow this protocol:

- Generate an object file in memory with symbols and other desired debug information. The file must include the virtual addresses of the sections.
- Create a code entry for the file, which gives the start and size of the symbol file.
- Add it to the linked list in the JIT descriptor.
- Point the relevant_entry field of the descriptor at the entry.
- Set action_flag to JIT_REGISTER and call __jit_debug_register_code.

When GDB is attached and the breakpoint fires, GDB uses the relevant_entry pointer so it doesn't have to walk the list looking for new code. However, the linked list must still be maintained in order to allow GDB to attach to a running process and still find the symbol files.

29.3 Unregistering Code

If code is freed, then the JIT should use the following protocol:

- Remove the code entry corresponding to the code from the linked list.
- Point the relevant_entry field of the descriptor at the code entry.
- Set action_flag to JIT_UNREGISTER and call __jit_debug_register_code.

If the JIT frees or recompiles code without unregistering it, then gdb and the JIT will leak the memory used for the associated symbol files.

30 Reporting Bugs in GDB

Your bug reports play an essential role in making gdb reliable.

Reporting a bug may help you by bringing a solution to your problem, or it may not. But in any case the principal function of a bug report is to help the entire community by making the next version of GDB work better. Bug reports are your contribution to the maintenance of GDB.

In order for a bug report to serve its purpose, you must include the information that enables us to fix the bug.

30.1 Have You Found a Bug?

If you are not sure whether you have found a bug, here are some guidelines:

- If the debugger gets a fatal signal, for any input whatever, that is a gdb bug. Reliable debuggers never crash.
- If GDB produces an error message for valid input, that is a bug. (Note that if you're cross debugging, the problem may also be somewhere in the connection to the target.)
- If GDB does not produce an error message for invalid input, that is a bug. However, you should note that your idea of "invalid input" might be our idea of "an extension" or "support for traditional practice".
- If you are an experienced user of debugging tools, your suggestions for improvement of GDB are welcome in any case.

30.2 How to Report Bugs

A number of companies and individuals offer support for GNU products. If you obtained GDB from a support organization, we recommend you contact that organization first.

You can find contact information for many support companies and individuals in the file 'etc/SERVICE' in the GNU Emacs distribution.

In any event, we also recommend that you submit bug reports for GDB to https://support.codesourcery.com/GNUToolchain/.

The fundamental principle of reporting bugs usefully is this: **report all the facts**. If you are not sure whether to state a fact or leave it out, state it!

Often people omit facts because they think they know what causes the problem and assume that some details do not matter. Thus, you might assume that the name of the variable you use in an example does not matter. Well, probably it does not, but one cannot be sure. Perhaps the bug is a stray memory reference which happens to fetch from the location where that name is stored in memory; perhaps, if the name were different, the contents of that location would fool the debugger into doing the right thing despite the bug. Play it safe and give a specific, complete example. That is the easiest thing for you to do, and the most helpful.

Keep in mind that the purpose of a bug report is to enable us to fix the bug. It may be that the bug has been reported previously, but neither you nor we can know that unless your bug report is complete and self-contained.

Sometimes people give a few sketchy facts and ask, "Does this ring a bell?" Those bug reports are useless, and we urge everyone to refuse to respond to them except to chide the sender to report bugs properly.

To enable us to fix the bug, you should include all these things:

- The version of GDB. GDB announces it if you start with no arguments; you can also print it at any time using show version.
 - Without this, we will not know whether there is any point in looking for the bug in the current version of GDB.
- The type of machine you are using, and the operating system name and version number.
- What compiler (and its version) was used to compile GDB—e.g. "gcc-2.8.1".
- What compiler (and its version) was used to compile the program you are debugging—e.g. "gcc-2.8.1", or "HP92453-01 A.10.32.03 HP C Compiler". For GCC, you can say gcc --version to get this information; for other compilers, see the documentation for those compilers.
- The command arguments you gave the compiler to compile your example and observe the bug. For example, did you use '-0'? To guarantee you will not omit something important, list them all. A copy of the Makefile (or the output from make) is sufficient.

 If we were to try to guess the arguments, we would probably guess wrong and then we
 - If we were to try to guess the arguments, we would probably guess wrong and then we might not encounter the bug.
- A complete input script, and all necessary source files, that will reproduce the bug.
- A description of what behavior you observe that you believe is incorrect. For example, "It gets a fatal signal."

Of course, if the bug is that GDB gets a fatal signal, then we will certainly notice it. But if the bug is incorrect output, we might not notice unless it is glaringly wrong. You might as well not give us a chance to make a mistake.

Even if the problem you experience is a fatal signal, you should still say so explicitly. Suppose something strange is going on, such as, your copy of GDB is out of synch, or you have encountered a bug in the C library on your system. (This has happened!) Your copy might crash and ours would not. If you told us to expect a crash, then when ours fails to crash, we would know that the bug was not happening for us. If you had not told us to expect a crash, then we would not be able to draw any conclusion from our observations.

To collect all this information, you can use a session recording program such as script, which is available on many Unix systems. Just run your GDB session inside script and then include the 'typescript' file with your bug report.

Another way to record a GDB session is to run GDB inside Emacs and then save the entire buffer to a file.

• If you wish to suggest changes to the GDB source, send us context diffs. If you even discuss something in the GDB source, refer to it by context, not by line number.

The line numbers in our development sources will not match those in your sources. Your line numbers would convey no useful information to us.

Here are some things that are not necessary:

• A description of the envelope of the bug.

Often people who encounter a bug spend a lot of time investigating which changes to the input file will make the bug go away and which changes will not affect it.

This is often time consuming and not very useful, because the way we will find the bug is by running a single example under the debugger with breakpoints, not by pure deduction from a series of examples. We recommend that you save your time for something else.

Of course, if you can find a simpler example to report *instead* of the original one, that is a convenience for us. Errors in the output will be easier to spot, running under the debugger will take less time, and so on.

However, simplification is not vital; if you do not want to do this, report the bug anyway and send us the entire test case you used.

• A patch for the bug.

A patch for the bug does help us if it is a good one. But do not omit the necessary information, such as the test case, on the assumption that a patch is all we need. We might see problems with your patch and decide to fix the problem another way, or we might not understand it at all.

Sometimes with a program as complicated as GDB it is very hard to construct an example that will make the program follow a certain path through the code. If you do not send us the example, we will not be able to construct one, so we will not be able to verify that the bug is fixed.

And if we cannot understand what bug you are trying to fix, or why your patch should be an improvement, we will not install it. A test case will help us to understand.

• A guess about what the bug is or what it depends on.

Such guesses are usually wrong. Even we cannot guess right about such things without first using the debugger to find the facts.

31 Command Line Editing

This chapter describes the basic features of the GNU command line editing interface.

31.1 Introduction to Line Editing

The following paragraphs describe the notation used to represent keystrokes.

The text C-k is read as 'Control-K' and describes the character produced when the K key is pressed while the Control key is depressed.

The text *M-k* is read as 'Meta-K' and describes the character produced when the Meta key (if you have one) is depressed, and the K key is pressed. The Meta key is labeled ALT on many keyboards. On keyboards with two keys labeled ALT (usually to either side of the space bar), the ALT on the left side is generally set to work as a Meta key. The ALT key on the right may also be configured to work as a Meta key or may be configured as some other modifier, such as a Compose key for typing accented characters.

If you do not have a Meta or ALT key, or another key working as a Meta key, the identical keystroke can be generated by typing ESC *first*, and then typing K. Either process is known as *metafying* the K key.

The text M-C-k is read as 'Meta-Control-k' and describes the character produced by metafying C-k.

In addition, several keys have their own names. Specifically, DEL, ESC, LFD, SPC, RET, and TAB all stand for themselves when seen in this text, or in an init file (see Section 31.3 [Readline Init File], page 368). If your keyboard lacks a LFD key, typing C-J will produce the desired character. The RET key may be labeled RETURN or ENTER on some keyboards.

31.2 Readline Interaction

Often during an interactive session you type in a long line of text, only to notice that the first word on the line is misspelled. The Readline library gives you a set of commands for manipulating the text as you type it in, allowing you to just fix your typo, and not forcing you to retype the majority of the line. Using these editing commands, you move the cursor to the place that needs correction, and delete or insert the text of the corrections. Then, when you are satisfied with the line, you simply press RET. You do not have to be at the end of the line to press RET; the entire line is accepted regardless of the location of the cursor within the line.

31.2.1 Readline Bare Essentials

In order to enter characters into the line, simply type them. The typed character appears where the cursor was, and then the cursor moves one space to the right. If you mistype a character, you can use your erase character to back up and delete the mistyped character.

Sometimes you may mistype a character, and not notice the error until you have typed several other characters. In that case, you can type C-b to move the cursor to the left, and then correct your mistake. Afterwards, you can move the cursor to the right with C-f.

When you add text in the middle of a line, you will notice that characters to the right of the cursor are 'pushed over' to make room for the text that you have inserted. Likewise,

when you delete text behind the cursor, characters to the right of the cursor are 'pulled back' to fill in the blank space created by the removal of the text. A list of the bare essentials for editing the text of an input line follows.

C-b Move back one character.

C-f Move forward one character.

DEL or BACKSPACE

Delete the character to the left of the cursor.

C-d Delete the character underneath the cursor.

Printing characters

Insert the character into the line at the cursor.

 $C-_{-}$ or C-x C-u

Undo the last editing command. You can undo all the way back to an empty line

(Depending on your configuration, the BACKSPACE key be set to delete the character to the left of the cursor and the DEL key set to delete the character underneath the cursor, like C-d, rather than the character to the left of the cursor.)

31.2.2 Readline Movement Commands

The above table describes the most basic keystrokes that you need in order to do editing of the input line. For your convenience, many other commands have been added in addition to C-b, C-f, C-d, and DEL. Here are some commands for moving more rapidly about the line.

C-a Move to the start of the line.

C-e Move to the end of the line.

M-f Move forward a word, where a word is composed of letters and digits.

M-b Move backward a word.

C-1 Clear the screen, reprinting the current line at the top.

Notice how C-f moves forward a character, while M-f moves forward a word. It is a loose convention that control keystrokes operate on characters while meta keystrokes operate on words.

31.2.3 Readline Killing Commands

Killing text means to delete the text from the line, but to save it away for later use, usually by yanking (re-inserting) it back into the line. ('Cut' and 'paste' are more recent jargon for 'kill' and 'yank'.)

If the description for a command says that it 'kills' text, then you can be sure that you can get the text back in a different (or the same) place later.

When you use a kill command, the text is saved in a *kill-ring*. Any number of consecutive kills save all of the killed text together, so that when you yank it back, you get it all. The kill ring is not line specific; the text that you killed on a previously typed line is available to be yanked back later, when you are typing another line.

Here is the list of commands for killing text.

C-k Kill the text from the current cursor position to the end of the line.

M-d Kill from the cursor to the end of the current word, or, if between words, to the end of the next word. Word boundaries are the same as those used by M-f.

M-DEL Kill from the cursor the start of the current word, or, if between words, to the start of the previous word. Word boundaries are the same as those used by M-b.

C-w Kill from the cursor to the previous whitespace. This is different than M-DEL because the word boundaries differ.

Here is how to yank the text back into the line. Yanking means to copy the most-recently-killed text from the kill buffer.

C-y Yank the most recently killed text back into the buffer at the cursor.

M-y Rotate the kill-ring, and yank the new top. You can only do this if the prior command is C-y or M-y.

31.2.4 Readline Arguments

You can pass numeric arguments to Readline commands. Sometimes the argument acts as a repeat count, other times it is the *sign* of the argument that is significant. If you pass a negative argument to a command which normally acts in a forward direction, that command will act in a backward direction. For example, to kill text back to the start of the line, you might type 'M-- C-k'.

The general way to pass numeric arguments to a command is to type meta digits before the command. If the first 'digit' typed is a minus sign ('-'), then the sign of the argument will be negative. Once you have typed one meta digit to get the argument started, you can type the remainder of the digits, and then the command. For example, to give the C-d command an argument of 10, you could type 'M-1 0 C-d', which will delete the next ten characters on the input line.

31.2.5 Searching for Commands in the History

Readline provides commands for searching through the command history for lines containing a specified string. There are two search modes: *incremental* and *non-incremental*.

Incremental searches begin before the user has finished typing the search string. As each character of the search string is typed, Readline displays the next entry from the history matching the string typed so far. An incremental search requires only as many characters as needed to find the desired history entry. To search backward in the history for a particular string, type C-r. Typing C-s searches forward through the history. The characters present in the value of the isearch-terminators variable are used to terminate an incremental search. If that variable has not been assigned a value, the ESC and C-J characters will terminate an incremental search. C-g will abort an incremental search and restore the original line. When the search is terminated, the history entry containing the search string becomes the current line.

To find other matching entries in the history list, type C-r or C-s as appropriate. This will search backward or forward in the history for the next entry matching the search string typed so far. Any other key sequence bound to a Readline command will terminate the

search and execute that command. For instance, a RET will terminate the search and accept the line, thereby executing the command from the history list. A movement command will terminate the search, make the last line found the current line, and begin editing.

Readline remembers the last incremental search string. If two *C-r*s are typed without any intervening characters defining a new search string, any remembered search string is used.

Non-incremental searches read the entire search string before starting to search for matching history lines. The search string may be typed by the user or be part of the contents of the current line.

31.3 Readline Init File

Although the Readline library comes with a set of Emacs-like keybindings installed by default, it is possible to use a different set of keybindings. Any user can customize programs that use Readline by putting commands in an *inputrc* file, conventionally in his home directory. The name of this file is taken from the value of the environment variable INPUTRC. If that variable is unset, the default is '~/.inputrc'.

When a program which uses the Readline library starts up, the init file is read, and the key bindings are set.

In addition, the C-x C-r command re-reads this init file, thus incorporating any changes that you might have made to it.

31.3.1 Readline Init File Syntax

There are only a few basic constructs allowed in the Readline init file. Blank lines are ignored. Lines beginning with a '#' are comments. Lines beginning with a '\$' indicate conditional constructs (see Section 31.3.2 [Conditional Init Constructs], page 373). Other lines denote variable settings and key bindings.

Variable Settings

You can modify the run-time behavior of Readline by altering the values of variables in Readline using the **set** command within the init file. The syntax is simple:

set variable value

Here, for example, is how to change from the default Emacs-like key binding to use vi line editing commands:

set editing-mode vi

Variable names and values, where appropriate, are recognized without regard to case. Unrecognized variable names are ignored.

Boolean variables (those that can be set to on or off) are set to on if the value is null or empty, on (case-insensitive), or 1. Any other value results in the variable being set to off.

A great deal of run-time behavior is changeable with the following variables.

bell-style

Controls what happens when Readline wants to ring the terminal bell. If set to 'none', Readline never rings the bell. If set to

'visible', Readline uses a visible bell if one is available. If set to 'audible' (the default), Readline attempts to ring the terminal's bell.

bind-tty-special-chars

If set to 'on', Readline attempts to bind the control characters treated specially by the kernel's terminal driver to their Readline equivalents.

comment-begin

The string to insert at the beginning of the line when the insert-comment command is executed. The default value is "#".

completion-ignore-case

If set to 'on', Readline performs filename matching and completion in a case-insensitive fashion. The default value is 'off'.

completion-query-items

The number of possible completions that determines when the user is asked whether the list of possibilities should be displayed. If the number of possible completions is greater than this value, Readline will ask the user whether or not he wishes to view them; otherwise, they are simply listed. This variable must be set to an integer value greater than or equal to 0. A negative value means Readline should never ask. The default limit is 100.

convert-meta

If set to 'on', Readline will convert characters with the eighth bit set to an ASCII key sequence by stripping the eighth bit and prefixing an ESC character, converting them to a meta-prefixed key sequence. The default value is 'on'.

disable-completion

If set to 'On', Readline will inhibit word completion. Completion characters will be inserted into the line as if they had been mapped to self-insert. The default is 'off'.

editing-mode

The editing-mode variable controls which default set of key bindings is used. By default, Readline starts up in Emacs editing mode, where the keystrokes are most similar to Emacs. This variable can be set to either 'emacs' or 'vi'.

enable-keypad

When set to 'on', Readline will try to enable the application keypad when it is called. Some systems need this to enable the arrow keys. The default is 'off'.

expand-tilde

If set to 'on', tilde expansion is performed when Readline attempts word completion. The default is 'off'.

history-preserve-point

If set to 'on', the history code attempts to place point at the same location on each history line retrieved with previous-history or next-history. The default is 'off'.

horizontal-scroll-mode

This variable can be set to either 'on' or 'off'. Setting it to 'on' means that the text of the lines being edited will scroll horizontally on a single screen line when they are longer than the width of the screen, instead of wrapping onto a new screen line. By default, this variable is set to 'off'.

input-meta

If set to 'on', Readline will enable eight-bit input (it will not clear the eighth bit in the characters it reads), regardless of what the terminal claims it can support. The default value is 'off'. The name meta-flag is a synonym for this variable.

isearch-terminators

The string of characters that should terminate an incremental search without subsequently executing the character as a command (see Section 31.2.5 [Searching], page 367). If this variable has not been given a value, the characters ESC and C-J will terminate an incremental search.

Sets Readline's idea of the current keymap for key binding commands. Acceptable keymap names are emacs, emacs-standard, emacs-meta, emacs-ctlx, vi, vi-move, vi-command, and vi-insert. vi is equivalent to vi-command; emacs is equivalent to emacs-standard. The default value is emacs. The value of the editing-mode variable also affects the default keymap.

mark-directories

If set to 'on', completed directory names have a slash appended. The default is 'on'.

mark-modified-lines

This variable, when set to 'on', causes Readline to display an asterisk ('*') at the start of history lines which have been modified. This variable is 'off' by default.

mark-symlinked-directories

If set to 'on', completed names which are symbolic links to directories have a slash appended (subject to the value of mark-directories). The default is 'off'.

match-hidden-files

This variable, when set to 'on', causes Readline to match files whose names begin with a '.' (hidden files) when performing filename completion, unless the leading '.' is supplied by the user in the filename to be completed. This variable is 'on' by default.

output-meta

If set to 'on', Readline will display characters with the eighth bit set directly rather than as a meta-prefixed escape sequence. The default is 'off'.

page-completions

If set to 'on', Readline uses an internal more-like pager to display a screenful of possible completions at a time. This variable is 'on' by default.

print-completions-horizontally

If set to 'on', Readline will display completions with matches sorted horizontally in alphabetical order, rather than down the screen. The default is 'off'.

show-all-if-ambiguous

This alters the default behavior of the completion functions. If set to 'on', words which have more than one possible completion cause the matches to be listed immediately instead of ringing the bell. The default value is 'off'.

show-all-if-unmodified

This alters the default behavior of the completion functions in a fashion similar to *show-all-if-ambiguous*. If set to 'on', words which have more than one possible completion without any possible partial completion (the possible completions don't share a common prefix) cause the matches to be listed immediately instead of ringing the bell. The default value is 'off'.

visible-stats

If set to 'on', a character denoting a file's type is appended to the filename when listing possible completions. The default is 'off'.

Key Bindings

The syntax for controlling key bindings in the init file is simple. First you need to find the name of the command that you want to change. The following sections contain tables of the command name, the default keybinding, if any, and a short description of what the command does.

Once you know the name of the command, simply place on a line in the init file the name of the key you wish to bind the command to, a colon, and then the name of the command. The name of the key can be expressed in different ways, depending on what you find most comfortable.

In addition to command names, readline allows keys to be bound to a string that is inserted when the key is pressed (a macro).

keyname: function-name or macro

keyname is the name of a key spelled out in English. For example:

Control-u: universal-argument Meta-Rubout: backward-kill-word

Control-o: "> output"

In the above example, C-u is bound to the function universal-argument, M-DEL is bound to the function backward-kill-word, and C-o is bound to run the macro expressed on the right hand side (that is, to insert the text '> output' into the line).

A number of symbolic character names are recognized while processing this key binding syntax: *DEL*, *ESC*, *ESCAPE*, *LFD*, *NEW-LINE*, *RET*, *RETURN*, *RUBOUT*, *SPACE*, *SPC*, and *TAB*.

"keyseq": function-name or macro

keyseq differs from keyname above in that strings denoting an entire key sequence can be specified, by placing the key sequence in double quotes. Some GNU Emacs style key escapes can be used, as in the following example, but the special character names are not recognized.

```
"\C-u": universal-argument
"\C-x\C-r": re-read-init-file
"\e[11~": "Function Key 1"
```

In the above example, C-u is again bound to the function universal-argument (just as it was in the first example), 'C-x C-r' is bound to the function re-read-init-file, and 'ESC [1 1 \sim ' is bound to insert the text 'Function Key 1'.

The following GNU Emacs style escape sequences are available when specifying key sequences:

\C- control prefix
\M- meta prefix
\e an escape character
\\ backslash
\" ", a double quotation mark
\', a single quote or apostrophe

In addition to the GNU Emacs style escape sequences, a second set of backslash escapes is available:

alert (bell) \a ****b backspace \d delete form feed \f newline \n carriage return \r \t horizontal tab vertical tab \v

\nnn the eight-bit character whose value is the octal value nnn (one to three digits)

\x\H\H\ the eight-bit character whose value is the hexadecimal value \(H\H\)
(one or two hex digits)

When entering the text of a macro, single or double quotes must be used to indicate a macro definition. Unquoted text is assumed to be a function name. In the macro body, the backslash escapes described above are expanded. Backslash will quote any other character in the macro text, including '"' and '''. For example, the following binding will make C-x insert a single '\' into the line:

31.3.2 Conditional Init Constructs

Readline implements a facility similar in spirit to the conditional compilation features of the C preprocessor which allows key bindings and variable settings to be performed as the result of tests. There are four parser directives used.

\$if The \$if construct allows bindings to be made based on the editing mode, the terminal being used, or the application using Readline. The text of the test extends to the end of the line; no characters are required to isolate it.

The mode= form of the \$if directive is used to test whether Readline is in emacs or vi mode. This may be used in conjunction with the 'set keymap' command, for instance, to set bindings in the emacs-standard and emacs-ctlx keymaps only if Readline is starting out in emacs mode.

The term= form may be used to include terminal-specific key bindings, perhaps to bind the key sequences output by the terminal's function keys. The word on the right side of the '=' is tested against both the full name of the terminal and the portion of the terminal name before the first '-'. This allows sun to match both sun and sun-cmd, for instance.

application

The application construct is used to include application-specific settings. Each program using the Readline library sets the application name, and you can test for a particular value. This could be used to bind key sequences to functions useful for a specific program. For instance, the following command adds a key sequence that quotes the current or previous word in Bash:

```
$if Bash
# Quote the current or previous word
"\C-xq": "\eb\"\ef\""
$endif
```

\$endif This command, as seen in the previous example, terminates an \$if command.

\$else Commands in this branch of the \$if directive are executed if the test fails.

\$include This directive takes a single filename as an argument and reads commands and bindings from that file. For example, the following directive reads from '/etc/inputrc':

\$include /etc/inputrc

31.3.3 Sample Init File

Here is an example of an inputrc file. This illustrates key binding, variable assignment, and conditional syntax.

```
# This file controls the behaviour of line input editing for
# programs that use the GNU Readline library. Existing
# programs include FTP, Bash, and GDB.
# You can re-read the inputrc file with C-x C-r.
# Lines beginning with '#' are comments.
# First, include any systemwide bindings and variable
# assignments from /etc/Inputrc
$include /etc/Inputrc
# Set various bindings for emacs mode.
set editing-mode emacs
$if mode=emacs
Meta-Control-h: backward-kill-word Text after the function name is ignored
# Arrow keys in keypad mode
#"\M-OD":
               backward-char
#"\M-OC":
               forward-char
# Arrow keys in ANSI mode
"\M-[D":
              backward-char
"\M-[C":
             forward-char
"\M-[A":
              previous-history
"\M-[B": next-history
# Arrow keys in 8 bit keypad mode
#"\M-\C-OD":
                 backward-char
#"\M-\C-OC":
                 forward-char
#"\M-\C-OA":
                previous-history
#"\M-\C-OB":
                 next-history
# Arrow keys in 8 bit ANSI mode
#"\M-\C-[D":
                 backward-char
#"\M-\C-[C":
                 forward-char
```

```
#"\M-\C-[A":
                   previous-history
#"\M-\C-[B":
                   next-history
C-q: quoted-insert
$endif
# An old-style binding. This happens to be the default.
TAB: complete
# Macros that are convenient for shell interaction
$if Bash
# edit the path
\label{eq:c-e-c-a-ef-c-f} $$ \C-xp": "PATH=${PATH}\e\C-e\C-a\ef\C-f"} $$
# prepare to type a quoted word --
# insert open and close double quotes
# and move to just after the open quote
"\C-x\"": "\"\"\C-b"
# insert a backslash (testing backslash escapes
# in sequences and macros)
"\C-x\\": "\\"
# Quote the current or previous word
"\C-xq": "\eb\"\ef\""
# Add a binding to refresh the line, which is unbound
"\C-xr": redraw-current-line
# Edit variable on current line.
"\M-\C-v": "\C-a\C-k\C-y\M-\C-e\C-a\C-y="
$endif
# use a visible bell if one is available
set bell-style visible
# don't strip characters to 7 bits when reading
set input-meta on
# allow iso-latin1 characters to be inserted rather
# than converted to prefix-meta sequences
set convert-meta off
# display characters with the eighth bit set directly
# rather than as meta-prefixed characters
set output-meta on
# if there are more than 150 possible completions for
# a word, ask the user if he wants to see all of them
set completion-query-items 150
```

```
# For FTP
$if Ftp
"\C-xg": "get \M-?"
"\C-xt": "put \M-?"
"\M-.": yank-last-arg
$endif
```

31.4 Bindable Readline Commands

This section describes Readline commands that may be bound to key sequences. Command names without an accompanying key sequence are unbound by default.

In the following descriptions, *point* refers to the current cursor position, and *mark* refers to a cursor position saved by the **set-mark** command. The text between the point and mark is referred to as the *region*.

31.4.1 Commands For Moving

```
beginning-of-line (C-a)
```

Move to the start of the current line.

end-of-line (C-e)

Move to the end of the line.

forward-char (C-f)

Move forward a character.

backward-char (C-b)

Move back a character.

forward-word (M-f)

Move forward to the end of the next word. Words are composed of letters and digits.

backward-word (M-b)

Move back to the start of the current or previous word. Words are composed of letters and digits.

clear-screen (C-1)

Clear the screen and redraw the current line, leaving the current line at the top of the screen.

redraw-current-line ()

Refresh the current line. By default, this is unbound.

31.4.2 Commands For Manipulating The History

accept-line (Newline or Return)

Accept the line regardless of where the cursor is. If this line is non-empty, it may be added to the history list for future recall with add_history(). If this line is a modified history line, the history line is restored to its original state.

previous-history (C-p)

Move 'back' through the history list, fetching the previous command.

next-history (C-n)

Move 'forward' through the history list, fetching the next command.

beginning-of-history (M-<)

Move to the first line in the history.

end-of-history (M->)

Move to the end of the input history, i.e., the line currently being entered.

reverse-search-history (C-r)

Search backward starting at the current line and moving 'up' through the history as necessary. This is an incremental search.

forward-search-history (C-s)

Search forward starting at the current line and moving 'down' through the the history as necessary. This is an incremental search.

non-incremental-reverse-search-history (M-p)

Search backward starting at the current line and moving 'up' through the history as necessary using a non-incremental search for a string supplied by the user.

non-incremental-forward-search-history (M-n)

Search forward starting at the current line and moving 'down' through the the history as necessary using a non-incremental search for a string supplied by the user.

history-search-forward ()

Search forward through the history for the string of characters between the start of the current line and the point. This is a non-incremental search. By default, this command is unbound.

history-search-backward ()

Search backward through the history for the string of characters between the start of the current line and the point. This is a non-incremental search. By default, this command is unbound.

yank-nth-arg (M-C-y)

Insert the first argument to the previous command (usually the second word on the previous line) at point. With an argument n, insert the nth word from the previous command (the words in the previous command begin with word 0). A negative argument inserts the nth word from the end of the previous command. Once the argument n is computed, the argument is extracted as if the '!n' history expansion had been specified.

yank-last-arg (M-. or M-_)

Insert last argument to the previous command (the last word of the previous history entry). With an argument, behave exactly like yank-nth-arg. Successive calls to yank-last-arg move back through the history list, inserting the last argument of each line in turn. The history expansion facilities are used to extract the last argument, as if the '!\$' history expansion had been specified.

31.4.3 Commands For Changing Text

delete-char (C-d)

Delete the character at point. If point is at the beginning of the line, there are no characters in the line, and the last character typed was not bound to delete-char, then return EOF.

backward-delete-char (Rubout)

Delete the character behind the cursor. A numeric argument means to kill the characters instead of deleting them.

forward-backward-delete-char ()

Delete the character under the cursor, unless the cursor is at the end of the line, in which case the character behind the cursor is deleted. By default, this is not bound to a key.

quoted-insert (C-q or C-v)

Add the next character typed to the line verbatim. This is how to insert key sequences like C-q, for example.

tab-insert (M-TAB)

Insert a tab character.

self-insert (a, b, A, 1, !, ...)

Insert yourself.

transpose-chars (C-t)

Drag the character before the cursor forward over the character at the cursor, moving the cursor forward as well. If the insertion point is at the end of the line, then this transposes the last two characters of the line. Negative arguments have no effect.

transpose-words (M-t)

Drag the word before point past the word after point, moving point past that word as well. If the insertion point is at the end of the line, this transposes the last two words on the line.

upcase-word (M-u)

Uppercase the current (or following) word. With a negative argument, uppercase the previous word, but do not move the cursor.

downcase-word (M-1)

Lowercase the current (or following) word. With a negative argument, lowercase the previous word, but do not move the cursor.

capitalize-word (M-c)

Capitalize the current (or following) word. With a negative argument, capitalize the previous word, but do not move the cursor.

overwrite-mode ()

Toggle overwrite mode. With an explicit positive numeric argument, switches to overwrite mode. With an explicit non-positive numeric argument, switches to insert mode. This command affects only emacs mode; vi mode does overwrite differently. Each call to readline() starts in insert mode.

In overwrite mode, characters bound to self-insert replace the text at point rather than pushing the text to the right. Characters bound to backward-delete-char replace the character before point with a space.

By default, this command is unbound.

31.4.4 Killing And Yanking

kill-line (C-k)

Kill the text from point to the end of the line.

backward-kill-line (C-x Rubout)

Kill backward to the beginning of the line.

unix-line-discard (C-u)

Kill backward from the cursor to the beginning of the current line.

kill-whole-line ()

Kill all characters on the current line, no matter where point is. By default, this is unbound.

kill-word (M-d)

Kill from point to the end of the current word, or if between words, to the end of the next word. Word boundaries are the same as forward-word.

backward-kill-word (M-DEL)

Kill the word behind point. Word boundaries are the same as backward-word.

unix-word-rubout (C-w)

Kill the word behind point, using white space as a word boundary. The killed text is saved on the kill-ring.

unix-filename-rubout ()

Kill the word behind point, using white space and the slash character as the word boundaries. The killed text is saved on the kill-ring.

delete-horizontal-space ()

Delete all spaces and tabs around point. By default, this is unbound.

kill-region ()

Kill the text in the current region. By default, this command is unbound.

copy-region-as-kill ()

Copy the text in the region to the kill buffer, so it can be yanked right away. By default, this command is unbound.

copy-backward-word ()

Copy the word before point to the kill buffer. The word boundaries are the same as backward-word. By default, this command is unbound.

copy-forward-word ()

Copy the word following point to the kill buffer. The word boundaries are the same as forward-word. By default, this command is unbound.

yank (C-y)

Yank the top of the kill ring into the buffer at point.

yank-pop (M-y)

Rotate the kill-ring, and yank the new top. You can only do this if the prior command is yank or yank-pop.

31.4.5 Specifying Numeric Arguments

digit-argument (*M*-0, *M*-1, ... *M*--)

Add this digit to the argument already accumulating, or start a new argument. M-- starts a negative argument.

universal-argument ()

This is another way to specify an argument. If this command is followed by one or more digits, optionally with a leading minus sign, those digits define the argument. If the command is followed by digits, executing universal-argument again ends the numeric argument, but is otherwise ignored. As a special case, if this command is immediately followed by a character that is neither a digit or minus sign, the argument count for the next command is multiplied by four. The argument count is initially one, so executing this function the first time makes the argument count four, a second time makes the argument count sixteen, and so on. By default, this is not bound to a key.

31.4.6 Letting Readline Type For You

complete (TAB)

Attempt to perform completion on the text before point. The actual completion performed is application-specific. The default is filename completion.

possible-completions (M-?)

List the possible completions of the text before point.

insert-completions (M-*)

Insert all completions of the text before point that would have been generated by possible-completions.

menu-complete ()

Similar to complete, but replaces the word to be completed with a single match from the list of possible completions. Repeated execution of menu-complete steps through the list of possible completions, inserting each match in turn. At the end of the list of completions, the bell is rung (subject to the setting of bell-style) and the original text is restored. An argument of n moves n positions forward in the list of matches; a negative argument may be used to move backward through the list. This command is intended to be bound to TAB, but is unbound by default.

delete-char-or-list ()

Deletes the character under the cursor if not at the beginning or end of the line (like delete-char). If at the end of the line, behaves identically to possible-completions. This command is unbound by default.

31.4.7 Keyboard Macros

start-kbd-macro (C-x ()

Begin saving the characters typed into the current keyboard macro.

end-kbd-macro (C-x))

Stop saving the characters typed into the current keyboard macro and save the definition.

call-last-kbd-macro (C-x e)

Re-execute the last keyboard macro defined, by making the characters in the macro appear as if typed at the keyboard.

31.4.8 Some Miscellaneous Commands

re-read-init-file (C-x C-r)

Read in the contents of the *inputrc* file, and incorporate any bindings or variable assignments found there.

abort (C-g)

Abort the current editing command and ring the terminal's bell (subject to the setting of bell-style).

do-uppercase-version (M-a, M-b, M-x, ...)

If the metafied character x is lowercase, run the command that is bound to the corresponding uppercase character.

prefix-meta (ESC)

Metafy the next character typed. This is for keyboards without a meta key. Typing 'ESC f' is equivalent to typing M-f.

undo (C-_ or C-x C-u)

Incremental undo, separately remembered for each line.

revert-line (M-r)

Undo all changes made to this line. This is like executing the undo command enough times to get back to the beginning.

tilde-expand (M-~)

Perform tilde expansion on the current word.

set-mark (C-@)

Set the mark to the point. If a numeric argument is supplied, the mark is set to that position.

exchange-point-and-mark (C-x C-x)

Swap the point with the mark. The current cursor position is set to the saved position, and the old cursor position is saved as the mark.

character-search (C-])

A character is read and point is moved to the next occurrence of that character. A negative count searches for previous occurrences.

character-search-backward (M-C-])

A character is read and point is moved to the previous occurrence of that character. A negative count searches for subsequent occurrences.

insert-comment (M-#)

Without a numeric argument, the value of the comment-begin variable is inserted at the beginning of the current line. If a numeric argument is supplied,

this command acts as a toggle: if the characters at the beginning of the line do not match the value of comment-begin, the value is inserted, otherwise the characters in comment-begin are deleted from the beginning of the line. In either case, the line is accepted as if a newline had been typed.

dump-functions ()

Print all of the functions and their key bindings to the Readline output stream. If a numeric argument is supplied, the output is formatted in such a way that it can be made part of an *inputrc* file. This command is unbound by default.

dump-variables ()

Print all of the settable variables and their values to the Readline output stream. If a numeric argument is supplied, the output is formatted in such a way that it can be made part of an *inputrc* file. This command is unbound by default.

dump-macros ()

Print all of the Readline key sequences bound to macros and the strings they output. If a numeric argument is supplied, the output is formatted in such a way that it can be made part of an *inputrc* file. This command is unbound by default.

emacs-editing-mode (C-e)

When in vi command mode, this causes a switch to emacs editing mode.

vi-editing-mode (M-C-j)

When in emacs editing mode, this causes a switch to vi editing mode.

31.5 Readline vi Mode

While the Readline library does not have a full set of vi editing functions, it does contain enough to allow simple editing of the line. The Readline vi mode behaves as specified in the POSIX 1003.2 standard.

In order to switch interactively between emacs and vi editing modes, use the command M-C-j (bound to emacs-editing-mode when in vi mode and to vi-editing-mode in emacs mode). The Readline default is emacs mode.

When you enter a line in vi mode, you are already placed in 'insertion' mode, as if you had typed an 'i'. Pressing ESC switches you into 'command' mode, where you can edit the text of the line with the standard vi movement keys, move to previous history lines with 'k' and subsequent lines with 'j', and so forth.

32 Using History Interactively

This chapter describes how to use the GNU History Library interactively, from a user's standpoint. It should be considered a user's guide. For information on using the GNU History Library in other programs, see the GNU Readline Library Manual.

32.1 History Expansion

The History library provides a history expansion feature that is similar to the history expansion provided by csh. This section describes the syntax used to manipulate the history information.

History expansions introduce words from the history list into the input stream, making it easy to repeat commands, insert the arguments to a previous command into the current input line, or fix errors in previous commands quickly.

History expansion takes place in two parts. The first is to determine which line from the history list should be used during substitution. The second is to select portions of that line for inclusion into the current one. The line selected from the history is called the event, and the portions of that line that are acted upon are called words. Various modifiers are available to manipulate the selected words. The line is broken into words in the same fashion that Bash does, so that several words surrounded by quotes are considered one word. History expansions are introduced by the appearance of the history expansion character, which is '!' by default.

32.1.1 Event Designators

An event designator is a reference to a command line entry in the history list.

! Start a history substitution, except when followed by a space, tab, the end of the line, or '='.

!n Refer to command line n.

!-n Refer to the command n lines back.

!! Refer to the previous command. This is a synonym for '!-1'.

!string Refer to the most recent command starting with string.

!?string[?]

Refer to the most recent command containing *string*. The trailing '?' may be omitted if the *string* is followed immediately by a newline.

^string1^string2^

Quick Substitution. Repeat the last command, replacing string1 with string2. Equivalent to !!:s/string1/string2/.

!# The entire command line typed so far.

32.1.2 Word Designators

Word designators are used to select desired words from the event. A ':' separates the event specification from the word designator. It may be omitted if the word designator begins with a '-', '\$', '*', '-', or '%'. Words are numbered from the beginning of the line, with the

first word being denoted by 0 (zero). Words are inserted into the current line separated by single spaces.

For example,

- !! designates the preceding command. When you type this, the preceding command is repeated in toto.
- !!:\$ designates the last argument of the preceding command. This may be shortened to !\$.
- !fi:2 designates the second argument of the most recent command starting with the letters fi.

Here are the word designators:

- O (zero) The Oth word. For many applications, this is the command word.
- n The nth word.
- ^ The first argument; that is, word 1.
- \$ The last argument.
- The word matched by the most recent '?string?' search.
- x-y A range of words; '-y' abbreviates '0-y'.
- * All of the words, except the 0th. This is a synonym for '1-\$'. It is not an error to use '*' if there is just one word in the event; the empty string is returned in that case.
- x* Abbreviates 'x-\$'
- x- Abbreviates 'x-\$' like 'x*', but omits the last word.

If a word designator is supplied without an event specification, the previous command is used as the event.

32.1.3 Modifiers

After the optional word designator, you can add a sequence of one or more of the following modifiers, each preceded by a ':'.

- h Remove a trailing pathname component, leaving only the head.
- t Remove all leading pathname components, leaving the tail.
- r Remove a trailing suffix of the form '.suffix', leaving the basename.
- e Remove all but the trailing suffix.
- Print the new command but do not execute it.

s/old/new/

Substitute new for the first occurrence of old in the event line. Any delimiter may be used in place of '/'. The delimiter may be quoted in old and new with a single backslash. If '&' appears in new, it is replaced by old. A single backslash will quote the '&'. The final delimiter is optional if it is the last character on the input line.

- & Repeat the previous substitution.
- g
 a Cause changes to be applied over the entire event line. Used in conjunction with 's', as in gs/old/new/, or with '&'.
- G Apply the following 's' modifier once to each word in the event.

Appendix A Formatting Documentation

The gdb 4 release includes an already-formatted reference card, ready for printing with PostScript or Ghostscript, in the 'gdb' subdirectory of the main source directory¹. If you can use PostScript or Ghostscript with your printer, you can print the reference card immediately with 'refcard.ps'.

The release also includes the source for the reference card. You can format it, using TeX, by typing:

```
make refcard.dvi
```

The GDB reference card is designed to print in *landscape* mode on US "letter" size paper; that is, on a sheet 11 inches wide by 8.5 inches high. You will need to specify this form of printing as an option to your DVI output program.

All the documentation for GDB comes as part of the machine-readable distribution. The documentation is written in Texinfo format, which is a documentation system that uses a single source file to produce both on-line information and a printed manual. You can use one of the Info formatting commands to create the on-line version of the documentation and TeX (or texi2roff) to typeset the printed version.

GDB includes an already formatted copy of the on-line Info version of this manual in the 'gdb' subdirectory. The main Info file is 'gdb-7.0.50.20100218-cvs/gdb/gdb.info', and it refers to subordinate files matching 'gdb.info*' in the same directory. If necessary, you can print out these files, or read them with any editor; but they are easier to read using the info subsystem in GNU Emacs or the standalone info program, available as part of the GNU Texinfo distribution.

If you want to format these Info files yourself, you need one of the Info formatting programs, such as texinfo-format-buffer or makeinfo.

If you have makeinfo installed, and are in the top level GDB source directory ('gdb-7.0.50.20100218-cvs', in the case of version 7.0.50.20100218-cvs), you can make the Info file by typing:

```
cd gdb
make gdb.info
```

If you want to typeset and print copies of this manual, you need TEX, a program to print its DVI output files, and 'texinfo.tex', the Texinfo definitions file.

TeX is a typesetting program; it does not print files directly, but produces output files called DVI files. To print a typeset document, you need a program to print DVI files. If your system has TeX installed, chances are it has such a program. The precise command to use depends on your system; lpr - d is common; another (for PostScript devices) is dvips. The DVI print command may require a file name without any extension or a '.dvi' extension.

T_EX also requires a macro definitions file called 'texinfo.tex'. This file tells T_EX how to typeset a document written in Texinfo format. On its own, T_EX cannot either read or typeset a Texinfo file. 'texinfo.tex' is distributed with GDB and is located in the 'gdb-version-number/texinfo' directory.

If you have TEX and a DVI printer program installed, you can typeset and print this manual. First switch to the 'gdb' subdirectory of the main source directory (for example, to 'gdb-7.0.50.20100218-cvs/gdb') and type:

 $^{^{1}}$ In 'gdb-7.0.50.20100218-cvs/gdb/refcard.ps' of the version 7.0.50.20100218-cvs release.

make gdb.dvi

Then give 'gdb.dvi' to your DVI printing program.

Appendix B Installing GDB

B.1 Requirements for Building gdb

Building GDB requires various tools and packages to be available. Other packages will be used only if they are found.

Tools/Packages Necessary for Building gdb

ISO C90 compiler

GDB is written in ISO C90. It should be buildable with any working C90 compiler, e.g. GCC.

Tools/Packages Optional for Building gdb

Expat

GDB can use the Expat XML parsing library. This library may be included with your operating system distribution; if it is not, you can get the latest version from http://expat.sourceforge.net. The 'configure' script will search for this library in several standard locations; if it is installed in an unusual path, you can use the '--with-libexpat-prefix' option to specify its location.

Expat is used for:

- Remote protocol memory maps (see Section D.16 [Memory Map Format], page 450)
- Target descriptions (see Appendix G [Target Descriptions], page 467)
- Remote shared library lists (see Section D.15 [Library List Format], page 449)
- MS-Windows shared libraries (see [Shared Libraries], page 191)

zlib

GDB will use the 'zlib' library, if available, to read compressed debug sections. Some linkers, such as GNU gold, are capable of producing binaries with compressed debug sections. If GDB is compiled with 'zlib', it will be able to read the debug information in such binaries.

The 'zlib' library is likely included with your operating system distribution; if it is not, you can get the latest version from http://zlib.net.

iconv

GDB's features related to character sets (see Section 10.18 [Character Sets], page 116) require a functioning iconv implementation. If you are on a GNU system, then this is provided by the GNU C Library. Some other systems also provide a working iconv.

On systems with iconv, you can install GNU Libiconv. If you have previously installed Libiconv, you can use the '--with-libiconv-prefix' option to configure.

GDB's top-level 'configure' and 'Makefile' will arrange to build Libiconv if a directory named 'libiconv' appears in the top-most source directory. If Libiconv is built this way, and if the operating system does not provide a suitable iconv implementation, then the just-built library will automatically be used by

GDB. One easy way to set this up is to download GNU Libiconv, unpack it, and then rename the directory holding the Libiconv source code to 'libiconv'.

B.2 Invoking the GDB 'configure' Script

GDB comes with a 'configure' script that automates the process of preparing GDB for installation; you can then use make to build the gdb program.¹

The GDB distribution includes all the source code you need for GDB in a single directory, whose name is usually composed by appending the version number to 'gdb'.

For example, the GDB version 7.0.50.20100218-cvs distribution is in the 'gdb-7.0.50.20100218-cvs' directory. That directory contains:

gdb-7.0.50.20100218-cvs/configure (and supporting files) script for configuring GDB and all its supporting libraries

gdb-7.0.50.20100218-cvs/gdb the source specific to GDB itself

gdb-7.0.50.20100218-cvs/bfd source for the Binary File Descriptor library

gdb-7.0.50.20100218-cvs/include GNU include files

gdb-7.0.50.20100218-cvs/libiberty source for the '-liberty' free software library

gdb-7.0.50.20100218-cvs/opcodes source for the library of opcode tables and disassemblers

gdb-7.0.50.20100218-cvs/readline source for the GNU command-line interface

gdb-7.0.50.20100218-cvs/glob source for the GNU filename pattern-matching subroutine

gdb-7.0.50.20100218-cvs/mmalloc source for the GNU memory-mapped malloc package

The simplest way to configure and build GDB is to run 'configure' from the 'gdb-version-number' source directory, which in this example is the 'gdb-7.0.50.20100218-cvs' directory.

First switch to the 'gdb-version-number' source directory if you are not already in it; then run 'configure'. Pass the identifier for the platform on which GDB will run as an argument.

For example:

cd gdb-7.0.50.20100218-cvs ./configure *host* make

¹ If you have a more recent version of gdb than 7.0.50.20100218-cvs, look at the 'README' file in the sources; we may have improved the installation procedures since publishing this manual.

where *host* is an identifier such as 'sun4' or 'decstation', that identifies the platform where GDB will run. (You can often leave off *host*; 'configure' tries to guess the correct value by examining your system.)

Running 'configure host' and then running make builds the 'bfd', 'readline', 'mmalloc', and 'libiberty' libraries, then gdb itself. The configured source files, and the binaries, are left in the corresponding source directories.

'configure' is a Bourne-shell (/bin/sh) script; if your system does not recognize this automatically when you run a different shell, you may need to run sh on it explicitly:

```
sh configure host
```

If you run 'configure' from a directory that contains source directories for multiple libraries or programs, such as the 'gdb-7.0.50.20100218-cvs' source directory for version 7.0.50.20100218-cvs, 'configure' creates configuration files for every directory level underneath (unless you tell it not to, with the '--norecursion' option).

You should run the 'configure' script from the top directory in the source tree, the 'gdb-version-number' directory. If you run 'configure' from one of the subdirectories, you will configure only that subdirectory. That is usually not what you want. In particular, if you run the first 'configure' from the 'gdb' subdirectory of the 'gdb-version-number' directory, you will omit the configuration of 'bfd', 'readline', and other sibling directories of the 'gdb' subdirectory. This leads to build errors about missing include files such as 'bfd/bfd.h'.

You can install gdb anywhere; it has no hardwired paths. However, you should make sure that the shell on your path (named by the 'SHELL' environment variable) is publicly readable. Remember that GDB uses the shell to start your program—some systems refuse to let GDB debug child processes whose programs are not readable.

B.3 Compiling GDB in Another Directory

If you want to run GDB versions for several host or target machines, you need a different gdb compiled for each combination of host and target. 'configure' is designed to make this easy by allowing you to generate each configuration in a separate subdirectory, rather than in the source directory. If your make program handles the 'VPATH' feature (GNU make does), running make in each of these directories builds the gdb program specified there.

To build gdb in a separate directory, run 'configure' with the '--srcdir' option to specify where to find the source. (You also need to specify a path to find 'configure' itself from your working directory. If the path to 'configure' would be the same as the argument to '--srcdir', you can leave out the '--srcdir' option; it is assumed.)

For example, with version 7.0.50.20100218-cvs, you can build gdb in a separate directory for a Sun 4 like this:

```
cd gdb-7.0.50.20100218-cvs
mkdir ../gdb-sun4
cd ../gdb-sun4
../gdb-7.0.50.20100218-cvs/configure sun4
make
```

When 'configure' builds a configuration using a remote source directory, it creates a tree for the binaries with the same structure (and using the same names) as the tree under the source directory. In the example, you'd find the Sun 4 library 'libiberty.a' in the directory 'gdb-sun4/libiberty', and GDB itself in 'gdb-sun4/gdb'.

Make sure that your path to the 'configure' script has just one instance of 'gdb' in it. If your path to 'configure' looks like '../gdb-7.0.50.20100218-cvs/gdb/configure', you are configuring only one subdirectory of GDB, not the whole package. This leads to build errors about missing include files such as 'bfd/bfd.h'.

One popular reason to build several GDB configurations in separate directories is to configure GDB for cross-compiling (where GDB runs on one machine—the *host*—while debugging programs that run on another machine—the *target*). You specify a cross-debugging target by giving the '--target=target' option to 'configure'.

When you run make to build a program or library, you must run it in a configured directory—whatever directory you were in when you called 'configure' (or one of its subdirectories).

The Makefile that 'configure' generates in each source directory also runs recursively. If you type make in a source directory such as 'gdb-7.0.50.20100218-cvs' (or in a separate configured directory configured with '--srcdir=dirname/gdb-7.0.50.20100218-cvs'), you will build all the required libraries, and then build GDB.

When you have multiple hosts or targets configured in separate directories, you can run make on them in parallel (for example, if they are NFS-mounted on each of the hosts); they will not interfere with each other.

B.4 Specifying Names for Hosts and Targets

The specifications used for hosts and targets in the 'configure' script are based on a three-part naming scheme, but some short predefined aliases are also supported. The full naming scheme encodes three pieces of information in the following pattern:

```
architecture-vendor-os
```

For example, you can use the alias sun4 as a host argument, or as the value for target in a --target=target option. The equivalent full name is 'sparc-sun-sunos4'.

The 'configure' script accompanying GDB does not provide any query facility to list all supported host and target names or aliases. 'configure' calls the Bourne shell script config.sub to map abbreviations to full names; you can read the script, if you wish, or you can use it to test your guesses on abbreviations—for example:

```
% sh config.sub i386-linux
i386-pc-linux-gnu
% sh config.sub alpha-linux
alpha-unknown-linux-gnu
% sh config.sub hp9k700
hppa1.1-hp-hpux
% sh config.sub sun4
sparc-sun-sunos4.1.1
% sh config.sub sun3
m68k-sun-sunos4.1.1
% sh config.sub i986v
Invalid configuration 'i986v': machine 'i986v' not recognized
```

config.sub is also distributed in the GDB source directory ('gdb-7.0.50.20100218-cvs', for version 7.0.50.20100218-cvs).

B.5 'configure' Options

Here is a summary of the 'configure' options and arguments that are most often useful for building GDB. 'configure' also has several other options not listed here. See Info file 'configure.info', node 'What Configure Does', for a full explanation of 'configure'.

```
configure [--help]
        [--prefix=dir]
        [--exec-prefix=dir]
        [--srcdir=dirname]
        [--norecursion] [--rm]
        [--target=target]
        host
```

You may introduce options with a single '-' rather than '--' if you prefer; but you may abbreviate option names if you use '--'.

--help Display a quick summary of how to invoke 'configure'.

--prefix=dir

Configure the source to install programs and files under directory 'dir'.

--exec-prefix=dir

Configure the source to install programs under directory 'dir'.

--srcdir=dirname

Warning: using this option requires GNU make, or another make that implements the VPATH feature.

Use this option to make configurations in directories separate from the GDB source directories. Among other things, you can use this to build (or maintain) several configurations simultaneously, in separate directories. 'configure' writes configuration-specific files in the current directory, but arranges for them to use the source in the directory dirname. 'configure' creates directories under the working directory in parallel to the source directories below dirname.

--norecursion

Configure only the directory level where 'configure' is executed; do not propagate configuration to subdirectories.

--target=target

Configure GDB for cross-debugging programs running on the specified *target*. Without this option, GDB is configured to debug programs that run on the same machine (*host*) as GDB itself.

There is no convenient way to generate a list of all available targets.

host ... Configure GDB to run on the specified host.

There is no convenient way to generate a list of all available hosts.

There are many other options available as well, but they are generally needed for special purposes only.

B.6 System-wide configuration and settings

GDB can be configured to have a system-wide init file; this file will be read and executed at startup (see Section 2.1.3 [What GDB does during startup], page 15).

Here is the corresponding configure option:

--with-system-gdbinit=file

Specify that the default location of the system-wide init file is file.

If gdb has been configured with the option '--prefix=\$prefix', it may be subject to relocation. Two possible cases:

- If the default location of this init file contains '\$prefix', it will be subject to relocation. Suppose that the configure options are '--prefix=\$prefix --with-system-gdbinit=\$prefix/etc/gdbinit'; if GDB is moved from '\$prefix' to '\$install', the system init file is looked for as '\$install/etc/gdbinit' instead of '\$prefix/etc/gdbinit'.
- By contrast, if the default location does not contain the prefix, it will not be relocated. E.g. if GDB has been configured with '--prefix=/usr/local --with-system-gdbinit=/usr/share/gdb/gdbinit', then GDB will always look for '/usr/share/gdb/gdbinit', wherever GDB is installed.

Appendix C Maintenance Commands

In addition to commands intended for GDB users, GDB includes a number of commands intended for GDB developers, that are not documented elsewhere in this manual. These commands are provided here for reference. (For commands that turn on debugging messages, see Section 22.8 [Debugging Output], page 251.)

maint agent expression maint agent-eval expression

Translate the given expression into remote agent bytecodes. This command is useful for debugging the Agent Expression mechanism (see Appendix E [Agent Expressions], page 453). The 'agent' version produces an expression useful for data collection, such as by tracepoints, while 'maint agent-eval' produces an expression that evaluates directly to a result. For instance, a collection expression for globa + globb will include bytecodes to record four bytes of memory at each of the addresses of globa and globb, while discarding the result of the addition, while an evaluation expression will do the addition and return the sum.

maint info breakpoints

Using the same format as 'info breakpoints', display both the breakpoints you've set explicitly, and those GDB is using for internal purposes. Internal breakpoints are shown with negative breakpoint numbers. The type column identifies what kind of breakpoint is shown:

breakpoint

Normal, explicitly set breakpoint.

watchpoint

Normal, explicitly set watchpoint.

longjmp Internal breakpoint, used to handle correctly stepping through longjmp calls.

longjmp resume

Internal breakpoint at the target of a long imp.

until Temporary internal breakpoint used by the GDB until command.

finish Temporary internal breakpoint used by the gdb finish command.

shlib events

Shared library events.

set displaced-stepping show displaced-stepping

Control whether or not GDB will do displaced stepping if the target supports it. Displaced stepping is a way to single-step over breakpoints without removing them from the inferior, by executing an out-of-line copy of the instruction that was originally at the breakpoint location. It is also known as out-of-line single-stepping.

set displaced-stepping on

If the target architecture supports it, GDB will use displaced stepping to step over breakpoints.

set displaced-stepping off

GDB will not use displaced stepping to step over breakpoints, even if such is supported by the target architecture.

set displaced-stepping auto

This is the default mode. GDB will use displaced stepping only if non-stop mode is active (see Section 5.4.2 [Non-Stop Mode], page 67) and the target architecture supports displaced stepping.

maint check-symtabs

Check the consistency of psymtabs and symtabs.

maint cplus first_component name

Print the first C++ class/namespace component of name.

maint cplus namespace

Print the list of possible C++ namespaces.

maint demangle name

Demangle a C++ or Objective-C mangled name.

maint deprecate command [replacement] maint undeprecate command

Deprecate or undeprecate the named *command*. Deprecated commands cause GDB to issue a warning when you use them. The optional argument *replacement* says which newer command should be used in favor of the deprecated one; if it is given, GDB will mention the replacement as part of the warning.

maint dump-me

Cause a fatal signal in the debugger and force it to dump its core. This is supported only on systems which support aborting a program with the SIGQUIT signal.

```
maint internal-error [message-text]
maint internal-warning [message-text]
```

Cause GDB to call the internal function internal_error or internal_warning and hence behave as though an internal error or internal warning has been detected. In addition to reporting the internal problem, these functions give the user the opportunity to either quit GDB or create a core file of the current GDB session.

These commands take an optional parameter message-text that is used as the text of the error or warning message.

Here's an example of using internal-error:

```
(gdb) maint internal-error testing, 1, 2 .../maint.c:121: internal-error: testing, 1, 2 A problem internal to GDB has been detected. Further debugging may prove unreliable. Quit this debugging session? (y or n) n
```

```
Create a core file? (y or n) n (gdb)
```

maint set internal-error action [ask|yes|no]
maint show internal-error action
maint set internal-warning action [ask|yes|no]
maint show internal-warning action

When GDB reports an internal problem (error or warning) it gives the user the opportunity to both quit GDB and create a core file of the current GDB session. These commands let you override the default behaviour for each particular action, described in the table below.

'quit' You can specify that GDB should always (yes) or never (no) quit.

The default is to ask the user what to do.

'corefile'

You can specify that GDB should always (yes) or never (no) create a core file. The default is to ask the user what to do.

maint packet text

If GDB is talking to an inferior via the serial protocol, then this command sends the string text to the inferior, and displays the response packet. GDB supplies the initial '\$' character, the terminating '#' character, and the checksum.

maint print architecture [file]

Print the entire architecture configuration. The optional argument file names the file where the output goes.

maint print c-tdesc

Print the current target description (see Appendix G [Target Descriptions], page 467) as a C source file. The created source file can be used in GDB when an XML parser is not available to parse the description.

maint print dummy-frames

Prints the contents of gdb's internal dummy-frame stack.

```
(gdb) b add
...
(gdb) print add(2,3)
Breakpoint 2, add (a=2, b=3) at ...
58  return (a + b);
The program being debugged stopped while in a function called from GDB.
...
(gdb) maint print dummy-frames
0x1a57c80: pc=0x01014068 fp=0x0200bddc sp=0x0200bdd6
top=0x0200bdd4 id={stack=0x200bddc,code=0x101405c}
call_lo=0x01014000 call_hi=0x01014001
(gdb)
```

Takes an optional file parameter.

```
maint print registers [file]
maint print raw-registers [file]
maint print cooked-registers [file]
maint print register-groups [file]
```

Print GDB's internal register data structures.

The command maint print raw-registers includes the contents of the raw register cache; the command maint print cooked-registers includes the (cooked) value of all registers; and the command maint print register-groups includes the groups that each register is a member of. See Section "Registers" in GDB Internals.

These commands take an optional parameter, a file name to which to write the information.

maint print reggroups [file]

Print GDB's internal register group data structures. The optional argument file tells to what file to write the information.

The register groups info looks like this:

```
(gdb) maint print reggroups
Group
            Type
general
            user
float
            user
all
            user
vector
            user
system
            user
save
            internal
            internal
restore
```

flushregs

This command forces gdb to flush its internal register cache.

maint print objfiles

Print a dump of all known object files. For each object file, this command prints its name, address in memory, and all of its psymtabs and symtabs.

maint print statistics

This command prints, for each object file in the program, various data about that object file followed by the byte cache (bcache) statistics for the object file. The objfile data includes the number of minimal, partial, full, and stabs symbols, the number of types defined by the objfile, the number of as yet unexpanded psym tables, the number of line tables and string tables, and the amount of memory used by the various tables. The bcache statistics include the counts, sizes, and counts of duplicates of all and unique objects, max, average, and median entry size, total memory used and its overhead and savings, and various measures of the hash table size and chain lengths.

maint print target-stack

A target is an interface between the debugger and a particular kind of file or process. Targets can be stacked in *strata*, so that more than one target can potentially respond to a request. In particular, memory accesses will walk down the stack of targets until they find a target that is interested in handling that particular address.

This command prints a short description of each layer that was pushed on the target stack, starting from the top layer down to the bottom one.

maint print type expr

Print the type chain for a type specified by expr. The argument can be either a type name or a symbol. If it is a symbol, the type of that symbol is described.

The type chain produced by this command is a recursive definition of the data type as stored in GDB's data structures, including its flags and contained types.

maint set dwarf2 max-cache-age maint show dwarf2 max-cache-age

Control the DWARF 2 compilation unit cache.

In object files with inter-compilation-unit references, such as those produced by the GCC option '-feliminate-dwarf2-dups', the DWARF 2 reader needs to frequently refer to previously read compilation units. This setting controls how long a compilation unit will remain in the cache if it is not referenced. A higher limit means that cached compilation units will be stored in memory longer, and more total memory will be used. Setting it to zero disables caching, which will slow down gdb startup, but reduce memory consumption.

maint set profile maint show profile

Control profiling of GDB.

Profiling will be disabled until you use the 'maint set profile' command to enable it. When you enable profiling, the system will begin collecting timing and execution count data; when you disable profiling or exit GDB, the results will be written to a log file. Remember that if you use profiling, GDB will overwrite the profiling log file (often called 'gmon.out'). If you have a record of important profiling data in a 'gmon.out' file, be sure to move it to a safe location.

Configuring with '--enable-profiling' arranges for GDB to be compiled with the '-pg' compiler option.

maint set show-debug-regs maint show show-debug-regs

Control whether to show variables that mirror the hardware debug registers. Use ON to enable, OFF to disable. If enabled, the debug registers values are shown when GDB inserts or removes a hardware breakpoint or watchpoint, and when the inferior triggers a hardware-assisted breakpoint or watchpoint.

maint space

Control whether to display memory usage for each command. If set to a nonzero value, GDB will display how much memory each command took, following the command's own output. This can also be requested by invoking GDB with the '--statistics' command-line switch (see Section 2.1.2 [Mode Options], page 13).

maint time

Control whether to display the execution time for each command. If set to a nonzero value, GDB will display how much time it took to execute each command, following the command's own output. The time is not printed for the commands that run the target, since there's no mechanism currently to compute how much time was spend by GDB and how much time was spend by the program been debugged. it's not possibly currently This can also be requested by invoking GDB with the '--statistics' command-line switch (see Section 2.1.2 [Mode Options], page 13).

maint translate-address [section] addr

Find the symbol stored at the location specified by the address addr and an optional section name section. If found, GDB prints the name of the closest symbol and an offset from the symbol's location to the specified address. This is similar to the info address command (see Chapter 16 [Symbols], page 175), except that this command also allows to find symbols in other sections.

If section was not specified, the section in which the symbol was found is also printed. For dynamically linked executables, the name of executable or shared library containing the symbol is printed as well.

The following command is useful for non-interactive invocations of GDB, such as in the test suite.

set watchdog nsec

Set the maximum number of seconds GDB will wait for the target operation to finish. If this time expires, GDB reports and error and the command is aborted.

show watchdog

Show the current setting of the target wait timeout.

Appendix D GDB Remote Serial Protocol

D.1 Overview

There may be occasions when you need to know something about the protocol—for example, if there is only one serial port to your target machine, you might want your program to do something special if it recognizes a packet meant for GDB.

In the examples below, '->' and '<-' are used to indicate transmitted and received data, respectively.

All GDB commands and responses (other than acknowledgments and notifications, see Section D.10 [Notification Packets], page 434) are sent as a packet. A packet is introduced with the character '\$', the actual packet-data, and the terminating character '#' followed by a two-digit checksum:

```
$packet-data#checksum
```

The two-digit *checksum* is computed as the modulo 256 sum of all characters between the leading '\$' and the trailing '#' (an eight bit unsigned checksum).

Implementors should note that prior to GDB 5.0 the protocol specification also included an optional two-digit sequence-id:

```
$sequence-id:packet-data#checksum
```

That sequence-id was appended to the acknowledgment. GDB has never output sequence-ids. Stubs that handle packets added since GDB 5.0 must not accept sequence-id.

When either the host or the target machine receives a packet, the first response expected is an acknowledgment: either '+' (to indicate the package was received correctly) or '-' (to request retransmission):

```
-> $packet-data#checksum
<- +</pre>
```

The '+'/'-' acknowledgments can be disabled once a connection is established. See Section D.12 [Packet Acknowledgment], page 436, for details.

The host (GDB) sends commands, and the target (the debugging stub incorporated in your program) sends a response. In the case of step and continue commands, the response is only sent when the operation has completed, and the target has again stopped all threads in all attached processes. This is the default all-stop mode behavior, but the remote protocol also supports GDB's non-stop execution mode; see Section D.11 [Remote Non-Stop], page 434, for details.

packet-data consists of a sequence of characters with the exception of '#' and '\$' (see 'X' packet for additional exceptions).

Fields within the packet should be separated using ',' ';' or ':'. Except where otherwise noted all numbers are represented in HEX with leading zeros suppressed.

Implementors should note that prior to GDB 5.0, the character ':' could not appear as the third character in a packet (as it would potentially conflict with the sequence-id).

Binary data in most packets is encoded either as two hexadecimal digits per byte of binary data. This allowed the traditional remote protocol to work over connections which were only seven-bit clean. Some packets designed more recently assume an eight-bit clean connection, and use a more efficient encoding to send and receive binary data.

The binary data representation uses 7d (ASCII '}') as an escape character. Any escaped byte is transmitted as the escape character followed by the original character XORed with 0x20. For example, the byte 0x7d would be transmitted as the two bytes 0x7d 0x5d. The bytes 0x23 (ASCII '#'), 0x24 (ASCII '\$'), and 0x7d (ASCII '}') must always be escaped. Responses sent by the stub must also escape 0x2a (ASCII '*'), so that it is not interpreted as the start of a run-length encoded sequence (described next).

Response data can be run-length encoded to save space. Run-length encoding replaces runs of identical characters with one instance of the repeated character, followed by a '*' and a repeat count. The repeat count is itself sent encoded, to avoid binary characters in data: a value of n is sent as n+29. For a repeat count greater or equal to 3, this produces a printable ASCII character, e.g. a space (ASCII code 32) for a repeat count of 3. (This is because run-length encoding starts to win for counts 3 or more.) Thus, for example, '0*' is a run-length encoding of "0000": the space character after '*' means repeat the leading 0 32 - 29 = 3 more times.

The printable characters '#' and '\$' or with a numeric value greater than 126 must not be used. Runs of six repeats ('#') or seven repeats ('\$') can be expanded using a repeat count of only five ('"'). For example, '00000000' can be encoded as '0*"00'.

In describing packets (and responses), each description has a template showing the overall syntax, followed by an explanation of the packet's meaning. We include spaces in some of the templates for clarity; these are not part of the packet's syntax. No GDB packet uses spaces to separate its components. For example, a template like 'foo bar baz' describes a packet beginning with the three ASCII bytes 'foo', followed by a bar, followed directly by a baz. GDB does not transmit a space character between the 'foo' and the bar, or between the bar and the baz.

We place optional portions of a packet in [square brackets]; for example, a template like 'c [addr]' describes a packet beginning with the single ASCII character 'c', possibly followed by an addr.

D.2 Standard Replies

The remote protocol specifies a few standard replies. All commands support these, except as noted in the individual command descriptions.

- For any command not supported by the stub, an empty response ('\$#00') should be returned. That way it is possible to extend the protocol. A newer gdb can tell if a packet is supported based on that response (but see also [qSupported], page 419).
- 'E xx' An error has occurred; xx is a two-digit hexadecimal error number. In almost all cases, the protocol does not specify the meaning of the error numbers; GDB usually ignores the numbers, or displays them to the user without further interpretation.

'E.name[.message]'

An error has occurred; name is the name of the error. The name may contain letters, numbers, and '-' characters. If present, message is an error message, encoded using the escaped eight-bit conventions for binary data described above.

Except as noted, named errors may only be returned to commands documented to expect them; this ensures that older stubs can interact with newer versions of gdb, even when interpretations for named errors have been added to the protocol.

The protocol uses the following error names:

'fatal' A fatal error has occurred; the stub will be unable to interact further with gdb. Fatal errors should always include a message explaining their cause.

Any command may return this error.

'memtype' The memory addressed is of the wrong type for the given command. For example, a 'vFlashWrite' command applied to non-flash memory elicits an 'E.memtype' error response.

D.3 Packets

The following table provides a complete list of all currently defined *commands* and their corresponding response *data*. See Section D.14 [File-I/O Remote Protocol Extension], page 437, for details about the File I/O extension of the remote protocol.

Each packet's description has a template showing the packet's overall syntax, followed by an explanation of the packet's meaning. We include spaces in some of the templates for clarity; these are not part of the packet's syntax. No gdb packet uses spaces to separate its components. For example, a template like 'foo bar baz' describes a packet beginning with the three ASCII bytes 'foo', followed by a bar, followed directly by a baz. gdb does not transmit a space character between the 'foo' and the bar, or between the bar and the baz.

Several packets and replies include a *thread-id* field to identify a thread. Normally these are positive numbers with a target-specific interpretation, formatted as big-endian hex strings. A *thread-id* can also be a literal '-1' to indicate all threads, or '0' to pick any thread.

In addition, the remote protocol supports a multiprocess feature in which the thread-id syntax is extended to optionally include both process and thread ID fields, as 'ppid.tid'. The pid (process) and tid (thread) components each have the format described above: a positive number with target-specific interpretation formatted as a big-endian hex string, literal '-1' to indicate all processes or threads (respectively), or '0' to indicate an arbitrary process or thread. Specifying just a process, as 'ppid', is equivalent to 'ppid.-1'. It is an error to specify all processes but a specific thread, such as 'p-1.tid'. Note that the 'p' prefix is not used for those packets and replies explicitly documented to include a process ID, rather than a thread-id.

The multiprocess thread-id syntax extensions are only used if both gdb and the stub report support for the 'multiprocess' feature using 'qSupported'. See [multiprocess extensions], page 423, for more information.

Note that all packet forms beginning with an upper- or lower-case letter, other than those described here, are reserved for future use.

A stub is required to support the 'g', 'G', 'm', 'M', 'c', and 's' commands. All other commands are optional.

Here are the packet descriptions.

'!' Enable extended mode. In extended mode, the remote server is made persistent. The 'R' packet is used to restart the program being debugged.

Reply:

'OK' The remote target both supports and has enabled extended mode.

'?' Indicate the reason the target halted. The reply is the same as for step and continue. This packet has a special interpretation when the target is in non-stop mode; see Section D.11 [Remote Non-Stop], page 434.

Reply: See Section D.4 [Stop Reply Packets], page 413, for the reply specifications.

'A arglen, argnum, arg, ...'

Initialized argv[] array passed into program. arglen specifies the number of bytes in the hex encoded byte stream arg. See gdbserver for more details.

Reply:

'OK' The arguments were set.

'b baud' (Don't use this packet; its behavior is not well-defined.) Change the serial line speed to baud.

JTC: When does the transport layer state change? When it's received, or after the ACK is transmitted. In either case, there are problems if the command or the acknowledgment packet is dropped.

Stan: If people really wanted to add something like this, and get it working for the first time, they ought to modify ser-unix.c to send some kind of out-of-band message to a specially-setup stub and have the switch happen "in between" packets, so that from remote protocol's point of view, nothing actually happened.

'B addr, mode'

Set (mode is 'S') or clear (mode is 'C') a breakpoint at addr.

Don't use this packet. Use the 'Z' and 'z' packets instead (see [insert breakpoint or watchpoint packet], page 412).

'bc' Backward continue. Execute the target system in reverse. No parameter. See Chapter 6 [Reverse Execution], page 71, for more information.

Reply: See Section D.4 [Stop Reply Packets], page 413, for the reply specifications.

'bs' Backward single step. Execute one instruction in reverse. No parameter. See Chapter 6 [Reverse Execution], page 71, for more information.

Reply: See Section D.4 [Stop Reply Packets], page 413, for the reply specifications.

'c [addr]' Continue. addr is address to resume. If addr is omitted, resume at current address.

Reply: See Section D.4 [Stop Reply Packets], page 413, for the reply specifications.

'C sig[;addr]'

Continue with signal sig (hex signal number). If '; addr' is omitted, resume at same address.

Reply: See Section D.4 [Stop Reply Packets], page 413, for the reply specifications

'd' Toggle debug flag.

Don't use this packet; instead, define a general set packet (see Section D.5 [General Query Packets], page 415).

۲D'

'D; pid' The first form of the packet is used to detach GDB from the remote system. It is sent to the remote target before GDB disconnects via the detach command. The second form, including a process ID, is used when multiprocess protocol extensions are enabled (see [multiprocess extensions], page 423), to detach only a specific process. The pid is specified as a big-endian hex string.

Reply:

'OK' for success

'F RC, EE, CF; XX'

A reply from GDB to an 'F' packet sent by the target. This is part of the File-I/O protocol extension. See Section D.14 [File-I/O Remote Protocol Extension], page 437, for the specification.

'g' Read general registers.

Reply:

'XX...' Each byte of register data is described by two hex digits. The bytes with the register are transmitted in target byte order. The size of each register and their position within the 'g' packet are determined by the GDB internal gdbarch functions DEPRECATED_REGISTER_RAW_SIZE and gdbarch_register_name. The specification of several standard 'g' packets is specified below.

'G XX...' Write general registers. See [read registers packet], page 407, for a description of the XX... data.

Reply:

'OK' for success

'H c thread-id'

Set thread for subsequent operations ('m', 'M', 'g', 'G', et.al.). c depends on the operation to be performed: it should be 'c' for step and continue operations, 'g' for other operations. The thread designator thread-id has the format and interpretation described in [thread-id syntax], page 405.

Reply:

'OK' for success

'i [addr[,nnn]]'

Step the remote target by a single clock cycle. If ',nnn' is present, cycle step nnn cycles. If addr is present, cycle step starting at that address.

'I' Signal, then cycle step. See [step with signal packet], page 409. See [cycle step packet], page 407.

'k' Kill request.

FIXME: There is no description of how to operate when a specific thread context has been selected (i.e. does 'k' kill only that thread?).

'm addr, length'

Read *length* bytes of memory starting at address *addr*. Note that *addr* may not be aligned to any particular boundary.

The stub need not use any particular size or alignment when gathering data from memory for the response; even if addr is word-aligned and length is a multiple of the word size, the stub is free to use byte accesses, or not. For this reason, this packet may not be suitable for accessing memory-mapped I/O devices.

Reply:

'XX...' Memory contents; each byte is transmitted as a two-digit hexadecimal number. The reply may contain fewer bytes than requested if the server was able to read only part of the region of memory.

'M addr, length: XX...'

Write *length* bytes of memory starting at address *addr. XX...* is the data; each byte is transmitted as a two-digit hexadecimal number.

Reply:

'OK' All the data was written successfully. (If only part of the data was written, this command returns an error.)

'p n' Read the value of register n; n is in hex. See [read registers packet], page 407, for a description of how the returned register value is encoded.

Reply:

'XX...' the register's value

'P n...=r...'

Write register n... with value r... The register number n is in hexadecimal, and r... contains two hex digits for each byte in the register (target byte order). Reply:

'OK' for success

'q name params...'

'Q name params...'

General query ('q') and set ('Q'). These packets are described fully in Section D.5 [General Query Packets], page 415.

'r' Reset the entire system.

Don't use this packet; use the 'R' packet instead.

'R XX' Restart the program being debugged. XX, while needed, is ignored. This packet is only available in extended mode (see [extended mode], page 406).

The 'R' packet has no reply.

's [addr]' Single step. addr is the address at which to resume. If addr is omitted, resume at same address.

Reply: See Section D.4 [Stop Reply Packets], page 413, for the reply specifications.

'S sig[;addr]'

Step with signal. This is analogous to the 'C' packet, but requests a single-step, rather than a normal resumption of execution.

Reply: See Section D.4 [Stop Reply Packets], page 413, for the reply specifications.

't addr:PP.MM'

Search backwards starting at address addr for a match with pattern PP and mask MM. PP and MM are 4 bytes. addr must be at least 3 digits.

'T thread-id'

Find out if the thread thread-id is alive. See [thread-id syntax], page 405. Reply:

'OK' thread is still alive

'E NN' thread is dead

'v' Packets starting with 'v' are identified by a multi-letter name, up to the first ';' or '?' (or the end of the packet).

'vAttach; pid'

Attach to a new process with the specified process ID *pid*. The process ID is a hexadecimal integer identifying the process. In all-stop mode, all threads in the attached process are stopped; in non-stop mode, it may be attached without being stopped if that is supported by the target.

This packet is only available in extended mode (see [extended mode], page 406). Reply:

'E nn' for an error

'Any stop packet'

for success in all-stop mode (see Section D.4 [Stop Reply Packets], page 413)

'OK' for success in non-stop mode (see Section D.11 [Remote Non-Stop], page 434)

'vCont[;action[:thread-id]]...'

Resume the inferior, specifying different actions for each thread. If an action is specified with no *thread-id*, then it is applied to any threads that don't have a specific action specified; if no default action is specified then other threads should remain stopped in all-stop mode and in their current state in non-stop mode. Specifying multiple default actions is an error; specifying no actions is also an error. Thread IDs are specified using the syntax described in [thread-id syntax], page 405.

Currently supported actions are:

'c' Continue.

'C sig' Continue with signal sig. The signal sig should be two hex digits.

's' Step.

'S sig' Step with signal sig. The signal sig should be two hex digits.

't' Stop.

The optional argument addr normally associated with the 'c', 'C', 's', and 'S' packets is not supported in 'vCont'.

The 't' action is only relevant in non-stop mode (see Section D.11 [Remote Non-Stop], page 434) and may be ignored by the stub otherwise. A stop reply should be generated for any affected thread not already stopped. When a thread is stopped by means of a 't' action, the corresponding stop reply should indicate that the thread has stopped with signal '0', regardless of whether the target uses some other signal as an implementation detail.

Reply: See Section D.4 [Stop Reply Packets], page 413, for the reply specifications.

'vCont?' Request a list of actions supported by the 'vCont' packet.

Reply:

'vCont[; action...]'

The 'vCont' packet is supported. Each action is a supported command in the 'vCont' packet.

' The 'vCont' packet is not supported.

'vFile: operation: parameter . . . '

Perform a file operation on the target system. For details, see Section D.8 [Host I/O Packets], page 432.

'vFlashErase: addr, length'

Direct the stub to erase *length* bytes of flash starting at *addr*. The region may enclose any number of flash blocks, but its start and end must fall on block boundaries, as indicated by the flash block size appearing in the memory map (see Section D.16 [Memory Map Format], page 450). GDB groups flash memory programming operations together, and sends a 'vFlashDone' request after each group; the stub is allowed to delay erase operation until the 'vFlashDone' packet is received.

The stub must support 'vCont' if it reports support for multiprocess extensions (see [multiprocess extensions], page 423). Note that in this case 'vCont' actions can be specified to apply to all threads in a process by using the 'ppid.-1' form of the thread-id.

Reply:

'OK' for success

'E.memtype'

for vFlashErase addressing non-flash memory

'vFlashWrite:addr:XX...'

Direct the stub to write data to flash address addr. The data is passed in binary form using the same encoding as for the 'X' packet (see [Binary Data], page 403). The memory ranges specified by 'vFlashWrite' packets preceding a 'vFlashDone' packet must not overlap, and must appear in order of increasing addresses (although 'vFlashErase' packets for higher addresses may already have been received; the ordering is guaranteed only between 'vFlashWrite' packets). If a packet writes to an address that was neither erased by a preceding 'vFlashErase' packet nor by some other target-specific method, the results are unpredictable.

Reply:

'OK' for success

'E.memtype'

for vFlashWrite addressing non-flash memory

'vFlashDone'

Indicate to the stub that flash programming operation is finished. The stub is permitted to delay or batch the effects of a group of 'vFlashErase' and 'vFlashWrite' packets until a 'vFlashDone' packet is received. The contents of the affected regions of flash memory are unpredictable until the 'vFlashDone' request is completed.

'vKill; pid'

Kill the process with the specified process ID. *pid* is a hexadecimal integer identifying the process. This packet is used in preference to 'k' when multiprocess protocol extensions are supported; see [multiprocess extensions], page 423.

Reply:

'E nn' for an error
'OK' for success

'vRun; filename[; argument]...'

Run the program filename, passing it each argument on its command line. The file and arguments are hex-encoded strings. If filename is an empty string, the stub may use a default program (e.g. the last program run). The program is created in the stopped state.

This packet is only available in extended mode (see [extended mode], page 406). Reply:

'E nn' for an error

'Any stop packet'

for success (see Section D.4 [Stop Reply Packets], page 413)

'vStopped'

In non-stop mode (see Section D.11 [Remote Non-Stop], page 434), acknowledge a previous stop reply and prompt for the stub to report another one.

Reply:

Debugging with GDB

'Any stop packet'

if there is another unreported stop event (see Section D.4 [Stop Reply Packets], page 413)

'OK' if there are no unreported stop events

'X addr, length: XX...'

Write data to memory, where the data is transmitted in binary. *addr* is address, *length* is number of bytes, 'XX...' is binary data (see [Binary Data], page 403). Reply:

'OK' for success

'z type, addr, kind'

'Z type, addr, kind'

Insert ('Z') or remove ('z') a type breakpoint or watchpoint starting at address address of kind kind.

Each breakpoint and watchpoint packet type is documented separately.

Implementation notes: A remote target shall return an empty string for an unrecognized breakpoint or watchpoint packet type. A remote target shall support either both or neither of a given 'Ztype...' and 'ztype...' packet pair. To avoid potential problems with duplicate packets, the operations should be implemented in an idempotent way.

'z0, addr, kind'

'Z0, addr, kind'

Insert ('20') or remove ('20') a memory breakpoint at address addr of type kind.

A memory breakpoint is implemented by replacing the instruction at addr with a software breakpoint or trap instruction. The kind is target-specific and typically indicates the size of the breakpoint in bytes that should be inserted. E.g., the ARM and MIPS can insert either a 2 or 4 byte breakpoint. Some architectures have additional meanings for kind; see Section D.6 [Architecture-Specific Protocol Details], page 428.

Implementation note: It is possible for a target to copy or move code that contains memory breakpoints (e.g., when implementing overlays). The behavior of this packet, in the presence of such a target, is not defined.

Reply:

'OK' success

'z1, addr, kind'

'Z1, addr, kind'

Insert ('Z1') or remove ('z1') a hardware breakpoint at address addr.

A hardware breakpoint is implemented using a mechanism that is not dependent on being able to modify the target's memory. *kind* has the same meaning as in 'Z0' packets.

Implementation note: A hardware breakpoint is not affected by code movement. Reply:

```
'OK'
                        success
'z2, addr, kind'
'Z2, addr, kind'
            Insert ('Z2') or remove ('Z2') a write watchpoint at addr. kind is interpreted
            as the number of bytes to watch.
            Reply:
            'OK'
                        success
'z3, addr, kind'
'Z3, addr, kind'
            Insert ('Z3') or remove ('z3') a read watchpoint at addr. kind is interpreted as
            the number of bytes to watch.
            Reply:
            'OK'
                        success
'z4, addr, kind'
'Z4, addr, kind'
            Insert ('Z4') or remove ('z4') an access watchpoint at addr. kind is interpreted
            as the number of bytes to watch.
            Reply:
            'OK'
                        success
```

Error responses use one of the following formats:

```
'Enn'
'E.text'
'Enn :text'
```

Error numbers are not well defined, any non-zero error number can be used to indicate an error. nn is a 2 character hex encoded value. The text string may provide additional information about the error.

D.4 Stop Reply Packets

The 'C', 'c', 'S', 's', 'vCont', 'vAttach', 'vRun', 'vStopped', and '?' packets can receive any of the below as a reply. Except for '?' and 'vStopped', that reply is only returned when the target halts. In the below the exact meaning of signal number is defined by the header 'include/gdb/signals.h' in the GDB source code.

As in the description of request packets, we include spaces in the reply templates for clarity; these are not part of the reply packet's syntax. No GDB stop reply packet uses spaces to separate its components.

'S AA' The program received signal number AA (a two-digit hexadecimal number). This is equivalent to a 'T' response with no n:r pairs.

```
'T AA n1:r1;n2:r2;...'
```

The program received signal number AA (a two-digit hexadecimal number). This is equivalent to an 'S' response, except that the 'n:r' pairs can carry values

of important registers and other information directly in the stop reply packet, reducing round-trip latency. Single-step and breakpoint traps are reported this way. Each n:r pair is interpreted as follows:

- If n is a hexadecimal number, it is a register number, and the corresponding r gives that register's value. r is a series of bytes in target byte order, with each byte given by a two-digit hex number.
- If n is 'thread', then r is the thread-id of the stopped thread, as specified in [thread-id syntax], page 405.
- If n is 'core', then r is the hexadecimal number of the core on which the stop event was detected.
- If n is a recognized stop reason, it describes a more specific event that stopped the target. The currently defined stop reasons are listed below. as should be '05', the trap signal. At most one stop reason should be present.
- Otherwise, GDB should ignore this 'n:r' pair and go on to the next; this allows us to extend the protocol in the future.

The currently defined stop reasons are:

'watch'

'rwatch'

'awatch' The packet indicates a watchpoint hit, and r is the data address, in hex.

'library' The packet indicates that the loaded libraries have changed. GDB should use 'qXfer:libraries:read' to fetch a new list of loaded libraries. r is ignored.

'replaylog'

The packet indicates that the target cannot continue replaying logged execution events, because it has reached the end (or the beginning when executing backward) of the log. The value of r will be either 'begin' or 'end'. See Chapter 6 [Reverse Execution], page 71, for more information.

'W AA'

'W AA ; process:pid'

The process exited, and AA is the exit status. This is only applicable to certain targets.

The second form of the response, including the process ID of the exited process, can be used only when GDB has reported support for multiprocess protocol extensions; see [multiprocess extensions], page 423. The *pid* is formatted as a big-endian hex string.

'X AA'

'X AA ; process:pid'

The process terminated with signal AA.

The second form of the response, including the process ID of the terminated process, can be used only when GDB has reported support for multiprocess protocol extensions; see [multiprocess extensions], page 423. The *pid* is formatted as a big-endian hex string.

'O XX...' is hex encoding of ASCII data, to be written as the program's console output. This can happen at any time while the program is running and the debugger should continue to wait for 'W', 'T', etc. This reply is not permitted in non-stop mode.

'F call-id, parameter...'

call-id is the identifier which says which host system call should be called. This is just the name of the function. Translation into the correct system call is only applicable as it's defined in GDB. See Section D.14 [File-I/O Remote Protocol Extension], page 437, for a list of implemented system calls.

'parameter...' is a list of parameters as defined for this very system call.

The target replies with this packet when it expects GDB to call a host system call on behalf of the target. GDB replies with an appropriate 'F' packet and keeps up waiting for the next reply packet from the target. The latest 'C', 'c', 'S' or 's' action is expected to be continued. See Section D.14 [File-I/O Remote Protocol Extension], page 437, for more details.

D.5 General Query Packets

Packets starting with 'q' are general query packets; packets starting with 'Q' are general set packets. General query and set packets are a semi-unified form for retrieving and sending information to and from the stub.

The initial letter of a query or set packet is followed by a name indicating what sort of thing the packet applies to. For example, GDB may use a 'qSymbol' packet to exchange symbol definitions with the stub. These packet names follow some conventions:

- The name must not contain commas, colons or semicolons.
- Most GDB query and set packets have a leading upper case letter.
- The names of custom vendor packets should use a company prefix, in lower case, followed by a period. For example, packets designed at the Acme Corporation might begin with 'qacme.foo' (for querying foos) or 'Qacme.bar' (for setting bars).

The name of a query or set packet should be separated from any parameters by a ':'; the parameters themselves should be separated by ',' or ';'. Stubs must be careful to match the full packet name, and check for a separator or the end of the packet, in case two packet names share a common prefix. New packets should not begin with 'qC', 'qP', or 'qL'¹.

Like the descriptions of the other packets, each description here has a template showing the packet's overall syntax, followed by an explanation of the packet's meaning. We include spaces in some of the templates for clarity; these are not part of the packet's syntax. No GDB packet uses spaces to separate its components.

Here are the currently defined query and set packets:

'qAuth: username, password'

Request permission to debug the target. The username and password are the values from set remote username and set remote password, both encoded as

¹ The 'qP' and 'qL' packets predate these conventions, and have arguments without any terminator for the packet name; we suspect they are in widespread use in places that are difficult to upgrade. The 'qC' packet has no arguments, but some existing stubs (e.g. RedBoot) are known to not check for the end of the packet.

hex strings. If the target accepts the username and password, then it replies with 'OK'; otherwise it replies with an error and may disconnect on its own. If the target requires a login ('MustAuth' feature), then if this packet is not supplied soon after connection, then the target may choose to return errors to all packets and/or disconnect.

'qC' Return the current thread ID.

Reply:

'QC thread-id'

Where thread-id is a thread ID as documented in [thread-id syntax], page 405.

'(anything else)'

Any other reply implies the old thread ID.

'qCRC: addr, length'

Compute the CRC checksum of a block of memory using CRC-32 defined in IEEE 802.3. The CRC is computed byte at a time, taking the most significant bit of each byte first. The initial pattern code <code>Oxffffffff</code> is used to ensure leading zeros affect the CRC.

Note: This is the same CRC used in validating separate debug files (see Section 18.2 [Debugging Information in Separate Files], page 193). However the algorithm is slightly different. When validating separate debug files, the CRC is computed taking the *least* significant bit of each byte first, and the final result is inverted to detect trailing zeros.

Reply:

'C crc32' The specified memory region's checksum is crc32.

```
'qfThreadInfo' 'qsThreadInfo'
```

Obtain a list of all active thread IDs from the target (OS). Since there may be too many active threads to fit into one reply packet, this query works iteratively: it may require more than one query/reply sequence to obtain the entire list of threads. The first query of the sequence will be the 'qfThreadInfo' query; subsequent queries in the sequence will be the 'qsThreadInfo' query.

NOTE: This packet replaces the 'qL' query (see below).

Reply:

'm thread-id'

A single thread ID

'm thread-id, thread-id...'

a comma-separated list of thread IDs

'1' (lower case letter 'L') denotes end of list.

In response to each query, the target will reply with a list of one or more thread IDs, separated by commas. GDB will respond to each reply with a request for more thread ids (using the 'qs' form of the query), until the target responds

with '1' (lower-case el, for last). Refer to [thread-id syntax], page 405, for the format of the thread-id fields.

'qGetTLSAddr: thread-id, offset, lm'

Fetch the address associated with thread local storage specified by thread-id, offset, and lm.

thread-id is the thread ID associated with the thread for which to fetch the TLS address. See [thread-id syntax], page 405.

offset is the (big endian, hex encoded) offset associated with the thread local variable. (This offset is obtained from the debug information associated with the variable.)

lm is the (big endian, hex encoded) OS/ABI-specific encoding of the the load module associated with the thread local storage. For example, a GNU/Linux system will pass the link map address of the shared object associated with the thread local storage under consideration. Other operating environments may choose to represent the load module differently, so the precise meaning of this parameter will vary.

Reply:

'XX...' Hex encoded (big endian) bytes representing the address of the thread local storage requested.

'qL startflag threadcount nextthread'

Obtain thread information from RTOS. Where: startflag (one hex digit) is one to indicate the first query and zero to indicate a subsequent query; threadcount (two hex digits) is the maximum number of threads the response packet can contain; and nextthread (eight hex digits), for subsequent queries (startflag is zero), is returned in the response as argthread.

Don't use this packet; use the 'qfThreadInfo' query instead (see above). Reply:

'qM count done argthread thread...'

Where: count (two hex digits) is the number of threads being returned; done (one hex digit) is zero to indicate more threads and one indicates no further threads; argthreadid (eight hex digits) is nextthread from the request packet; thread... is a sequence of thread IDs from the target. threadid (eight hex digits). See remote.c:parse_threadlist_response().

'qOffsets'

Get section offsets that the target used when relocating the downloaded image. Reply:

'Text=xxx; Data=yyy[; Bss=zzz]'

Relocate the Text section by xxx from its original address. Relocate the Data section by yyy from its original address. If the object file format provides segment information (e.g. ELF 'PT_LOAD' program headers), GDB will relocate entire segments by the supplied offsets.

Note: while a Bss offset may be included in the response, GDB ignores this and instead applies the Data offset to the Bss section.

'TextSeg=xxx[;DataSeg=yyy]'

Relocate the first segment of the object file, which conventionally contains program code, to a starting address of xxx. If 'DataSeg' is specified, relocate the second segment, which conventionally contains modifiable data, to a starting address of yyy. GDB will report an error if the object file does not contain segment information, or does not contain at least as many segments as mentioned in the reply. Extra segments are kept at fixed offsets relative to the last relocated segment.

'qP mode thread-id'

Returns information on thread-id. Where: mode is a hex encoded 32 bit mode; thread-id is a thread ID (see [thread-id syntax], page 405).

Don't use this packet; use the 'qThreadExtraInfo' query instead (see below).

Reply: see remote.c:remote_unpack_thread_info_response().

'QNonStop:1' 'QNonStop:0'

Enter non-stop ('QNonStop:1') or all-stop ('QNonStop:0') mode. See Section D.11 [Remote Non-Stop], page 434, for more information.

Reply:

'OK' The request succeeded.

'E nn' An error occurred. nn are hex digits.

'' An empty reply indicates that 'QNonStop' is not supported by the stub.

This packet is not probed by default; the remote stub must request it, by supplying an appropriate 'qSupported' response (see [qSupported], page 419). Use of this packet is controlled by the set non-stop command; see Section 5.4.2 [Non-Stop Mode], page 67.

'QPassSignals: signal [; signal]...'

Each listed signal should be passed directly to the inferior process. Signals are numbered identically to continue packets and stop replies (see Section D.4 [Stop Reply Packets], page 413). Each signal list item should be strictly greater than the previous item. These signals do not need to stop the inferior, or be reported to GDB. All other signals should be reported to GDB. Multiple 'QPassSignals' packets do not combine; any earlier 'QPassSignals' list is completely replaced by the new list. This packet improves performance when using 'handle signal nostop noprint pass'.

Reply:

'OK' The request succeeded.

Use of this packet is controlled by the set remote pass-signals command (see Section 20.4 [Remote Configuration], page 208). This packet is not probed

by default; the remote stub must request it, by supplying an appropriate 'qSupported' response (see [qSupported], page 419).

'qRcmd, command'

command (hex encoded) is passed to the local interpreter for execution. Invalid commands should be reported using the output string. Before the final result packet, the target may also respond with a number of intermediate 'Ooutput' console output packets. Implementors should note that providing access to a stubs's interpreter may have security implications.

Reply:

'OK' A command response with no output.

'OUTPUT' A command response with the hex encoded output string OUT-PUT.

(Note that the qRcmd packet's name is separated from the command by a ',', not a ':', contrary to the naming conventions above. Please don't use this packet as a model for new packets.)

'qSearch:memory:address;length;search-pattern'

Search length bytes at address for search-pattern. address and length are encoded in hex. search-pattern is a sequence of bytes, hex encoded.

Reply:

'0' The pattern was not found.

'1, address'

The pattern was found at address.

'E NN' A badly formed request or an error was encountered while searching memory.

An empty reply indicates that 'qSearch:memory' is not recognized.

'QStartNoAckMode'

Request that the remote stub disable the normal '+'/'-' protocol acknowledgments (see Section D.12 [Packet Acknowledgment], page 436).

Reply:

'OK' The stub has switched to no-acknowledgment mode. GDB acknowledges this reponse, but neither the stub nor GDB shall send or expect further '+'/'-' acknowledgments in the current connection.

An empty reply indicates that the stub does not support no-acknowledgment mode.

'qSupported [:gdbfeature [;gdbfeature]...]'

Tell the remote stub about features supported by GDB, and query the stub for features it supports. This packet allows GDB and the remote stub to take advantage of each others' features. 'qSupported' also consolidates multiple feature probes at startup, to improve GDB performance—a single larger packet performs better than multiple smaller probe packets on high-latency links. Some

features may enable behavior which must not be on by default, e.g. because it would confuse older clients or stubs. Other features may describe packets which could be automatically probed for, but are not. These features must be reported before GDB will use them. This "default unsupported" behavior is not appropriate for all packets, but it helps to keep the initial connection time under control with new versions of GDB which support increasing numbers of packets.

Reply:

'stubfeature [;stubfeature]...'

The stub supports or does not support each returned *stubfeature*, depending on the form of each *stubfeature* (see below for the possible forms).

The allowed forms for each feature (either a gdbfeature in the 'qSupported' packet, or a stubfeature in the response) are:

'name=value'

The remote protocol feature *name* is supported, and associated with the specified *value*. The format of *value* depends on the feature, but it must not include a semicolon.

'name+' The remote protocol feature name is supported, and does not need an associated value.

'name-' The remote protocol feature name is not supported.

'name?' The remote protocol feature name may be supported, and GDB should auto-detect support in some other way when it is needed. This form will not be used for gdbfeature notifications, but may be used for stubfeature responses.

Whenever the stub receives a 'qSupported' request, the supplied set of GDB features should override any previous request. This allows GDB to put the stub in a known state, even if the stub had previously been communicating with a different version of GDB.

The following values of gdbfeature (for the packet sent by GDB) are defined:

'multiprocess'

This feature indicates whether GDB supports multiprocess extensions to the remote protocol. GDB does not use such extensions unless the stub also reports that it supports them by including 'multiprocess+' in its 'qSupported' reply. See [multiprocess extensions], page 423, for details.

Stubs should ignore any unknown values for *gdbfeature*. Any GDB which sends a 'qSupported' packet supports receiving packets of unlimited length (earlier versions of GDB may reject overly long responses). Additional values for *gdbfeature* may be defined in the future to let the stub take advantage of new features in GDB, e.g. incompatible improvements in the remote protocol—the 'multiprocess' feature is an example of such a feature. The stub's reply should

be independent of the *gdbfeature* entries sent by gdb; first gdb describes all the features it supports, and then the stub replies with all the features it supports.

Similarly, GDB will silently ignore unrecognized stub feature responses, as long as each response uses one of the standard forms.

Some features are flags. A stub which supports a flag feature should respond with a '+' form response. Other features require values, and the stub should respond with an '=' form response.

Each feature has a default value, which GDB will use if 'qSupported' is not available or if the feature is not mentioned in the 'qSupported' response. The default values are fixed; a stub is free to omit any feature responses that match the defaults.

Not all features can be probed, but for those which can, the probing mechanism is useful: in some cases, a stub's internal architecture may not allow the protocol layer to know some information about the underlying target in advance. This is especially common in stubs which may be configured for multiple targets.

These are the currently defined stub features and their properties:

Feature Name	Value Required	Default	Probe Allowed
'PacketSize'	Yes	۷,	No
'qXfer:auxv:read'	No	٠_;	Yes
'qXfer:features:read'	No	·_;	Yes
'qXfer:libraries:read'	No	۷,	Yes
'qXfer:memory-map:read'	No	٠_;	Yes
'qXfer:spu:read'	No	۷,	Yes
'qXfer:spu:write'	No	۷,	Yes
'qXfer:siginfo:read'	No	·_;	Yes
'qXfer:siginfo:write'	No	٠_;	Yes
'qXfer:threads:read'	No	۷,	Yes
'QNonStop'	No	۷,	Yes
'QPassSignals'	No	۷,	Yes
'QStartNoAckMode'	No	·_,	Yes

Debugging with GDB

'multiprocess'	No	·_,	No
'MustAuth'	No	٠_,	No
`Conditional Trace points'	No	٠_,	No
'ReverseContinue'	No	٠_;	No
'ReverseStep'	No	·_,	No

These are the currently defined stub features, in more detail:

'PacketSize=bytes'

The remote stub can accept packets up to at least bytes in length. GDB will send packets up to this size for bulk transfers, and will never send larger packets. This is a limit on the data characters in the packet, including the frame and checksum. There is no trailing NUL byte in a remote protocol packet; if the stub stores packets in a NUL-terminated format, it should allow an extra byte in its buffer for the NUL. If this stub feature is not supported, GDB guesses based on the size of the 'g' packet response.

'qXfer:auxv:read'

The remote stub understands the 'qXfer:auxv:read' packet (see [qXfer auxiliary vector read], page 425).

'qXfer:features:read'

The remote stub understands the 'qXfer:features:read' packet (see [qXfer target description read], page 425).

'qXfer:libraries:read'

The remote stub understands the 'qXfer:libraries:read' packet (see [qXfer library list read], page 425).

'qXfer:memory-map:read'

The remote stub understands the 'qXfer:memory-map:read' packet (see [qXfer memory map read], page 425).

'qXfer:spu:read'

The remote stub understands the 'qXfer:spu:read' packet (see [qXfer spu read], page 426).

'qXfer:spu:write'

The remote stub understands the 'qXfer:spu:write' packet (see [qXfer spu write], page 427).

'qXfer:siginfo:read'

The remote stub understands the 'qXfer:siginfo:read' packet (see [qXfer siginfo read], page 426).

'qXfer:siginfo:write'

The remote stub understands the 'qXfer:siginfo:write' packet (see [qXfer siginfo write], page 427).

'qXfer:threads:read'

The remote stub understands the 'qXfer:threads:read' packet (see [qXfer threads read], page 426).

'QNonStop'

The remote stub understands the 'QNonStop' packet (see [QNonStop], page 418).

'QPassSignals'

The remote stub understands the 'QPassSignals' packet (see [QPassSignals], page 418).

'QStartNoAckMode'

The remote stub understands the 'QStartNoAckMode' packet and prefers to operate in no-acknowledgment mode. See Section D.12 [Packet Acknowledgment], page 436.

'multiprocess'

The remote stub understands the multiprocess extensions to the remote protocol syntax. The multiprocess extensions affect the syntax of thread IDs in both packets and replies (see [thread-id syntax], page 405), and add process IDs to the 'D' packet and 'W' and 'X' replies. Note that reporting this feature indicates support for the syntactic extensions only, not that the stub necessarily supports debugging of more than one process at a time. The stub must not use multiprocess extensions in packet replies unless GDB has also indicated it supports them in its 'qSupported' request.

'MustAuth'

GDB must send an authentication packet 'qAuth' and the username/password pair must be accepted, otherwise the target will return an error and disconnect.

'qXfer:osdata:read'

The remote stub understands the 'qXfer:osdata:read' packet ((see [qXfer osdata read], page 426).

'ConditionalTracepoints'

The remote stub accepts and implements conditional expressions defined for tracepoints (see Section 13.1.4 [Tracepoint Conditions], page 131).

'ReverseContinue'

The remote stub accepts and implements the reverse continue packet (see [bc], page 406).

'ReverseStep'

The remote stub accepts and implements the reverse step packet (see [bs], page 406).

'qSymbol::'

Notify the target that GDB is prepared to serve symbol lookup requests. Accept requests from the target for the values of symbols.

Debugging with GDB

Reply:

'OK' The target does not need to look up any (more) symbols.

'qSymbol:sym_name'

The target requests the value of symbol sym_name (hex encoded). GDB may provide the value by using the 'qSymbol: $sym_value:sym_name$ ' message, described below.

'qSymbol: sym_value: sym_name'

Set the value of sym_name to sym_value.

sym_name (hex encoded) is the name of a symbol whose value the target has previously requested.

sym_value (hex) is the value for symbol sym_name. If GDB cannot supply a value for sym_name, then this field will be empty.

Reply:

'OK' The target does not need to look up any (more) symbols.

'qSymbol: sym_name'

The target requests the value of a new symbol sym_name (hex encoded). GDB will continue to supply the values of symbols (if available), until the target ceases to request them.

```
'qTBuffer'
```

'QTDisconnected'

'QTDP'

'QTDV'

'qTfP'

'qTfV'

'QTFrame' See Section D.7 [Tracepoint Packets], page 429.

'qThreadExtraInfo, thread-id'

Obtain a printable string description of a thread's attributes from the target OS. thread-id is a thread ID; see [thread-id syntax], page 405. This string may contain anything that the target OS thinks is interesting for GDB to tell the user about the thread. The string is displayed in GDB's info threads display. Some examples of possible thread extra info strings are 'Runnable', or 'Blocked on Mutex'.

Reply:

'XX...' Where 'XX...' is a hex encoding of ASCII data, comprising the printable string containing the extra information about the thread's attributes

(Note that the qThreadExtraInfo packet's name is separated from the command by a ',', not a ':', contrary to the naming conventions above. Please don't use this packet as a model for new packets.)

'QTSave'

'qTsP'

```
'qTsV'
'QTStart'
'QTStop'
'QTinit'
'QTro'
'qTStatus'
'qTV' See Section D.7 [Tracepoint Packets], page 429.
```

'qXfer:object:read:annex:offset,length'

Read uninterpreted bytes from the target's special data area identified by the keyword *object*. Request *length* bytes starting at *offset* bytes into the data. The content and encoding of *annex* is specific to *object*; it can supply additional details about what data to access.

Here are the specific requests of this form defined so far. All 'qXfer:object:read:...' requests use the same reply formats, listed below.

'qXfer:auxv:read::offset,length'

Access the target's auxiliary vector. See Section 10.14 [OS Information], page 112. Note annex must be empty.

This packet is not probed by default; the remote stub must request it, by supplying an appropriate 'qSupported' response (see [qSupported], page 419).

'qXfer:features:read:annex:offset,length'

Access the target description. See Appendix G [Target Descriptions], page 467. The annex specifies which XML document to access. The main description is always loaded from the 'target.xml' annex.

This packet is not probed by default; the remote stub must request it, by supplying an appropriate 'qSupported' response (see [qSupported], page 419).

'qXfer:libraries:read:annex:offset,length'

Access the target's list of loaded libraries. See Section D.15 [Library List Format], page 449. The annex part of the generic 'qXfer' packet must be empty (see [qXfer read], page 425).

Targets which maintain a list of libraries in the program's memory do not need to implement this packet; it is designed for platforms where the operating system manages the list of loaded libraries.

This packet is not probed by default; the remote stub must request it, by supplying an appropriate 'qSupported' response (see [qSupported], page 419).

'qXfer:memory-map:read::offset,length'

Access the target's memory-map. See Section D.16 [Memory Map Format], page 450. The annex part of the generic 'qXfer' packet must be empty (see [qXfer read], page 425).

This packet is not probed by default; the remote stub must request it, by supplying an appropriate 'qSupported' response (see [qSupported], page 419).

'qXfer:siginfo:read::offset,length'

Read contents of the extra signal information on the target system. The annex part of the generic 'qXfer' packet must be empty (see [qXfer read], page 425).

This packet is not probed by default; the remote stub must request it, by supplying an appropriate 'qSupported' response (see [qSupported], page 419).

'qXfer:spu:read:annex:offset,length'

Read contents of an spufs file on the target system. The annex specifies which file to read; it must be of the form 'id/name', where id specifies an SPU context ID in the target process, and name identifies the spufs file in that context to be accessed.

This packet is not probed by default; the remote stub must request it, by supplying an appropriate 'qSupported' response (see [qSupported], page 419).

'qXfer:threads:read::offset,length'

Access the list of threads on target. See Section D.17 [Thread List Format], page 451. The annex part of the generic 'qXfer' packet must be empty (see [qXfer read], page 425).

This packet is not probed by default; the remote stub must request it, by supplying an appropriate 'qSupported' response (see [qSupported], page 419).

'qXfer:osdata:read::offset,length'

Access the target's operating system information. See Appendix H [Operating System Information], page 475.

Reply:

'm data' Data data (see [Binary Data], page 403) has been read from the target. There may be more data at a higher address (although it is permitted to return 'm' even for the last valid block of data, as long as at least one byte of data was read). data may have fewer bytes than the length in the request.

'1 data' Data data (see [Binary Data], page 403) has been read from the target. There is no more data to be read. data may have fewer bytes than the length in the request.

'1' The offset in the request is at the end of the data. There is no more data to be read.

'E00' The request was malformed, or annex was invalid.

'E nn' The offset was invalid, or there was an error encountered reading the data. nn is a hex-encoded errno value.

'Ye An empty reply indicates the *object* string was not recognized by the stub, or that the object does not support reading.

'qXfer:object:write:annex:offset:data...'

Write uninterpreted bytes into the target's special data area identified by the keyword *object*, starting at *offset* bytes into the data. *data*... is the binary-encoded data (see [Binary Data], page 403) to be written. The content and encoding of *annex* is specific to *object*; it can supply additional details about what data to access.

Here are the specific requests of this form defined so far. All 'qXfer:object:write:...' requests use the same reply formats, listed below.

'qXfer:siginfo:write::offset:data...'

Write data to the extra signal information on the target system. The annex part of the generic 'qXfer' packet must be empty (see [qXfer write], page 427).

This packet is not probed by default; the remote stub must request it, by supplying an appropriate 'qSupported' response (see [qSupported], page 419).

'qXfer:spu:write:annex:offset:data...'

Write data to an spufs file on the target system. The annex specifies which file to write; it must be of the form 'id/name', where id specifies an SPU context ID in the target process, and name identifies the spufs file in that context to be accessed.

This packet is not probed by default; the remote stub must request it, by supplying an appropriate 'qSupported' response (see [qSupported], page 419).

Reply:

'nn' nn (hex encoded) is the number of bytes written. This may be fewer bytes than supplied in the request.

'E00' The request was malformed, or annex was invalid.

'E nn' The offset was invalid, or there was an error encountered writing the data. nn is a hex-encoded errno value.

An empty reply indicates the *object* string was not recognized by the stub, or that the object does not support writing.

'qXfer:object:operation:...'

Requests of this form may be added in the future. When a stub does not recognize the *object* keyword, or its support for *object* does not recognize the *operation* keyword, the stub must respond with an empty packet.

'qAttached: pid'

Return an indication of whether the remote server attached to an existing process or created a new process. When the multiprocess protocol extensions are supported (see [multiprocess extensions], page 423), pid is an integer in hexadecimal format identifying the target process. Otherwise, GDB will omit the pid field and the query packet will be simplified as 'qAttached'.

This query is used, for example, to know whether the remote process should be detached or killed when a GDB session is ended with the quit command.

Reply:

'1' The remote server attached to an existing process.

'0' The remote server created a new process.

'E NN' A badly formed request or an error was encountered.

D.6 Architecture-Specific Protocol Details

This section describes how the remote protocol is applied to specific target architectures. Also see Section G.4 [Standard Target Features], page 471, for details of XML target descriptions for each architecture.

D.6.1 ARM

D.6.1.1 Breakpoint Kinds

These breakpoint kinds are defined for the 'Z0' and 'Z1' packets.

- 2 16-bit Thumb mode breakpoint.
- 3 32-bit Thumb mode (Thumb-2) breakpoint.
- 4 32-bit ARM mode breakpoint.

D.6.2 MIPS

D.6.2.1 Register Packet Format

The following g/G packets have previously been defined. In the below, some thirty-two bit registers are transferred as sixty-four bits. Those registers should be zero/sign extended (which?) to fill the space allocated. Register bytes are transferred in target byte order. The two nibbles within a register byte are transferred most-significant - least-significant.

MIPS32

All registers are transferred as thirty-two bit quantities in the order: 32 general-purpose; sr; lo; hi; bad; cause; pc; 32 floating-point registers; fsr; fir; fp.

MIPS64

All registers are transferred as sixty-four bit quantities (including thirty-two bit registers such as sr). The ordering is the same as MIPS32.

D.6.2.2 Breakpoint Kinds

These breakpoint kinds are defined for the 'Z0' and 'Z1' packets.

- 2 16-bit MIPS16 mode breakpoint.
- 3 16-bit microMIPS mode breakpoint.
- 4 32-bit standard MIPS mode breakpoint.
- 5 32-bit microMIPS mode breakpoint.

D.7 Tracepoint Packets

Here we describe the packets GDB uses to implement tracepoints (see Chapter 13 [Tracepoints], page 129).

'QTDP:n:addr:ena:step:pass[:Fflen][:Xlen,bytes][-]'

Create a new tracepoint, number n, at addr. If ena is 'E', then the tracepoint is enabled; if it is 'D', then the tracepoint is disabled. step is the tracepoint's step count, and pass is its pass count. If an 'F' is present, then the tracepoint is to be a fast tracepoint, and the flen is the number of bytes that the target should copy elsewhere to make room for the tracepoint. If an 'X' is present, it introduces a tracepoint condition, which consists of a hexadecimal length, followed by a comma and hex-encoded bytes, in a manner similar to action encodings as described below. If the trailing '-' is present, further 'QTDP' packets will follow to specify this tracepoint's actions.

Replies:

'OK' The packet was understood and carried out.

'QTDP: -n: addr: [S] action...[-]'

Define actions to be taken when a tracepoint is hit. *n* and *addr* must be the same as in the initial 'QTDP' packet for this tracepoint. This packet may only be sent immediately after another 'QTDP' packet that ended with a '-'. If the trailing '-' is present, further 'QTDP' packets will follow, specifying more actions for this tracepoint.

In the series of action packets for a given tracepoint, at most one can have an 'S' before its first action. If such a packet is sent, it and the following packets define "while-stepping" actions. Any prior packets define ordinary actions — that is, those taken when the tracepoint is first hit. If no action packet has an 'S', then all the packets in the series specify ordinary tracepoint actions.

The 'action...' portion of the packet is a series of actions, concatenated without separators. Each action has one of the following forms:

'R mask' Collect the registers whose bits are set in mask. mask is a hexadecimal number whose i'th bit is set if register number i should be collected. (The least significant bit is numbered zero.) Note that mask may be any number of digits long; it may not fit in a 32-bit word.

'M basereg, offset, len'

Collect len bytes of memory starting at the address in register number basereg, plus offset. If basereg is '-1', then the range has a fixed address: offset is the address of the lowest byte to collect. The basereg, offset, and len parameters are all unsigned hexadecimal values (the '-1' value for basereg is a special case).

'X len, expr'

Evaluate expr, whose length is len, and collect memory as it directs. expr is an agent expression, as described in Appendix E [Agent Expressions], page 453. Each byte of the expression is encoded as

a two-digit hex number in the packet; *len* is the number of bytes in the expression (and thus one-half the number of hex digits in the packet).

Any number of actions may be packed together in a single 'QTDP' packet, as long as the packet does not exceed the maximum packet length (400 bytes, for many stubs). There may be only one 'R' action per tracepoint, and it must precede any 'M' or 'X' actions. Any registers referred to by 'M' and 'X' actions must be collected by a preceding 'R' action. (The "while-stepping" actions are treated as if they were attached to a separate tracepoint, as far as these restrictions are concerned.)

Replies:

'OK' The packet was understood and carried out.

'QTDV:n:value'

Create a new trace state variable, number n, with an initial value of value, which is a 64-bit signed integer. Both n and value are encoded as hexadecimal values. GDB has the option of not using this packet for initial values of zero; the target should simply create the trace state variables as they are mentioned in expressions.

'QTFrame:n'

Select the *n*'th tracepoint frame from the buffer, and use the register and memory contents recorded there to answer subsequent request packets from GDB.

A successful reply from the stub indicates that the stub has found the requested frame. The response is a series of parts, concatenated without separators, describing the frame we selected. Each part has one of the following forms:

- 'F f' The selected frame is number n in the trace frame buffer; f is a hexadecimal number. If f is '-1', then there was no frame matching the criteria in the request packet.
- 'T t' The selected trace frame records a hit of tracepoint number t; t is a hexadecimal number.

'QTFrame:pc:addr'

Like 'QTFrame:n', but select the first tracepoint frame after the currently selected frame whose PC is addr; addr is a hexadecimal number.

'QTFrame:tdp:t'

Like 'QTFrame:n', but select the first tracepoint frame after the currently selected frame that is a hit of tracepoint t; t is a hexadecimal number.

'QTFrame:range:start:end'

Like 'QTFrame:n', but select the first tracepoint frame after the currently selected frame whose PC is between *start* (inclusive) and *end* (inclusive); *start* and *end* are hexadecimal numbers.

'QTFrame: outside: start: end'

Like 'QTFrame:range:start:end', but select the first frame outside the given range of addresses (exclusive).

'QTStart' Begin the tracepoint experiment. Begin collecting data from tracepoint hits in the trace frame buffer.

'QTStop' End the tracepoint experiment. Stop collecting trace frames.

'QTinit' Clear the table of tracepoints, and empty the trace frame buffer.

'QTro:start1,end1:start2,end2:...'

Establish the given ranges of memory as "transparent". The stub will answer requests for these ranges from memory's current contents, if they were not collected as part of the tracepoint hit.

GDB uses this to mark read-only regions of memory, like those containing program code. Since these areas never change, they should still have the same contents they did when the tracepoint was hit, so there's no reason for the stub to refuse to provide their contents.

'QTDisconnected: value'

Set the choice to what to do with the tracing run when GDB disconnects from the target. A *value* of 1 directs the target to continue the tracing run, while 0 tells the target to stop tracing if GDB is no longer in the picture.

'qTStatus'

Ask the stub if there is a trace experiment running right now.

Replies:

'T0' There is no trace experiment running.

'T1' There is a trace experiment running.

'qTV: var' Ask the stub for the value of the trace state variable number var.

Replies:

'Vvalue' The value of the variable is value. This will be the current value of the variable if the user is examining a running target, or a saved value if the variable was collected in the trace frame that the user is looking at. Note that multiple requests may result in different reply values, such as when requesting values while the program is running.

'U' The value of the variable is unknown. This would occur, for example, if the user is examining a trace frame in which the requested variable was not collected.

'qTfP' 'qTsP'

These packets request data about tracepoints that are being used by the target. GDB sends qTfP to get the first piece of data, and multiple qTsP to get additional pieces. Replies to these packets generally take the form of the QTDP packets that define tracepoints. (FIXME add detailed syntax)

'qTfV'

'qTsV' These packets request data about trace state variables that are on the target.

GDB sends qTfV to get the first vari of data, and multiple qTsV to get additional variables. Replies to these packets follow the syntax of the QTDV packets that define trace state variables.

'QTSave:filename'

This packet directs the target to save trace data to the file name filename in the target's filesystem. filename is encoded as a hex string; the interpretation of the file name (relative vs absolute, wild cards, etc) is up to the target.

'qTBuffer: offset, len'

Return up to *len* bytes of the current contents of trace buffer, starting at *offset*. The trace buffer is treated as if it were a contiguous collection of traceframes, as per the trace file format. The reply consists as many hex-encoded bytes as the target can deliver in a packet; it is not an error to return fewer than were asked for. A reply consisting of just 1 indicates that no bytes are available.

D.8 Host I/O Packets

The Host I/O packets allow GDB to perform I/O operations on the far side of a remote link. For example, Host I/O is used to upload and download files to a remote target with its own filesystem. Host I/O uses the same constant values and data structure layout as the target-initiated File-I/O protocol. However, the Host I/O packets are structured differently. The target-initiated protocol relies on target memory to store parameters and buffers. Host I/O requests are initiated by GDB, and the target's memory is not involved. See Section D.14 [File-I/O Remote Protocol Extension], page 437, for more details on the target-initiated protocol.

The Host I/O request packets all encode a single operation along with its arguments. They have this format:

'vFile:operation: parameter...'

operation is the name of the particular request; the target should compare the entire packet name up to the second colon when checking for a supported operation. The format of parameter depends on the operation. Numbers are always passed in hexadecimal. Negative numbers have an explicit minus sign (i.e. two's complement is not used). Strings (e.g. filenames) are encoded as a series of hexadecimal bytes. The last argument to a system call may be a buffer of escaped binary data (see [Binary Data], page 403).

The valid responses to Host I/O packets are:

'F result [, errno] [; attachment]'

result is the integer value returned by this operation, usually non-negative for success and -1 for errors. If an error has occured, errno will be included in the result. errno will have a value defined by the File-I/O protocol (see [Errno Values], page 447). For operations which return data, attachment supplies the data as a binary buffer. Binary buffers in response packets are escaped in the normal way (see [Binary Data], page 403). See the individual packet documentation for the interpretation of result and attachment.

An empty response indicates that this operation is not recognized.

These are the supported Host I/O operations:

'vFile:open: pathname, flags, mode'

٤,

Open a file at *pathname* and return a file descriptor for it, or return -1 if an error occurs. *pathname* is a string, *flags* is an integer indicating a mask of open

flags (see [Open Flags], page 447), and *mode* is an integer indicating a mask of mode bits to use if the file is created (see [mode_t Values], page 447). See [open], page 440, for details of the open flags and mode values.

'vFile:close: fd'

Close the open file corresponding to fd and return 0, or -1 if an error occurs.

'vFile:pread: fd, count, offset'

Read data from the open file corresponding to fd. Up to count bytes will be read from the file, starting at offset relative to the start of the file. The target may read fewer bytes; common reasons include packet size limits and an end-of-file condition. The number of bytes read is returned. Zero should only be returned for a successful read at the end of the file, or if count was zero.

The data read should be returned as a binary attachment on success. If zero bytes were read, the response should include an empty binary attachment (i.e. a trailing semicolon). The return value is the number of target bytes read; the binary attachment may be longer if some characters were escaped.

'vFile:pwrite: fd, offset, data'

Write data (a binary buffer) to the open file corresponding to fd. Start the write at offset from the start of the file. Unlike many write system calls, there is no separate count argument; the length of data in the packet is used. 'vFile:write' returns the number of bytes written, which may be shorter than the length of data, or -1 if an error occurred.

'vFile:unlink: pathname'

Delete the file at *pathname* on the target. Return 0, or -1 if an error occurs. *pathname* is a string.

D.9 Interrupts

When a program on the remote target is running, GDB may attempt to interrupt it by sending a 'Ctrl-C', BREAK or a BREAK followed by g, control of which is specified via GDB's 'interrupt-sequence'.

The precise meaning of BREAK is defined by the transport mechanism and may, in fact, be undefined. GDB does not currently define a BREAK mechanism for any of the network interfaces except for TCP, in which case GDB sends the telnet BREAK sequence.

'Ctrl-C', on the other hand, is defined and implemented for all transport mechanisms. It is represented by sending the single byte 0x03 without any of the usual packet overhead described in the Overview section (see Section D.1 [Overview], page 403). When a 0x03 byte is transmitted as part of a packet, it is considered to be packet data and does *not* represent an interrupt. E.g., an 'X' packet (see [X packet], page 412), used for binary downloads, may include an unescaped 0x03 as part of its packet.

BREAK followed by g is also known as Magic SysRq g. When Linux kernel receives this sequence from serial port, it stops execution and connects to gdb.

Stubs are not required to recognize these interrupt mechanisms and the precise meaning associated with receipt of the interrupt is implementation defined. If the target supports

debugging of multiple threads and/or processes, it should attempt to interrupt all currently-executing threads and processes. If the stub is successful at interrupting the running program, it should send one of the stop reply packets (see Section D.4 [Stop Reply Packets], page 413) to GDB as a result of successfully stopping the program in all-stop mode, and a stop reply for each stopped thread in non-stop mode. Interrupts received while the program is stopped are discarded.

D.10 Notification Packets

The GDB remote serial protocol includes *notifications*, packets that require no acknowledgment. Both the GDB and the stub may send notifications (although the only notifications defined at present are sent by the stub). Notifications carry information without incurring the round-trip latency of an acknowledgment, and so are useful for low-impact communications where occasional packet loss is not a problem.

A notification packet has the form '% data # checksum', where data is the content of the notification, and checksum is a checksum of data, computed and formatted as for ordinary GDB packets. A notification's data never contains '\$', '%' or '#' characters. Upon receiving a notification, the recipient sends no '+' or '-' to acknowledge the notification's receipt or to report its corruption.

Every notification's data begins with a name, which contains no colon characters, followed by a colon character.

Recipients should silently ignore corrupted notifications and notifications they do not understand. Recipients should restart timeout periods on receipt of a well-formed notification, whether or not they understand it.

Senders should only send the notifications described here when this protocol description specifies that they are permitted. In the future, we may extend the protocol to permit existing notifications in new contexts; this rule helps older senders avoid confusing newer recipients.

(Older versions of GDB ignore bytes received until they see the '\$' byte that begins an ordinary packet, so new stubs may transmit notifications without fear of confusing older clients. There are no notifications defined for GDB to send at the moment, but we assume that most older stubs would ignore them, as well.)

The following notification packets from the stub to GDB are defined:

'Stop: reply'

Report an asynchronous stop event in non-stop mode. The *reply* has the form of a stop reply, as described in Section D.4 [Stop Reply Packets], page 413. Refer to Section D.11 [Remote Non-Stop], page 434, for information on how these notifications are acknowledged by GDB.

D.11 Remote Protocol Support for Non-Stop Mode

GDB's remote protocol supports non-stop debugging of multi-threaded programs, as described in Section 5.4.2 [Non-Stop Mode], page 67. If the stub supports non-stop mode, it should report that to GDB by including 'QNonStop+' in its 'qSupported' response (see [qSupported], page 419).

GDB typically sends a 'QNonStop' packet only when establishing a new connection with the stub. Entering non-stop mode does not alter the state of any currently-running threads, but targets must stop all threads in any already-attached processes when entering all-stop mode. GDB uses the '?' packet as necessary to probe the target state after a mode change.

In non-stop mode, when an attached process encounters an event that would otherwise be reported with a stop reply, it uses the asynchronous notification mechanism (see Section D.10 [Notification Packets], page 434) to inform GDB. In contrast to all-stop mode, where all threads in all processes are stopped when a stop reply is sent, in non-stop mode only the thread reporting the stop event is stopped. That is, when reporting a 'S' or 'T' response to indicate completion of a step operation, hitting a breakpoint, or a fault, only the affected thread is stopped; any other still-running threads continue to run. When reporting a 'W' or 'X' response, all running threads belonging to other attached processes continue to run.

Only one stop reply notification at a time may be pending; if additional stop events occur before GDB has acknowledged the previous notification, they must be queued by the stub for later synchronous transmission in response to 'vStopped' packets from GDB. Because the notification mechanism is unreliable, the stub is permitted to resend a stop reply notification if it believes GDB may not have received it. GDB ignores additional stop reply notifications received before it has finished processing a previous notification and the stub has completed sending any queued stop events.

Otherwise, GDB must be prepared to receive a stop reply notification at any time. Specifically, they may appear when GDB is not otherwise reading input from the stub, or when GDB is expecting to read a normal synchronous response or a '+'/'-' acknowledgment to a packet it has sent. Notification packets are distinct from any other communication from the stub so there is no ambiguity.

After receiving a stop reply notification, GDB shall acknowledge it by sending a 'vStopped' packet (see [vStopped packet], page 411) as a regular, synchronous request to the stub. Such acknowledgment is not required to happen immediately, as GDB is permitted to send other, unrelated packets to the stub first, which the stub should process normally.

Upon receiving a 'vStopped' packet, if the stub has other queued stop events to report to GDB, it shall respond by sending a normal stop reply response. GDB shall then send another 'vStopped' packet to solicit further responses; again, it is permitted to send other, unrelated packets as well which the stub should process normally.

If the stub receives a 'vStopped' packet and there are no additional stop events to report, the stub shall return an 'OK' response. At this point, if further stop events occur, the stub shall send a new stop reply notification, GDB shall accept the notification, and the process shall be repeated.

In non-stop mode, the target shall respond to the '?' packet as follows. First, any incomplete stop reply notification/'vStopped' sequence in progress is abandoned. The target must begin a new sequence reporting stop events for all stopped threads, whether or not it has previously reported those events to GDB. The first stop reply is sent as a synchronous reply to the '?' packet, and subsequent stop replies are sent as responses to 'vStopped' packets using the mechanism described above. The target must not send asynchronous stop

reply notifications until the sequence is complete. If all threads are running when the target receives the '?' packet, or if the target is not attached to any process, it shall respond 'OK'.

D.12 Packet Acknowledgment

By default, when either the host or the target machine receives a packet, the first response expected is an acknowledgment: either '+' (to indicate the package was received correctly) or '-' (to request retransmission). This mechanism allows the GDB remote protocol to operate over unreliable transport mechanisms, such as a serial line.

In cases where the transport mechanism is itself reliable (such as a pipe or TCP connection), the '+'/'-' acknowledgments are redundant. It may be desirable to disable them in that case to reduce communication overhead, or for other reasons. This can be accomplished by means of the 'QStartNoAckMode' packet; see [QStartNoAckMode], page 419.

When in no-acknowledgment mode, neither the stub nor GDB shall send or expect '+'/'-' protocol acknowledgments. The packet and response format still includes the normal checksum, as described in Section D.1 [Overview], page 403, but the checksum may be ignored by the receiver.

If the stub supports 'QStartNoAckMode' and prefers to operate in no-acknowledgment mode, it should report that to GDB by including 'QStartNoAckMode+' in its response to 'qSupported'; see [qSupported], page 419. If GDB also supports 'QStartNoAckMode' and it has not been disabled via the set remote noack-packet off command (see Section 20.4 [Remote Configuration], page 208), GDB may then send a 'QStartNoAckMode' packet to the stub. Only then may the stub actually turn off packet acknowledgments. GDB sends a final '+' acknowledgment of the stub's 'OK' response, which can be safely ignored by the stub.

Note that set remote noack-packet command only affects negotiation between GDB and the stub when subsequent connections are made; it does not affect the protocol acknowledgment state for any current connection. Since '+'/'-' acknowledgments are enabled by default when a new connection is established, there is also no protocol request to reenable the acknowledgments for the current connection, once disabled.

D.13 Examples

Example sequence of a target being re-started. Notice how the restart does not get any direct output:

```
-> R00
<- +
target restarts
-> ?
<- +
<- T001:1234123412341234
```

Example sequence of a target being stepped by a single instruction:

```
-> G1445...

<- +

-> s

<- +

time passes

<- T001:1234123412341234

-> +
```

```
-> g
<- +
<- 1455...
```

D.14 File-I/O Remote Protocol Extension

D.14.1 File-I/O Overview

The File I/O remote protocol extension (short: File-I/O) allows the target to use the host's file system and console I/O to perform various system calls. System calls on the target system are translated into a remote protocol packet to the host system, which then performs the needed actions and returns a response packet to the target system. This simulates file system operations even on targets that lack file systems.

The protocol is defined to be independent of both the host and target systems. It uses its own internal representation of datatypes and values. Both gdb and the target's gdb stub are responsible for translating the system-dependent value representations into the internal protocol representations when data is transmitted.

The communication is synchronous. A system call is possible only when GDB is waiting for a response from the 'C', 'c', 'S' or 's' packets. While GDB handles the request for a system call, the target is stopped to allow deterministic access to the target's memory. Therefore File-I/O is not interruptible by target signals. On the other hand, it is possible to interrupt File-I/O by a user interrupt ('Ctrl-C') within GDB.

The target's request to perform a host system call does not finish the latest 'C', 'c', 'S' or 's' action. That means, after finishing the system call, the target returns to continuing the previous activity (continue, step). No additional continue or step request from GDB is required.

```
(gdb) continue
  <- target requests 'system call X'
  target is stopped, GDB executes system call
  -> GDB returns result
  ... target continues, GDB returns to wait for the target
  <- target hits breakpoint and sends a Txx packet</pre>
```

The protocol only supports I/O on the console and to regular files on the host file system. Character or block special devices, pipes, named pipes, sockets or any other communication method on the host system are not supported by this protocol.

File I/O is not supported in non-stop mode.

D.14.2 Protocol Basics

The File-I/O protocol uses the F packet as the request as well as reply packet. Since a File-I/O system call can only occur when GDB is waiting for a response from the continuing or stepping target, the File-I/O request is a reply that GDB has to expect as a result of a previous 'C', 'c', 'S' or 's' packet. This F packet contains all information needed to allow GDB to call the appropriate host system call:

- A unique identifier for the requested system call.
- All parameters to the system call. Pointers are given as addresses in the target memory address space. Pointers to strings are given as pointer/length pair. Numerical values are given as they are. Numerical control flags are given in a protocol-specific representation.

At this point, GDB has to perform the following actions.

• If the parameters include pointer values to data needed as input to a system call, GDB requests this data from the target with a standard m packet request. This additional communication has to be expected by the target implementation and is handled as any other m packet.

- GDB translates all value from protocol representation to host representation as needed. Datatypes are coerced into the host types.
- GDB calls the system call.
- It then coerces datatypes back to protocol representation.
- If the system call is expected to return data in buffer space specified by pointer parameters to the call, the data is transmitted to the target using a M or X packet. This packet has to be expected by the target implementation and is handled as any other M or X packet.

Eventually GDB replies with another F packet which contains all necessary information for the target to continue. This at least contains

- Return value.
- errno, if has been changed by the system call.
- "Ctrl-C" flag.

After having done the needed type and value coercion, the target continues the latest continue or step action.

D.14.3 The F Request Packet

The F request packet has the following format:

'Fcall-id, parameter...'

call-id is the identifier to indicate the host system call to be called. This is just the name of the function.

parameter... are the parameters to the system call. Parameters are hexadecimal integer values, either the actual values in case of scalar datatypes, pointers to target buffer space in case of compound datatypes and unspecified memory areas, or pointer/length pairs in case of string parameters. These are appended to the *call-id* as a comma-delimited list. All values are transmitted in ASCII string representation, pointer/length pairs separated by a slash.

D.14.4 The F Reply Packet

The F reply packet has the following format:

'Fretcode, errno, Ctrl-C flag; call-specific attachment'

retcode is the return code of the system call as hexadecimal value.

errno is the errno set by the call, in protocol-specific representation. This parameter can be omitted if the call was successful.

Ctrl-C flag is only sent if the user requested a break. In this case, errno must be sent as well, even if the call was successful. The Ctrl-C flag itself consists of the character 'C':

F0,0,C

or, if the call was interrupted before the host call has been performed:

F-1.4.0

assuming 4 is the protocol-specific representation of EINTR.

D.14.5 The 'Ctrl-C' Message

If the 'Ctrl-C' flag is set in the GDB reply packet (see Section D.14.4 [The F Reply Packet], page 438), the target should behave as if it had gotten a break message. The meaning for the target is "system call interrupted by SIGINT". Consequentially, the target should actually stop (as with a break message) and return to GDB with a TO2 packet.

It's important for the target to know in which state the system call was interrupted. There are two possible cases:

- The system call hasn't been performed on the host yet.
- The system call on the host has been finished.

These two states can be distinguished by the target by the value of the returned errno. If it's the protocol representation of EINTR, the system call hasn't been performed. This is equivalent to the EINTR handling on POSIX systems. In any other case, the target may presume that the system call has been finished — successfully or not — and should behave as if the break message arrived right after the system call.

GDB must behave reliably. If the system call has not been called yet, GDB may send the F reply immediately, setting EINTR as errno in the packet. If the system call on the host has been finished before the user requests a break, the full action must be finished by GDB. This requires sending M or X packets as necessary. The F packet may only be sent when either nothing has happened or the full action has been completed.

D.14.6 Console I/O

By default and if not explicitly closed by the target system, the file descriptors 0, 1 and 2 are connected to the GDB console. Output on the GDB console is handled as any other file output operation (write(1, ...) or write(2, ...)). Console input is handled by GDB so that after the target read request from file descriptor 0 all following typing is buffered until either one of the following conditions is met:

- The user types Ctrl-c. The behaviour is as explained above, and the read system call is treated as finished.
- The user presses RET. This is treated as end of input with a trailing newline.
- The user types Ctrl-d. This is treated as end of input. No trailing character (neither newline nor 'Ctrl-D') is appended to the input.

If the user has typed more characters than fit in the buffer given to the read call, the trailing characters are buffered in GDB until either another read(0, ...) is requested by the target, or debugging is stopped at the user's request.

D.14.7 List of Supported Calls

open

Synopsis:

```
int open(const char *pathname, int flags);
int open(const char *pathname, int flags, mode_t mode);
```

Request:

'Fopen, pathptr/len, flags, mode'

flags is the bitwise OR of the following values:

O_CREAT If the file does not exist it will be created. The host rules apply as far as file ownership and time stamps are concerned.

O_EXCL When used with O_CREAT, if the file already exists it is an error and open() fails.

O_TRUNC If the file already exists and the open mode allows writing (O_RDWR or O_WRONLY is given) it will be truncated to zero length.

O_APPEND The file is opened in append mode.

O_RDONLY The file is opened for reading only.

O_WRONLY The file is opened for writing only.

O_RDWR The file is opened for reading and writing.

Other bits are silently ignored.

mode is the bitwise OR of the following values:

S_IRUSR User has read permission.

S_IWUSR User has write permission.

S_IRGRP Group has read permission.

S_IWGRP Group has write permission.

S_IROTH Others have read permission.

S_IWOTH Others have write permission.

Other bits are silently ignored.

Return value:

open returns the new file descriptor or -1 if an error occurred.

Errors:

EEXIST pathname already exists and O_CREAT and O_EXCL were used.

EISDIR pathname refers to a directory.

EACCES The requested access is not allowed.

ENAMETOOLONG

pathname was too long.

ENOENT A directory component in pathname does not exist.

ENODEV pathname refers to a device, pipe, named pipe or socket.

EROFS pathname refers to a file on a read-only filesystem and write access

was requested.

EFAULT pathname is an invalid pointer value.

ENOSPC No space on device to create the file.

EMFILE The process already has the maximum number of files open.

ENFILE The limit on the total number of files open on the system has been

reached.

EINTR The call was interrupted by the user.

close

Synopsis:

int close(int fd);

Request: 'Fclose, fd'

Return value:

close returns zero on success, or -1 if an error occurred.

Errors:

EBADF fd isn't a valid open file descriptor.

EINTR The call was interrupted by the user.

read

Synopsis:

int read(int fd, void *buf, unsigned int count);

Request: 'Fread, fd, bufptr, count'

Return value:

On success, the number of bytes read is returned. Zero indicates end of file. If count is zero, read returns zero as well. On error, -1 is returned.

Errors:

EBADF fd is not a valid file descriptor or is not open for reading.

EFAULT bufptr is an invalid pointer value.

EINTR The call was interrupted by the user.

write

Synopsis:

int write(int fd, const void *buf, unsigned int count);

Request: 'Fwrite, fd, bufptr, count'

Return value:

On success, the number of bytes written are returned. Zero indicates nothing was written. On error, -1 is returned.

Errors:

EBADF fd is not a valid file descriptor or is not open for writing.

EFAULT bufptr is an invalid pointer value.

EFBIG An attempt was made to write a file that exceeds the host-specific

maximum file size allowed.

ENOSPC No space on device to write the data.

EINTR The call was interrupted by the user.

lseek

Synopsis:

long lseek (int fd, long offset, int flag);

Request: 'Flseek, fd, offset, flag'

flag is one of:

SEEK_SET The offset is set to offset bytes.

SEEK_CUR The offset is set to its current location plus offset bytes.

SEEK_END The offset is set to the size of the file plus offset bytes.

Return value:

On success, the resulting unsigned offset in bytes from the beginning of the file is returned. Otherwise, a value of -1 is returned.

Errors:

EBADF fd is not a valid open file descriptor.

ESPIPE fd is associated with the GDB console.

EINVAL flag is not a proper value.

EINTR The call was interrupted by the user.

rename

Synopsis:

int rename(const char *oldpath, const char *newpath);

Request: 'Frename, oldpathptr/len, newpathptr/len'

Return value:

On success, zero is returned. On error, -1 is returned.

Errors:

EISDIR newpath is an existing directory, but oldpath is not a directory.

EEXIST newpath is a non-empty directory.

EBUSY oldpath or newpath is a directory that is in use by some process.

EINVAL An attempt was made to make a directory a subdirectory of itself.

ENOTDIR A component used as a directory in oldpath or new path is not a

directory. Or oldpath is a directory and newpath exists but is not

a directory.

EFAULT oldpathptr or newpathptr are invalid pointer values.

EACCES No access to the file or the path of the file.

ENAMETOOLONG

oldpath or newpath was too long.

ENDENT A directory component in oldpath or newpath does not exist.

EROFS The file is on a read-only filesystem.

ENOSPC The device containing the file has no room for the new directory

entry.

EINTR The call was interrupted by the user.

unlink

Synopsis:

int unlink(const char *pathname);

Request: 'Funlink, pathnameptr/len'

Return value:

On success, zero is returned. On error, -1 is returned.

Errors:

EACCES No access to the file or the path of the file.

EPERM The system does not allow unlinking of directories.

EBUSY The file pathname cannot be unlinked because it's being used by

another process.

EFAULT pathnameptr is an invalid pointer value.

ENAMETOOLONG

pathname was too long.

ENOENT A directory component in pathname does not exist.

ENOTDIR A component of the path is not a directory.

EROFS The file is on a read-only filesystem.

EINTR The call was interrupted by the user.

stat/fstat

Synopsis:

int stat(const char *pathname, struct stat *buf);
int fstat(int fd, struct stat *buf);

Request: 'Fstat, pathnameptr/len, bufptr'

'Ffstat, fd, bufptr'

Return value:

On success, zero is returned. On error, -1 is returned.

Errors:

EBADF fd is not a valid open file.

ENOENT A directory component in pathname does not exist or the path is

an empty string.

ENOTDIR A component of the path is not a directory.

EFAULT pathnameptr is an invalid pointer value.

EACCES No access to the file or the path of the file.

ENAMETOOLONG

pathname was too long.

EINTR The call was interrupted by the user.

gettimeofday

Synopsis:

int gettimeofday(struct timeval *tv, void *tz);

Request: 'Fgettimeofday, tvptr, tzptr'

Return value:

On success, 0 is returned, -1 otherwise.

Errors:

EINVAL tz is a non-NULL pointer.

EFAULT typtr and/or typtr is an invalid pointer value.

isatty

Synopsis:

int isatty(int fd);

Request: 'Fisatty, fd'

Return value:

Returns 1 if fd refers to the GDB console, 0 otherwise.

Errors:

EINTR The call was interrupted by the user.

Note that the isatty call is treated as a special case: it returns 1 to the target if the file descriptor is attached to the GDB console, 0 otherwise. Implementing through system calls would require implementing ioctl and would be more complex than needed.

system

Synopsis:

int system(const char *command);

Request: 'Fsystem, commandptr/len'

Return value:

If len is zero, the return value indicates whether a shell is available. A zero return value indicates a shell is not available. For non-zero len, the value returned is -1 on error and the return status of the command otherwise. Only the exit status of the command is returned, which is extracted from the host's system return value by calling WEXITSTATUS(retval). In case '/bin/sh' could not be executed, 127 is returned.

Errors:

EINTR The call was interrupted by the user.

GDB takes over the full task of calling the necessary host calls to perform the system call. The return value of system on the host is simplified before it's returned to the target. Any termination signal information from the child process is discarded, and the return value consists entirely of the exit status of the called command.

Due to security concerns, the system call is by default refused by GDB. The user has to allow this call explicitly with the set remote system-call-allowed 1 command.

set remote system-call-allowed

Control whether to allow the system calls in the File I/O protocol for the remote target. The default is zero (disabled).

show remote system-call-allowed

Show whether the system calls are allowed in the File I/O protocol.

D.14.8 Protocol-specific Representation of Datatypes

Integral Datatypes

The integral datatypes used in the system calls are int, unsigned int, long, unsigned long, mode_t, and time_t.

int, unsigned int, mode_t and time_t are implemented as 32 bit values in this protocol. long and unsigned long are implemented as 64 bit types.

See [Limits], page 448, for corresponding MIN and MAX values (similar to those in 'limits.h') to allow range checking on host and target.

time_t datatypes are defined as seconds since the Epoch.

All integral datatypes transferred as part of a memory read or write of a structured datatype e.g. a struct stat have to be given in big endian byte order.

Pointer Values

Pointers to target data are transmitted as they are. An exception is made for pointers to buffers for which the length isn't transmitted as part of the function call, namely strings. Strings are transmitted as a pointer/length pair, both as hex values, e.g.

```
1aaf/12
```

which is a pointer to data of length 18 bytes at position 0x1aaf. The length is defined as the full string length in bytes, including the trailing null byte. For example, the string "hello world" at address 0x123456 is transmitted as

123456/d

Memory Transfer

Structured data which is transferred using a memory read or write (for example, a struct stat) is expected to be in a protocol-specific format with all scalar multibyte datatypes being big endian. Translation to this representation needs to be done both by the target before the F packet is sent, and by GDB before it transfers memory to the target. Transferred pointers to structured data should point to the already-coerced data at any time.

struct stat

The buffer of type struct stat used by the target and GDB is defined as follows:

```
struct stat {
   unsigned int st_dev;
                             /* device */
   unsigned int st_ino;
                             /* inode */
   mode_t
                            /* protection */
                st_mode;
   unsigned int st_nlink;
                             /* number of hard links */
   unsigned int st_uid;
                            /* user ID of owner */
   unsigned int st_gid;
                            /* group ID of owner */
   unsigned int st_rdev;
                            /* device type (if inode device) */
                            /* total size, in bytes */
   unsigned long st_size;
   unsigned long st_blksize; /* blocksize for filesystem I/O */
   unsigned long st_blocks; /* number of blocks allocated */
   time_t
               st_atime;
                            /* time of last access */
   time_t
                st_mtime;
                            /* time of last modification */
   time_t
                st_ctime;
                             /* time of last change */
};
```

The integral datatypes conform to the definitions given in the appropriate section (see [Integral Datatypes], page 445, for details) so this structure is of size 64 bytes.

The values of several fields have a restricted meaning and/or range of values.

```
st_dev A value of 0 represents a file, 1 the console.

st_ino No valid meaning for the target. Transmitted unchanged.

st_mode Valid mode bits are described in Section D.14.9 [Constants], page 447. Any other bits have currently no meaning for the target.

st_uid st_gid st_rdev No valid meaning for the target. Transmitted unchanged.

st_atime st_mtime st_ctime These values have a host and file system dependent accuracy. Especially on
```

Windows hosts, the file system may not support exact timing values.

The target gets a struct stat of the above representation and is responsible for coercing it to the target representation before continuing.

Note that due to size differences between the host, target, and protocol representations of struct stat members, these members could eventually get truncated on the target.

struct timeval

The buffer of type struct timeval used by the File-I/O protocol is defined as follows:

```
struct timeval {
    time_t tv_sec; /* second */
    long tv_usec; /* microsecond */
};
```

The integral datatypes conform to the definitions given in the appropriate section (see [Integral Datatypes], page 445, for details) so this structure is of size 8 bytes.

D.14.9 Constants

The following values are used for the constants inside of the protocol. GDB and target are responsible for translating these values before and after the call as needed.

Open Flags

All values are given in hexadecimal representation.

O_RDONLY	0x0
O_WRONLY	0x1
O_RDWR	0x2
O_APPEND	8x0
O_CREAT	0x200
O_TRUNC	0x400
O EXCL	0x800

mode_t Values

All values are given in octal representation.

```
0100000
S_IFREG
S_IFDIR
               040000
S_IRUSR
                 0400
S_IWUSR
                 0200
S_IXUSR
                 0100
S_IRGRP
                  040
S_{IWGRP}
                  020
S_IXGRP
                  010
S_IROTH
                   04
S_IWOTH
                   02
S_IXOTH
                   01
```

Errno Values

All values are given in decimal representation.

EPERM	1
ENOENT	2
EINTR	4
EBADF	9
EACCES	13
EFAULT	14
EBUSY	16
EEXIST	17
ENODEV	19

ENOTDIR	20
EISDIR	21
EINVAL	22
ENFILE	23
EMFILE	24
EFBIG	27
ENOSPC	28
ESPIPE	29
EROFS	30
ENAMETOOLONG	91
EUNKNOWN	9999

EUNKNOWN is used as a fallback error value if a host system returns any error value not in the list of supported error numbers.

Lseek Flags

```
SEEK_SET 0
SEEK_CUR 1
SEEK_END 2
```

Limits

All values are given in decimal representation.

```
INT_MIN -2147483648

INT_MAX 2147483647

UINT_MAX 4294967295

LONG_MIN -9223372036854775808

LONG_MAX 9223372036854775807

ULONG_MAX 18446744073709551615
```

D.14.10 File-I/O Examples

Example sequence of a write call, file descriptor 3, buffer is at target address 0x1234, 6 bytes should be written:

```
<- Fwrite,3,1234,6
request memory read from target
-> m1234,6
<- XXXXXX
return "6 bytes written"
-> F6
```

Example sequence of a read call, file descriptor 3, buffer is at target address 0x1234, 6 bytes should be read:

```
<- Fread,3,1234,6
request memory write to target
-> X1234,6:XXXXXX
return "6 bytes read"
-> F6
```

Example sequence of a read call, call fails on the host due to invalid file descriptor (EBADF):

```
<- Fread,3,1234,6
-> F-1,9
```

Example sequence of a read call, user presses Ctrl-c before syscall on host is called:

```
<- Fread,3,1234,6
-> F-1,4,C
```

```
<- T02
```

Example sequence of a read call, user presses Ctrl-c after syscall on host is called:

```
<- Fread,3,1234,6
-> X1234,6:XXXXXX
<- T02
```

D.15 Library List Format

On some platforms, a dynamic loader (e.g. 'ld.so') runs in the same process as your application to manage libraries. In this case, GDB can use the loader's symbol table and normal memory operations to maintain a list of shared libraries. On other platforms, the operating system manages loaded libraries. GDB can not retrieve the list of currently loaded libraries through memory operations, so it uses the 'qXfer:libraries:read' packet (see [qXfer library list read], page 425) instead. The remote stub queries the target's operating system and reports which libraries are loaded.

The 'qXfer:libraries:read' packet returns an XML document which lists loaded libraries and their offsets. Each library has an associated name and one or more segment or section base addresses, which report where the library was loaded in memory.

For the common case of libraries that are fully linked binaries, the library should have a list of segments. If the target supports dynamic linking of a relocatable object file, its library XML element should instead include a list of allocated sections. The segment or section bases are start addresses, not relocation offsets; they do not depend on the library's link-time base addresses.

GDB must be linked with the Expat library to support XML library lists. See [Expat], page 391.

A simple memory map, with one loaded library relocated by a single offset, looks like this:

```
library-list>
    library name="/lib/libc.so.6">
        <segment address="0x10000000"/>
        </library>
</library-list>
```

Another simple memory map, with one loaded library with three allocated sections (.text, .data, .bss), looks like this:

The format of a library list is described by this DTD:

```
<!-- library-list: Root element with versioning -->
<!ELEMENT library-list (library)*>
<!ATTLIST library-list version CDATA #FIXED "1.0">
<!ELEMENT library (segment*, section*)>
<!ATTLIST library name CDATA #REQUIRED>
<!ELEMENT segment EMPTY>
<!ATTLIST segment address CDATA #REQUIRED>
```

In addition, segments and section descriptors cannot be mixed within a single library element, and you must supply at least one segment or section for each library.

D.16 Memory Map Format

To be able to write into flash memory, GDB needs to obtain a memory map from the target. This section describes the format of the memory map.

The memory map is obtained using the 'qXfer:memory-map:read' (see [qXfer memory map read], page 425) packet and is an XML document that lists memory regions.

GDB must be linked with the Expat library to support XML memory maps. See [Expat], page 391.

The top-level structure of the document is shown below:

Each region can be either:

• A region of RAM starting at addr and extending for length bytes from there:

```
<memory type="ram" start="addr" length="length"/>
```

• A region of read-only memory:

```
<memory type="rom" start="addr" length="length"/>
```

• A region of flash memory, with erasure blocks blocksize bytes in length:

```
<memory type="flash" start="addr" length="length">
  cyproperty name="blocksize">blocksize
```

Regions must not overlap. GDB assumes that areas of memory not covered by the memory map are RAM, and uses the ordinary 'M' and 'X' packets to write to addresses in such ranges.

The formal DTD for memory map format is given below:

```
<!-- -->
<!-- Memory Map XML DTD ..... -->
<!-- File: memory-map.dtd ..... -->
<!-- ..... -->
<!-- memory-map.dtd -->
<!-- memory-map: Root element with versioning -->
<!ELEMENT memory-map (memory | property)>
<!ATTLIST memory-map
                 version CDATA #FIXED "1.0.0">
<!ELEMENT memory (property)>
<!-- memory: Specifies a memory region,
         and its type, or device. -->
<!ATTLIST memory
                  type
                       CDATA #REQUIRED
                  start CDATA #REQUIRED
                  length CDATA
                  device CDATA
<!-- property: Generic attribute tag -->
<!ELEMENT property (#PCDATA | property)*>
<!ATTLIST property
                 name CDATA #REQUIRED>
```

D.17 Thread List Format

To efficiently update the list of threads and their attributes, GDB issues the 'qXfer:threads:read' packet (see [qXfer threads read], page 426) and obtains the XML document with the following structure:

Each 'thread' element must have the 'id' attribute that identifies the thread (see [thread-id syntax], page 405). The 'core' attribute, if present, specifies which processor core the thread was last executing on. The content of the of 'thread' element is interpreted as human-readable auxilliary information.

Appendix E The GDB Agent Expression Mechanism

In some applications, it is not feasible for the debugger to interrupt the program's execution long enough for the developer to learn anything helpful about its behavior. If the program's correctness depends on its real-time behavior, delays introduced by a debugger might cause the program to fail, even when the code itself is correct. It is useful to be able to observe the program's behavior without interrupting it.

Using GDB's trace and collect commands, the user can specify locations in the program, and arbitrary expressions to evaluate when those locations are reached. Later, using the tfind command, she can examine the values those expressions had when the program hit the trace points. The expressions may also denote objects in memory — structures or arrays, for example — whose values GDB should record; while visiting a particular tracepoint, the user may inspect those objects as if they were in memory at that moment. However, because GDB records these values without interacting with the user, it can do so quickly and unobtrusively, hopefully not disturbing the program's behavior.

When GDB is debugging a remote target, the GDB agent code running on the target computes the values of the expressions itself. To avoid having a full symbolic expression evaluator on the agent, GDB translates expressions in the source language into a simpler bytecode language, and then sends the bytecode to the agent; the agent then executes the bytecode, and records the values for GDB to retrieve later.

The bytecode language is simple; there are forty-odd opcodes, the bulk of which are the usual vocabulary of C operands (addition, subtraction, shifts, and so on) and various sizes of literals and memory reference operations. The bytecode interpreter operates strictly on machine-level values — various sizes of integers and floating point numbers — and requires no information about types or symbols; thus, the interpreter's internal data structures are simple, and each bytecode requires only a few native machine instructions to implement it. The interpreter is small, and strict limits on the memory and time required to evaluate an expression are easy to determine, making it suitable for use by the debugging agent in real-time applications.

E.1 General Bytecode Design

The agent represents bytecode expressions as an array of bytes. Each instruction is one byte long (thus the term *bytecode*). Some instructions are followed by operand bytes; for example, the goto instruction is followed by a destination for the jump.

The bytecode interpreter is a stack-based machine; most instructions pop their operands off the stack, perform some operation, and push the result back on the stack for the next instruction to consume. Each element of the stack may contain either a integer or a floating point value; these values are as many bits wide as the largest integer that can be directly manipulated in the source language. Stack elements carry no record of their type; bytecode could push a value as an integer, then pop it as a floating point value. However, GDB will not generate code which does this. In C, one might define the type of a stack element as follows:

```
union agent_val {
  LONGEST 1;
  DOUBLEST d;
```

Debugging with GDB

};

where LONGEST and DOUBLEST are typedef names for the largest integer and floating point types on the machine.

By the time the bytecode interpreter reaches the end of the expression, the value of the expression should be the only value left on the stack. For tracing applications, trace bytecodes in the expression will have recorded the necessary data, and the value on the stack may be discarded. For other applications, like conditional breakpoints, the value may be useful.

Separate from the stack, the interpreter has two registers:

The address of the next bytecode to execute.

The address of the start of the bytecode expression, necessary for interpreting the goto and if_goto instructions.

Neither of these registers is directly visible to the bytecode language itself, but they are useful for defining the meanings of the bytecode operations.

There are no instructions to perform side effects on the running program, or call the program's functions; we assume that these expressions are only used for unobtrusive debugging, not for patching the running code.

Most bytecode instructions do not distinguish between the various sizes of values, and operate on full-width values; the upper bits of the values are simply ignored, since they do not usually make a difference to the value computed. The exceptions to this rule are:

memory reference instructions (refn)

There are distinct instructions to fetch different word sizes from memory. Once on the stack, however, the values are treated as full-size integers. They may need to be sign-extended; the ext instruction exists for this purpose.

the sign-extension instruction (ext n)

These clearly need to know which portion of their operand is to be extended to occupy the full length of the word.

If the interpreter is unable to evaluate an expression completely for some reason (a memory location is inaccessible, or a divisor is zero, for example), we say that interpretation "terminates with an error". This means that the problem is reported back to the interpreter's caller in some helpful way. In general, code using agent expressions should assume that they may attempt to divide by zero, fetch arbitrary memory locations, and misbehave in other ways.

Even complicated C expressions compile to a few bytecode instructions; for example, the expression x + y * z would typically produce code like the following, assuming that x and y live in registers, and z is a global variable holding a 32-bit int:

```
reg 1
reg 2
const32 address of z
ref32
ext 32
mul
add
```

end

In detail, these mean:

reg 1 Push the value of register 1 (presumably holding x) onto the stack.

reg 2 Push the value of register 2 (holding y).

const32 address of z

Push the address of **z** onto the stack.

Fetch a 32-bit word from the address at the top of the stack; replace the address on the stack with the value. Thus, we replace the address of z with z's value.

ext 32 Sign-extend the value on the top of the stack from 32 bits to full length. This is necessary because **z** is a signed integer.

Pop the top two numbers on the stack, multiply them, and push their product.

Now the top of the stack contains the value of the expression y * z.

add Pop the top two numbers, add them, and push the sum. Now the top of the stack contains the value of x + y * z.

end Stop executing; the value left on the stack top is the value to be recorded.

E.2 Bytecode Descriptions

Each bytecode description has the following form:

```
add (0x02): a b \Rightarrow a+b
```

Pop the top two stack items, a and b, as integers; push their sum, as an integer.

In this example, add is the name of the bytecode, and (0x02) is the one-byte value used to encode the bytecode, in hexadecimal. The phrase " $a \ b \Rightarrow a+b$ " shows the stack before and after the bytecode executes. Beforehand, the stack must contain at least two values, a and b; since the top of the stack is to the right, b is on the top of the stack, and a is underneath it. After execution, the bytecode will have popped a and b from the stack, and replaced them with a single value, a+b. There may be other values on the stack below those shown, but the bytecode affects only those shown.

Here is another example:

```
const8 (0x22) n: \Rightarrow n
```

Push the 8-bit integer constant n on the stack, without sign extension.

In this example, the bytecode const8 takes an operand n directly from the bytecode stream; the operand follows the const8 bytecode itself. We write any such operands immediately after the name of the bytecode, before the colon, and describe the exact encoding of the operand in the bytecode stream in the body of the bytecode description.

For the const8 bytecode, there are no stack items given before the \Rightarrow ; this simply means that the bytecode consumes no values from the stack. If a bytecode consumes no values, or produces no values, the list on either side of the \Rightarrow may be empty.

If a value is written as a, b, or n, then the bytecode treats it as an integer. If a value is written is addr, then the bytecode treats it as an address.

We do not fully describe the floating point operations here; although this design can be extended in a clean way to handle floating point values, they are not of immediate interest to the customer, so we avoid describing them, to save time.

float (0x01): \Rightarrow

Prefix for floating-point bytecodes. Not implemented yet.

add (0x02): $a b \Rightarrow a+b$

Pop two integers from the stack, and push their sum, as an integer.

sub (0x03): $ab \Rightarrow a-b$

Pop two integers from the stack, subtract the top value from the next-to-top value, and push the difference.

mul(0x04): $a b \Rightarrow a*b$

Pop two integers from the stack, multiply them, and push the product on the stack. Note that, when one multiplies two n-bit numbers yielding another n-bit number, it is irrelevant whether the numbers are signed or not; the results are the same.

div_signed (0x05): $a \ b \Rightarrow a/b$

Pop two signed integers from the stack; divide the next-to-top value by the top value, and push the quotient. If the divisor is zero, terminate with an error.

div_unsigned (0x06): $a \ b \Rightarrow a/b$

Pop two unsigned integers from the stack; divide the next-to-top value by the top value, and push the quotient. If the divisor is zero, terminate with an error.

rem_signed (0x07): $a b \Rightarrow a modulo b$

Pop two signed integers from the stack; divide the next-to-top value by the top value, and push the remainder. If the divisor is zero, terminate with an error.

rem_unsigned (0x08): $a b \Rightarrow a modulo b$

Pop two unsigned integers from the stack; divide the next-to-top value by the top value, and push the remainder. If the divisor is zero, terminate with an error.

lsh(0x09): $ab \Rightarrow a << b$

Pop two integers from the stack; let a be the next-to-top value, and b be the top value. Shift a left by b bits, and push the result.

 $rsh_signed (0x0a): a b \Rightarrow (signed) a >> b$

Pop two integers from the stack; let a be the next-to-top value, and b be the top value. Shift a right by b bits, inserting copies of the top bit at the high end, and push the result.

rsh_unsigned (0x0b): $a b \Rightarrow a >> b$

Pop two integers from the stack; let a be the next-to-top value, and b be the top value. Shift a right by b bits, inserting zero bits at the high end, and push the result.

 $log_not (0x0e): a \Rightarrow !a$

Pop an integer from the stack; if it is zero, push the value one; otherwise, push the value zero.

bit_and (0x0f): $a b \Rightarrow a\&b$

Pop two integers from the stack, and push their bitwise and.

bit_or (0x10): $a b \Rightarrow a \mid b$

Pop two integers from the stack, and push their bitwise or.

bit_xor (0x11): $a b \Rightarrow a^b$

Pop two integers from the stack, and push their bitwise exclusive-or.

bit_not (0x12): $a \Rightarrow a$

Pop an integer from the stack, and push its bitwise complement.

equal (0x13): $a b \Rightarrow a=b$

Pop two integers from the stack; if they are equal, push the value one; otherwise, push the value zero.

less_signed (0x14): $a b \Rightarrow a < b$

Pop two signed integers from the stack; if the next-to-top value is less than the top value, push the value one; otherwise, push the value zero.

less_unsigned (0x15): $a \ b \Rightarrow a < b$

Pop two unsigned integers from the stack; if the next-to-top value is less than the top value, push the value one; otherwise, push the value zero.

ext (0x16) n: $a \Rightarrow a$, sign-extended from n bits

Pop an unsigned value from the stack; treating it as an n-bit twos-complement value, extend it to full length. This means that all bits to the left of bit n-1 (where the least significant bit is bit 0) are set to the value of bit n-1. Note that n may be larger than or equal to the width of the stack elements of the bytecode engine; in this case, the bytecode should have no effect.

The number of source bits to preserve, n, is encoded as a single byte unsigned integer following the ext bytecode.

zero_ext (0x2a) n: $a \Rightarrow a$, zero-extended from n bits

Pop an unsigned value from the stack; zero all but the bottom n bits. This means that all bits to the left of bit n-1 (where the least significant bit is bit 0) are set to the value of bit n-1.

The number of source bits to preserve, n, is encoded as a single byte unsigned integer following the ${\tt zero_ext}$ bytecode.

ref8 (0x17): $addr \Rightarrow a$ ref16 (0x18): $addr \Rightarrow a$ ref32 (0x19): $addr \Rightarrow a$ ref64 (0x1a): $addr \Rightarrow a$

Pop an address addr from the stack. For bytecode **ref**n, fetch an n-bit value from addr, using the natural target endianness. Push the fetched value as an unsigned integer.

Note that addr may not be aligned in any particular way; the refn bytecodes should operate correctly for any address.

If attempting to access memory at addr would cause a processor exception of some sort, terminate with an error.

 $if_goto (0x20) offset: a \Rightarrow$

Pop an integer off the stack; if it is non-zero, branch to the given offset in the bytecode string. Otherwise, continue to the next instruction in the bytecode stream. In other words, if a is non-zero, set the pc register to start + offset. Thus, an offset of zero denotes the beginning of the expression.

The offset is stored as a sixteen-bit unsigned value, stored immediately following the if_goto bytecode. It is always stored most significant byte first, regardless of the target's normal endianness. The offset is not guaranteed to fall at any particular alignment within the bytecode stream; thus, on machines where fetching a 16-bit on an unaligned address raises an exception, you should fetch the offset one byte at a time.

```
goto (0x21) offset: \Rightarrow
```

Branch unconditionally to offset; in other words, set the pc register to start + offset.

The offset is stored in the same way as for the if_goto bytecode.

```
const8 (0x22) n: \Rightarrow n
const16 (0x23) n: \Rightarrow n
const32 (0x24) n: \Rightarrow n
const64 (0x25) n: \Rightarrow n
```

Push the integer constant n on the stack, without sign extension. To produce a small negative value, push a small twos-complement value, and then sign-extend it using the ext bytecode.

The constant n is stored in the appropriate number of bytes following the $\mathtt{const}b$ bytecode. The constant n is always stored most significant byte first, regardless of the target's normal endianness. The constant is not guaranteed to fall at any particular alignment within the bytecode stream; thus, on machines where fetching a 16-bit on an unaligned address raises an exception, you should fetch n one byte at a time.

```
reg (0x26) n: \Rightarrow a
```

Push the value of register number n, without sign extension. The registers are numbered following GDB's conventions.

The register number n is encoded as a 16-bit unsigned integer immediately following the **reg** bytecode. It is always stored most significant byte first, regardless of the target's normal endianness. The register number is not guaranteed to fall at any particular alignment within the bytecode stream; thus, on machines where fetching a 16-bit on an unaligned address raises an exception, you should fetch the register number one byte at a time.

 $getv (0x2c) n: \Rightarrow v$

Push the value of trace state variable number n, without sign extension.

The variable number n is encoded as a 16-bit unsigned integer immediately following the <code>getv</code> bytecode. It is always stored most significant byte first, regardless of the target's normal endianness. The variable number is not guaranteed to fall at any particular alignment within the bytecode stream; thus, on machines where fetching a 16-bit on an unaligned address raises an exception, you should fetch the register number one byte at a time.

 $setv (0x2d) n: \Rightarrow v$

Set trace state variable number n to the value found on the top of the stack. The stack is unchanged, so that the value is readily available if the assignment is part of a larger expression. The handling of n is as described for getv.

trace (0x0c): addr size \Rightarrow

Record the contents of the *size* bytes at *addr* in a trace buffer, for later retrieval by GDB.

 $trace_quick (0x0d) size: addr \Rightarrow addr$

Record the contents of the *size* bytes at *addr* in a trace buffer, for later retrieval by GDB. *size* is a single byte unsigned integer following the trace opcode.

This bytecode is equivalent to the sequence dup const8 size trace, but we provide it anyway to save space in bytecode strings.

trace16 (0x30) size: $addr \Rightarrow addr$

Identical to trace_quick, except that size is a 16-bit big-endian unsigned integer, not a single byte. This should probably have been named trace_quick16, for consistency.

tracev (0x2e) $n: \Rightarrow a$

Record the value of trace state variable number n in the trace buffer. The handling of n is as described for getv.

end (0x27): \Rightarrow

Stop executing bytecode; the result should be the top element of the stack. If the purpose of the expression was to compute an lvalue or a range of memory, then the next-to-top of the stack is the lvalue's address, and the top of the stack is the lvalue's size, in bytes.

E.3 Using Agent Expressions

Agent expressions can be used in several different ways by GDB, and the debugger can generate different bytecode sequences as appropriate.

One possibility is to do expression evaluation on the target rather than the host, such as for the conditional of a conditional tracepoint. In such a case, GDB compiles the source expression into a bytecode sequence that simply gets values from registers or memory, does arithmetic, and returns a result.

Another way to use agent expressions is for tracepoint data collection. GDB generates a different bytecode sequence for collection; in addition to bytecodes that do the calculation, GDB adds trace bytecodes to save the pieces of memory that were used.

- The user selects trace points in the program's code at which GDB should collect data.
- The user specifies expressions to evaluate at each trace point. These expressions may denote objects in memory, in which case those objects' contents are recorded as the program runs, or computed values, in which case the values themselves are recorded.
- GDB transmits the tracepoints and their associated expressions to the GDB agent, running on the debugging target.
- The agent arranges to be notified when a trace point is hit.
- When execution on the target reaches a trace point, the agent evaluates the expressions associated with that trace point, and records the resulting values and memory ranges.
- Later, when the user selects a given trace event and inspects the objects and expression values recorded, GDB talks to the agent to retrieve recorded data as necessary to meet the user's requests. If the user asks to see an object whose contents have not been recorded, GDB reports an error.

E.4 Varying Target Capabilities

Some targets don't support floating-point, and some would rather not have to deal with long long operations. Also, different targets will have different stack sizes, and different bytecode buffer lengths.

Thus, GDB needs a way to ask the target about itself. We haven't worked out the details yet, but in general, GDB should be able to send the target a packet asking it to describe itself. The reply should be a packet whose length is explicit, so we can add new information to the packet in future revisions of the agent, without confusing old versions of GDB, and it should contain a version number. It should contain at least the following information:

- whether floating point is supported
- whether long long is supported
- maximum acceptable size of bytecode stack
- maximum acceptable length of bytecode expressions
- which registers are actually available for collection
- whether the target supports disabled tracepoints

E.5 Rationale

Some of the design decisions apparent above are arguable.

What about stack overflow/underflow?

GDB should be able to query the target to discover its stack size. Given that information, GDB can determine at translation time whether a given expression

will overflow the stack. But this spec isn't about what kinds of error-checking GDB ought to do.

Why are you doing everything in LONGEST?

Speed isn't important, but agent code size is; using LONGEST brings in a bunch of support code to do things like division, etc. So this is a serious concern.

First, note that you don't need different bytecodes for different operand sizes. You can generate code without *knowing* how big the stack elements actually are on the target. If the target only supports 32-bit ints, and you don't send any 64-bit bytecodes, everything just works. The observation here is that the MIPS and the Alpha have only fixed-size registers, and you can still get C's semantics even though most instructions only operate on full-sized words. You just need to make sure everything is properly sign-extended at the right times. So there is no need for 32- and 64-bit variants of the bytecodes. Just implement everything using the largest size you support.

GDB should certainly check to see what sizes the target supports, so the user can get an error earlier, rather than later. But this information is not necessary for correctness.

Why don't you have > or <= operators?

I want to keep the interpreter small, and we don't need them. We can combine the $less_$ opcodes with $log_$ not, and swap the order of the operands, yielding all four asymmetrical comparison operators. For example, $(x \le y)$ is ! (x > y), which is ! (y < x).

Why do you have log_not?

Why do you have ext?

Why do you have zero_ext?

These are all easily synthesized from other instructions, but I expect them to be used frequently, and they're simple, so I include them to keep bytecode strings short.

log_not is equivalent to const8 0 equal; it's used in half the relational operators.

ext n is equivalent to const8 s-n 1sh const8 s-n rsh_signed, where s is the size of the stack elements; it follows refm and reg bytecodes when the value should be signed. See the next bulleted item.

zero_ext n is equivalent to constm mask log_and; it's used whenever we push the value of a register, because we can't assume the upper bits of the register aren't garbage.

Why not have sign-extending variants of the ref operators?

Because that would double the number of ref operators, and we need the ext bytecode anyway for accessing bitfields.

Why not have constant-address variants of the ref operators?

Because that would double the number of ref operators again, and const32 address ref32 is only one byte longer.

Why do the refn operators have to support unaligned fetches?

GDB will generate bytecode that fetches multi-byte values at unaligned addresses whenever the executable's debugging information tells it to. Furthermore, GDB does not know the value the pointer will have when GDB generates the bytecode, so it cannot determine whether a particular fetch will be aligned or not.

In particular, structure bitfields may be several bytes long, but follow no alignment rules; members of packed structures are not necessarily aligned either.

In general, there are many cases where unaligned references occur in correct C code, either at the programmer's explicit request, or at the compiler's discretion. Thus, it is simpler to make the GDB agent bytecodes work correctly in all circumstances than to make GDB guess in each case whether the compiler did the usual thing.

Why are there no side-effecting operators?

Because our current client doesn't want them? That's a cheap answer. I think the real answer is that I'm afraid of implementing function calls. We should re-visit this issue after the present contract is delivered.

Why aren't the goto ops PC-relative?

The interpreter has the base address around anyway for PC bounds checking, and it seemed simpler.

Why is there only one offset size for the goto ops?

Offsets are currently sixteen bits. I'm not happy with this situation either:

Suppose we have multiple branch ops with different offset sizes. As I generate code left-to-right, all my jumps are forward jumps (there are no loops in expressions), so I never know the target when I emit the jump opcode. Thus, I have to either always assume the largest offset size, or do jump relaxation on the code after I generate it, which seems like a big waste of time.

I can imagine a reasonable expression being longer than 256 bytes. I can't imagine one being longer than 64k. Thus, we need 16-bit offsets. This kind of reasoning is so bogus, but relaxation is pathetic.

The other approach would be to generate code right-to-left. Then I'd always know my offset size. That might be fun.

Where is the function call bytecode?

When we add side-effects, we should add this.

Why does the reg bytecode take a 16-bit register number?

Intel's IA-64 architecture has 128 general-purpose registers, and 128 floating-point registers, and I'm sure it has some random control registers.

Why do we need trace and trace_quick?

Because GDB needs to record all the memory contents and registers an expression touches. If the user wants to evaluate an expression x-y-z, the agent must record the values of z and z-y as well as the value of z-z-z.

Don't the trace bytecodes make the interpreter less general?

They do mean that the interpreter contains special-purpose code, but that doesn't mean the interpreter can only be used for that purpose. If an expression doesn't use the trace bytecodes, they don't get in its way.

Why doesn't trace_quick consume its arguments the way everything else does?

In general, you do want your operators to consume their arguments; it's consistent, and generally reduces the amount of stack rearrangement necessary. However, trace_quick is a kludge to save space; it only exists so we needn't write dup const8 SIZE trace before every memory reference. Therefore, it's okay for it not to consume its arguments; it's meant for a specific context in which we know exactly what it should do with the stack. If we're going to have a kludge, it should be an effective kludge.

Why does trace16 exist?

That opcode was added by the customer that contracted Cygnus for the data tracing work. I personally think it is unnecessary; objects that large will be quite rare, so it is okay to use dup const16 size trace in those cases.

Whatever we decide to do with trace16, we should at least leave opcode 0x30 reserved, to remain compatible with the customer who added it.

Appendix F Trace File Format

The trace file comes in three parts: a header, a textual description section, and a trace frame section with binary data.

The header has the form $\x7fTRACE0\n$. The first byte is 0x7f so as to indicate that the file contains binary data, while the 0 is a version number that may have different values in the future.

The description section consists of multiple lines of ASCII text separated by newline characters (0xa). The lines may include a variety of optional descriptive or context-setting information, such as tracepoint definitions or register set size. GDB will ignore any line that it does not recognize. An empty line marks the end of this section.

The trace frame section consists of a number of consecutive frames. Each frame begins with a two-byte tracepoint number, followed by a four-byte size giving the amount of data in the frame. The data in the frame consists of a number of blocks, each introduced by a character indicating its type (at least register, memory, and trace state variable). The data in this section is raw binary, not a hexadecimal or other encoding; its endianness matches the target's endianness.

R bytes Register block. The number and ordering of bytes matches that of a g packet in the remote protocol. Note that these are the actual bytes, in target order and GDB register order, not a hexadecimal encoding.

M address length bytes...

Memory block. This is a contiguous block of memory, at the 8-byte address address, with a 2-byte length length, followed by length bytes.

V number value

Trace state variable block. This records the 8-byte signed value value of trace state variable numbered number.

Future enhancements of the trace file format may include additional types of blocks.

Appendix G Target Descriptions

Warning: target descriptions are still under active development, and the contents and format may change between GDB releases. The format is expected to stabilize in the future.

One of the challenges of using GDB to debug embedded systems is that there are so many minor variants of each processor architecture in use. It is common practice for vendors to start with a standard processor core — ARM, PowerPC, or MIPS, for example — and then make changes to adapt it to a particular market niche. Some architectures have hundreds of variants, available from dozens of vendors. This leads to a number of problems:

- With so many different customized processors, it is difficult for the GDB maintainers to keep up with the changes.
- Since individual variants may have short lifetimes or limited audiences, it may not be worthwhile to carry information about every variant in the GDB source tree.
- When gdb does support the architecture of the embedded system at hand, the task of finding the correct architecture name to give the set architecture command can be error-prone.

To address these problems, the GDB remote protocol allows a target system to not only identify itself to GDB, but to actually describe its own features. This lets GDB support processor variants it has never seen before — to the extent that the descriptions are accurate, and that GDB understands them.

GDB must be linked with the Expat library to support XML target descriptions. See [Expat], page 391.

G.1 Retrieving Descriptions

Target descriptions can be read from the target automatically, or specified by the user manually. The default behavior is to read the description from the target. GDB retrieves it via the remote protocol using 'qXfer' requests (see Section D.5 [General Query Packets], page 415). The annex in the 'qXfer' packet will be 'target.xml'. The contents of the 'target.xml' annex are an XML document, of the form described in Section G.2 [Target Description Format], page 468.

Alternatively, you can specify a file to read for the target description. If a file is set, the target will not be queried. The commands to specify a file are:

set tdesc filename path

Read the target description from path.

unset tdesc filename

Do not read the XML target description from a file. GDB will use the description supplied by the current target.

show tdesc filename

Show the filename to read for a target description, if any.

G.2 Target Description Format

A target description annex is an XML document which complies with the Document Type Definition provided in the GDB sources in 'gdb/features/gdb-target.dtd'. This means you can use generally available tools like xmllint to check that your feature descriptions are well-formed and valid. However, to help people unfamiliar with XML write descriptions for their targets, we also describe the grammar here.

Target descriptions can identify the architecture of the remote target and (for some architectures) provide information about custom register sets. They can also identify the OS ABI of the remote target. GDB can use this information to autoconfigure for your target, or to warn you if you connect to an unsupported target.

Here is a simple target description:

```
<target version="1.0">
  <architecture>i386:x86-64</architecture>
</target>
```

This minimal description only says that the target uses the x86-64 architecture.

A target description has the following overall form, with [] marking optional elements and ... marking repeatable elements. The elements are explained further below.

```
<?xml version="1.0"?>
<!DOCTYPE target SYSTEM "gdb-target.dtd">
<target version="1.0">
    [architecture]
    [osabi]
    [compatible]
    [feature...]
```

The description is generally insensitive to whitespace and line breaks, under the usual common-sense rules. The XML version declaration and document type declaration can generally be omitted (GDB does not require them), but specifying them may be useful for XML validation tools. The 'version' attribute for '<target>' may also be omitted, but we recommend including it; if future versions of GDB use an incompatible revision of 'gdb-target.dtd', they will detect and report the version mismatch.

G.2.1 Inclusion

It can sometimes be valuable to split a target description up into several different annexes, either for organizational purposes, or to share files between different possible target descriptions. You can divide a description into multiple files by replacing any element of the target description with an inclusion directive of the form:

```
<xi:include href="document"/>
```

When GDB encounters an element of this form, it will retrieve the named XML document, and replace the inclusion directive with the contents of that document. If the current description was read using 'qXfer', then so will be the included document; document will be interpreted as the name of an annex. If the current description was read from a file, GDB will look for document as a file in the same directory where it found the original description.

G.2.2 Architecture

An '<architecture>' element has this form:

```
<architecture>arch</architecture>
```

arch is one of the architectures from the set accepted by set architecture (see Chapter 19 [Specifying a Debugging Target], page 199).

G.2.3 OS ABI

This optional field was introduced in gdb version 7.0. Previous versions of gdb ignore it.

An '<osabi>' element has this form:

```
<osabi>abi-name
```

abi-name is an OS ABI name from the same selection accepted by set osabi (see Section 22.6 [Configuring the Current ABI], page 248).

G.2.4 Compatible Architecture

This optional field was introduced in GDB version 7.0. Previous versions of GDB ignore it.

A '<compatible>' element has this form:

```
<compatible>arch</compatible>
```

arch is one of the architectures from the set accepted by set architecture (see Chapter 19 [Specifying a Debugging Target], page 199).

A '<compatible>' element is used to specify that the target is able to run binaries in some other than the main target architecture given by the '<architecture>' element. For example, on the Cell Broadband Engine, the main architecture is powerpc:common or powerpc:common64, but the system is able to run binaries in the spu architecture as well. The way to describe this capability with '<compatible>' is as follows:

```
<architecture>powerpc:common</architecture>
<compatible>spu</compatible>
```

G.2.5 Features

Each '<feature>' describes some logical portion of the target system. Features are currently used to describe available CPU registers and the types of their contents. A '<feature>' element has this form:

```
<freature name="name">
[type...]
reg...
</feature>
```

Each feature's name should be unique within the description. The name of a feature does not matter unless GDB has some special knowledge of the contents of that feature; if it does, the feature should have its standard name. See Section G.4 [Standard Target Features], page 471.

G.2.6 Types

Any register's value is a collection of bits which GDB must interpret. The default interpretation is a two's complement integer, but other types can be requested by name in the register description. Some predefined types are provided by GDB (see Section G.3 [Predefined Target Types], page 471), and the description can define additional composite types.

Each type element must have an 'id' attribute, which gives a unique (within the containing '<feature>') name to the type. Types must be defined before they are used.

Some targets offer vector registers, which can be treated as arrays of scalar elements. These types are written as '<vector>' elements, specifying the array element type, type, and the number of elements, count:

```
<vector id="id" type="type" count="count"/>
```

If a register's value is usefully viewed in multiple ways, define it with a union type containing the useful representations. The '<union>' element contains one or more '<field>' elements, each of which has a name and a type:

G.2.7 Registers

Each register is represented as an element with this form:

```
<reg name="name"
   bitsize="size"
   [regnum="num"]
   [save-restore="save-restore"]
   [type="type"]
   [group="group"]/>
```

The components are as follows:

name The register's name; it must be unique within the target description.

bitsize The register's size, in bits.

The register's number. If omitted, a register's number is one greater than that of the previous register (either in the current feature or in a preceeding feature); the first register in the target description defaults to zero. This register number is used to read or write the register; e.g. it is used in the remote p and P packets, and registers appear in the g and G packets in order of increasing register number.

save-restore

Whether the register should be preserved across inferior function calls; this must be either yes or no. The default is yes, which is appropriate for most registers except for some system control registers; this is not related to the target's ABI.

The type of the register. type may be a predefined type, a type defined in the current feature, or one of the special types int and float. int is an integer type of the correct size for bitsize, and float is a floating point type (in the architecture's normal floating point format) of the correct size for bitsize. The default is int.

group The register group to which this register belongs. group must be either general, float, or vector. If no group is specified, GDB will not display the register in info registers.

G.3 Predefined Target Types

Type definitions in the self-description can build up composite types from basic building blocks, but can not define fundamental types. Instead, standard identifiers are provided by GDB for the fundamental types. The currently supported types are:

int16 int32 int64 Signed integer types holding the specified number of bits. int128 uint8 uint16 uint32 uint64 uint128 Unsigned integer types holding the specified number of bits. code_ptr data_ptr Pointers to unspecified code and data. The program counter and any dedicated return address register may be marked as code pointers; printing a code pointer converts it into a symbolic address. The stack pointer and any dedicated address registers may be marked as data pointers. ieee_single Single precision IEEE floating point. ieee_double Double precision IEEE floating point.

arm_fpa_ext

int8

The 12-byte extended precision format used by ARM FPA registers.

i387_ext The 10-byte extended precision format used by x87 registers.

i386_eflags

32bit eflags register used by x86.

i386_mxcsr

32bit MXCSR register used by x86.

G.4 Standard Target Features

A target description must contain either no registers or all the target's registers. If the description contains no registers, then GDB will assume a default register layout, selected based on the architecture. If the description contains any registers, the default layout will not be used; the standard registers must be described in the target description, in such a way that GDB can recognize them.

This is accomplished by giving specific names to feature elements which contain standard registers. GDB will look for features with those names and verify that they contain the expected registers; if any known feature is missing required registers, or if any required feature is missing, GDB will reject the target description. You can add additional registers

to any of the standard features — GDB will display them just as if they were added to an unrecognized feature.

This section lists the known features and their expected contents. Sample XML documents for these features are included in the GDB source tree, in the directory 'gdb/features'.

Names recognized by GDB should include the name of the company or organization which selected the name, and the overall architecture to which the feature applies; so e.g. the feature containing ARM core registers is named 'org.gnu.gdb.arm.core'.

The names of registers are not case sensitive for the purpose of recognizing standard features, but gdb will only display registers using the capitalization used in the description.

G.4.1 ARM Features

The 'org.gnu.gdb.arm.core' feature is required for ARM targets. It should contain registers 'r0' through 'r13', 'sp', 'lr', 'pc', and 'cpsr'.

The 'org.gnu.gdb.arm.fpa' feature is optional. If present, it should contain registers 'f0' through 'f7' and 'fps'.

The 'org.gnu.gdb.xscale.iwmmxt' feature is optional. If present, it should contain at least registers 'wR0' through 'wR15' and 'wCGR0' through 'wCGR3'. The 'wCID', 'wCon', 'wCSSF', and 'wCASF' registers are optional.

The 'org.gnu.gdb.arm.vfp' feature is optional. If present, it should contain at least registers 'd0' through 'd15'. If they are present, 'd16' through 'd31' should also be included. GDB will synthesize the single-precision registers from halves of the double-precision registers.

The 'org.gnu.gdb.arm.neon' feature is optional. It does not need to contain registers; it instructs GDB to display the VFP double-precision registers as vectors and to synthesize the quad-precision registers from pairs of double-precision registers. If this feature is present, 'org.gnu.gdb.arm.vfp' must also be present and include 32 double-precision registers.

G.4.2 i386 Features

The 'org.gnu.gdb.i386.core' feature is required for i386/amd64 targets. It should describe the following registers:

- 'eax' through 'edi' plus 'eip' for i386
- 'rax' through 'r15' plus 'rip' for amd64
- 'eflags', 'cs', 'ss', 'ds', 'es', 'fs', 'gs'
- 'st0' through 'st7'
- 'fctrl', 'fstat', 'ftag', 'fiseg', 'fioff', 'foseg', 'fooff' and 'fop'

The register sets may be different, depending on the target.

The 'org.gnu.gdb.i386.sse' feature is required. It should describe registers:

- 'xmm0' through 'xmm7' for i386
- 'xmm0' through 'xmm15' for amd64
- 'mxcsr'

The 'org.gnu.gdb.i386.linux' feature is optional. It should describe a single register, 'orig_eax'.

G.4.3 MIPS Features

The 'org.gnu.gdb.mips.cpu' feature is required for MIPS targets. It should contain registers 'r0' through 'r31', 'lo', 'hi', and 'pc'. They may be 32-bit or 64-bit depending on the target.

The 'org.gnu.gdb.mips.cp0' feature is also required. It should contain at least the 'status', 'badvaddr', and 'cause' registers. They may be 32-bit or 64-bit depending on the target.

The 'org.gnu.gdb.mips.fpu' feature is currently required, though it may be optional in a future version of GDB. It should contain registers 'f0' through 'f31', 'fcsr', and 'fir'. They may be 32-bit or 64-bit depending on the target.

The 'org.gnu.gdb.mips.linux' feature is optional. It should contain a single register, 'restart', which is used by the Linux kernel to control restartable syscalls.

G.4.4 M68K Features

```
'org.gnu.gdb.m68k.core'
```

One of those features must be always present. The feature that is present determines which flavor of m68k is used. The feature that is present should contain registers 'd0' through 'd7', 'a0' through 'a5', 'fp', 'sp', 'ps' and 'pc'.

'org.gnu.gdb.coldfire.fp'

This feature is optional. If present, it should contain registers 'fp0' through 'fp7', 'fpcontrol', 'fpstatus' and 'fpiaddr'.

G.4.5 PowerPC Features

The 'org.gnu.gdb.power.core' feature is required for PowerPC targets. It should contain registers 'r0' through 'r31', 'pc', 'msr', 'cr', 'lr', 'ctr', and 'xer'. They may be 32-bit or 64-bit depending on the target.

The 'org.gnu.gdb.power.fpu' feature is optional. It should contain registers 'f0' through 'f31' and 'fpscr'.

The 'org.gnu.gdb.power.altivec' feature is optional. It should contain registers 'vr0' through 'vr31', 'vscr', and 'vrsave'.

The 'org.gnu.gdb.power.vsx' feature is optional. It should contain registers 'vs0h' through 'vs31h'. GDB will combine these registers with the floating point registers ('f0' through 'f31') and the altivec registers ('vr0' through 'vr31') to present the 128-bit wide registers 'vs0' through 'vs63', the set of vector registers for POWER7.

The 'org.gnu.gdb.power.spe' feature is optional. It should contain registers 'ev0h' through 'ev31h', 'acc', and 'spefscr'. SPE targets should provide 32-bit registers in 'org.gnu.gdb.power.core' and provide the upper halves in 'ev0h' through 'ev31h'. GDB will combine these to present registers 'ev0' through 'ev31' to the user.

^{&#}x27;org.gnu.gdb.coldfire.core'

^{&#}x27;org.gnu.gdb.fido.core'

Appendix H Operating System Information

Users of GDB often wish to obtain information about the state of the operating system running on the target—for example the list of processes, or the list of open files. This section describes the mechanism that makes it possible. This mechanism is similar to the target features mechanism (see Appendix G [Target Descriptions], page 467), but focuses on a different aspect of target.

Operating system information is retrived from the target via the remote protocol, using 'qXfer' requests (see [qXfer osdata read], page 426). The object name in the request should be 'osdata', and the *annex* identifies the data to be fetched.

H.1 Process list

When requesting the process list, the annex field in the 'qXfer' request should be 'processes'. The returned data is an XML document. The formal syntax of this document is defined in 'gdb/features/osdata.dtd'.

```
An example document is:
```

Each item should include a column whose name is 'pid'. The value of that column should identify the process on the target. The 'user' and 'command' columns are optional, and will be displayed by GDB. The 'cores' column, if present, should contain a comma-separated list of cores that this process is running on. Target may provide additional columns, which GDB currently ignores.

Appendix I GNU GENERAL PUBLIC LICENSE

Version 3, 29 June 2007

Copyright © 2007 Free Software Foundation, Inc. http://fsf.org/

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The GNU General Public License is a free, copyleft license for software and other kinds of works.

The licenses for most software and other practical works are designed to take away your freedom to share and change the works. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change all versions of a program—to make sure it remains free software for all its users. We, the Free Software Foundation, use the GNU General Public License for most of our software; it applies also to any other work released this way by its authors. You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for them if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs, and that you know you can do these things.

To protect your rights, we need to prevent others from denying you these rights or asking you to surrender the rights. Therefore, you have certain responsibilities if you distribute copies of the software, or if you modify it: responsibilities to respect the freedom of others.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must pass on to the recipients the same freedoms that you received. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

Developers that use the GNU GPL protect your rights with two steps: (1) assert copyright on the software, and (2) offer you this License giving you legal permission to copy, distribute and/or modify it.

For the developers' and authors' protection, the GPL clearly explains that there is no warranty for this free software. For both users' and authors' sake, the GPL requires that modified versions be marked as changed, so that their problems will not be attributed erroneously to authors of previous versions.

Some devices are designed to deny users access to install or run modified versions of the software inside them, although the manufacturer can do so. This is fundamentally incompatible with the aim of protecting users' freedom to change the software. The systematic pattern of such abuse occurs in the area of products for individuals to use, which is precisely where it is most unacceptable. Therefore, we have designed this version of the GPL to prohibit the practice for those products. If such problems arise substantially in other domains, we stand ready to extend this provision to those domains in future versions of the GPL, as needed to protect the freedom of users.

Finally, every program is threatened constantly by software patents. States should not allow patents to restrict development and use of software on general-purpose computers, but in those that do, we wish to avoid the special danger that patents applied to a free program could make it effectively proprietary. To prevent this, the GPL assures that patents cannot be used to render the program non-free.

The precise terms and conditions for copying, distribution and modification follow.

TERMS AND CONDITIONS

0. Definitions.

"This License" refers to version 3 of the GNU General Public License.

"Copyright" also means copyright-like laws that apply to other kinds of works, such as semiconductor masks.

"The Program" refers to any copyrightable work licensed under this License. Each licensee is addressed as "you". "Licensees" and "recipients" may be individuals or organizations.

To "modify" a work means to copy from or adapt all or part of the work in a fashion requiring copyright permission, other than the making of an exact copy. The resulting work is called a "modified version" of the earlier work or a work "based on" the earlier work.

A "covered work" means either the unmodified Program or a work based on the Program.

To "propagate" a work means to do anything with it that, without permission, would make you directly or secondarily liable for infringement under applicable copyright law, except executing it on a computer or modifying a private copy. Propagation includes copying, distribution (with or without modification), making available to the public, and in some countries other activities as well.

To "convey" a work means any kind of propagation that enables other parties to make or receive copies. Mere interaction with a user through a computer network, with no transfer of a copy, is not conveying.

An interactive user interface displays "Appropriate Legal Notices" to the extent that it includes a convenient and prominently visible feature that (1) displays an appropriate copyright notice, and (2) tells the user that there is no warranty for the work (except to the extent that warranties are provided), that licensees may convey the work under this License, and how to view a copy of this License. If the interface presents a list of user commands or options, such as a menu, a prominent item in the list meets this criterion.

1. Source Code.

The "source code" for a work means the preferred form of the work for making modifications to it. "Object code" means any non-source form of a work.

A "Standard Interface" means an interface that either is an official standard defined by a recognized standards body, or, in the case of interfaces specified for a particular programming language, one that is widely used among developers working in that language. The "System Libraries" of an executable work include anything, other than the work as a whole, that (a) is included in the normal form of packaging a Major Component, but which is not part of that Major Component, and (b) serves only to enable use of the work with that Major Component, or to implement a Standard Interface for which an implementation is available to the public in source code form. A "Major Component", in this context, means a major essential component (kernel, window system, and so on) of the specific operating system (if any) on which the executable work runs, or a compiler used to produce the work, or an object code interpreter used to run it.

The "Corresponding Source" for a work in object code form means all the source code needed to generate, install, and (for an executable work) run the object code and to modify the work, including scripts to control those activities. However, it does not include the work's System Libraries, or general-purpose tools or generally available free programs which are used unmodified in performing those activities but which are not part of the work. For example, Corresponding Source includes interface definition files associated with source files for the work, and the source code for shared libraries and dynamically linked subprograms that the work is specifically designed to require, such as by intimate data communication or control flow between those subprograms and other parts of the work.

The Corresponding Source need not include anything that users can regenerate automatically from other parts of the Corresponding Source.

The Corresponding Source for a work in source code form is that same work.

2. Basic Permissions.

All rights granted under this License are granted for the term of copyright on the Program, and are irrevocable provided the stated conditions are met. This License explicitly affirms your unlimited permission to run the unmodified Program. The output from running a covered work is covered by this License only if the output, given its content, constitutes a covered work. This License acknowledges your rights of fair use or other equivalent, as provided by copyright law.

You may make, run and propagate covered works that you do not convey, without conditions so long as your license otherwise remains in force. You may convey covered works to others for the sole purpose of having them make modifications exclusively for you, or provide you with facilities for running those works, provided that you comply with the terms of this License in conveying all material for which you do not control copyright. Those thus making or running the covered works for you must do so exclusively on your behalf, under your direction and control, on terms that prohibit them from making any copies of your copyrighted material outside their relationship with you.

Conveying under any other circumstances is permitted solely under the conditions stated below. Sublicensing is not allowed; section 10 makes it unnecessary.

3. Protecting Users' Legal Rights From Anti-Circumvention Law.

No covered work shall be deemed part of an effective technological measure under any applicable law fulfilling obligations under article 11 of the WIPO copyright treaty adopted on 20 December 1996, or similar laws prohibiting or restricting circumvention of such measures.

When you convey a covered work, you waive any legal power to forbid circumvention of technological measures to the extent such circumvention is effected by exercising rights under this License with respect to the covered work, and you disclaim any intention to limit operation or modification of the work as a means of enforcing, against the work's users, your or third parties' legal rights to forbid circumvention of technological measures.

4. Conveying Verbatim Copies.

You may convey verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice; keep intact all notices stating that this License and any non-permissive terms added in accord with section 7 apply to the code; keep intact all notices of the absence of any warranty; and give all recipients a copy of this License along with the Program.

You may charge any price or no price for each copy that you convey, and you may offer support or warranty protection for a fee.

5. Conveying Modified Source Versions.

You may convey a work based on the Program, or the modifications to produce it from the Program, in the form of source code under the terms of section 4, provided that you also meet all of these conditions:

- a. The work must carry prominent notices stating that you modified it, and giving a relevant date.
- b. The work must carry prominent notices stating that it is released under this License and any conditions added under section 7. This requirement modifies the requirement in section 4 to "keep intact all notices".
- c. You must license the entire work, as a whole, under this License to anyone who comes into possession of a copy. This License will therefore apply, along with any applicable section 7 additional terms, to the whole of the work, and all its parts, regardless of how they are packaged. This License gives no permission to license the work in any other way, but it does not invalidate such permission if you have separately received it.
- d. If the work has interactive user interfaces, each must display Appropriate Legal Notices; however, if the Program has interactive interfaces that do not display Appropriate Legal Notices, your work need not make them do so.

A compilation of a covered work with other separate and independent works, which are not by their nature extensions of the covered work, and which are not combined with it such as to form a larger program, in or on a volume of a storage or distribution medium, is called an "aggregate" if the compilation and its resulting copyright are not used to limit the access or legal rights of the compilation's users beyond what the individual works permit. Inclusion of a covered work in an aggregate does not cause this License to apply to the other parts of the aggregate.

6. Conveying Non-Source Forms.

You may convey a covered work in object code form under the terms of sections 4 and 5, provided that you also convey the machine-readable Corresponding Source under the terms of this License, in one of these ways:

- a. Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by the Corresponding Source fixed on a durable physical medium customarily used for software interchange.
- b. Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by a written offer, valid for at least three years and valid for as long as you offer spare parts or customer support for that product model, to give anyone who possesses the object code either (1) a copy of the Corresponding Source for all the software in the product that is covered by this License, on a durable physical medium customarily used for software interchange, for a price no more than your reasonable cost of physically performing this conveying of source, or (2) access to copy the Corresponding Source from a network server at no charge.
- c. Convey individual copies of the object code with a copy of the written offer to provide the Corresponding Source. This alternative is allowed only occasionally and noncommercially, and only if you received the object code with such an offer, in accord with subsection 6b.
- d. Convey the object code by offering access from a designated place (gratis or for a charge), and offer equivalent access to the Corresponding Source in the same way through the same place at no further charge. You need not require recipients to copy the Corresponding Source along with the object code. If the place to copy the object code is a network server, the Corresponding Source may be on a different server (operated by you or a third party) that supports equivalent copying facilities, provided you maintain clear directions next to the object code saying where to find the Corresponding Source. Regardless of what server hosts the Corresponding Source, you remain obligated to ensure that it is available for as long as needed to satisfy these requirements.
- e. Convey the object code using peer-to-peer transmission, provided you inform other peers where the object code and Corresponding Source of the work are being offered to the general public at no charge under subsection 6d.

A separable portion of the object code, whose source code is excluded from the Corresponding Source as a System Library, need not be included in conveying the object code work.

A "User Product" is either (1) a "consumer product", which means any tangible personal property which is normally used for personal, family, or household purposes, or (2) anything designed or sold for incorporation into a dwelling. In determining whether a product is a consumer product, doubtful cases shall be resolved in favor of coverage. For a particular product received by a particular user, "normally used" refers to a typical or common use of that class of product, regardless of the status of the particular user or of the way in which the particular user actually uses, or expects or is expected to use, the product. A product is a consumer product regardless of whether the product has substantial commercial, industrial or non-consumer uses, unless such uses represent the only significant mode of use of the product.

"Installation Information" for a User Product means any methods, procedures, authorization keys, or other information required to install and execute modified versions of a covered work in that User Product from a modified version of its Corresponding Source.

The information must suffice to ensure that the continued functioning of the modified object code is in no case prevented or interfered with solely because modification has been made.

If you convey an object code work under this section in, or with, or specifically for use in, a User Product, and the conveying occurs as part of a transaction in which the right of possession and use of the User Product is transferred to the recipient in perpetuity or for a fixed term (regardless of how the transaction is characterized), the Corresponding Source conveyed under this section must be accompanied by the Installation Information. But this requirement does not apply if neither you nor any third party retains the ability to install modified object code on the User Product (for example, the work has been installed in ROM).

The requirement to provide Installation Information does not include a requirement to continue to provide support service, warranty, or updates for a work that has been modified or installed by the recipient, or for the User Product in which it has been modified or installed. Access to a network may be denied when the modification itself materially and adversely affects the operation of the network or violates the rules and protocols for communication across the network.

Corresponding Source conveyed, and Installation Information provided, in accord with this section must be in a format that is publicly documented (and with an implementation available to the public in source code form), and must require no special password or key for unpacking, reading or copying.

7. Additional Terms.

"Additional permissions" are terms that supplement the terms of this License by making exceptions from one or more of its conditions. Additional permissions that are applicable to the entire Program shall be treated as though they were included in this License, to the extent that they are valid under applicable law. If additional permissions apply only to part of the Program, that part may be used separately under those permissions, but the entire Program remains governed by this License without regard to the additional permissions.

When you convey a copy of a covered work, you may at your option remove any additional permissions from that copy, or from any part of it. (Additional permissions may be written to require their own removal in certain cases when you modify the work.) You may place additional permissions on material, added by you to a covered work, for which you have or can give appropriate copyright permission.

Notwithstanding any other provision of this License, for material you add to a covered work, you may (if authorized by the copyright holders of that material) supplement the terms of this License with terms:

- a. Disclaiming warranty or limiting liability differently from the terms of sections 15 and 16 of this License; or
- b. Requiring preservation of specified reasonable legal notices or author attributions in that material or in the Appropriate Legal Notices displayed by works containing it: or
- c. Prohibiting misrepresentation of the origin of that material, or requiring that modified versions of such material be marked in reasonable ways as different from the original version; or

- d. Limiting the use for publicity purposes of names of licensors or authors of the material; or
- e. Declining to grant rights under trademark law for use of some trade names, trademarks, or service marks; or
- f. Requiring indemnification of licensors and authors of that material by anyone who conveys the material (or modified versions of it) with contractual assumptions of liability to the recipient, for any liability that these contractual assumptions directly impose on those licensors and authors.

All other non-permissive additional terms are considered "further restrictions" within the meaning of section 10. If the Program as you received it, or any part of it, contains a notice stating that it is governed by this License along with a term that is a further restriction, you may remove that term. If a license document contains a further restriction but permits relicensing or conveying under this License, you may add to a covered work material governed by the terms of that license document, provided that the further restriction does not survive such relicensing or conveying.

If you add terms to a covered work in accord with this section, you must place, in the relevant source files, a statement of the additional terms that apply to those files, or a notice indicating where to find the applicable terms.

Additional terms, permissive or non-permissive, may be stated in the form of a separately written license, or stated as exceptions; the above requirements apply either way.

8. Termination.

You may not propagate or modify a covered work except as expressly provided under this License. Any attempt otherwise to propagate or modify it is void, and will automatically terminate your rights under this License (including any patent licenses granted under the third paragraph of section 11).

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, you do not qualify to receive new licenses for the same material under section 10.

9. Acceptance Not Required for Having Copies.

You are not required to accept this License in order to receive or run a copy of the Program. Ancillary propagation of a covered work occurring solely as a consequence of using peer-to-peer transmission to receive a copy likewise does not require acceptance.

However, nothing other than this License grants you permission to propagate or modify any covered work. These actions infringe copyright if you do not accept this License. Therefore, by modifying or propagating a covered work, you indicate your acceptance of this License to do so.

10. Automatic Licensing of Downstream Recipients.

Each time you convey a covered work, the recipient automatically receives a license from the original licensors, to run, modify and propagate that work, subject to this License. You are not responsible for enforcing compliance by third parties with this License.

An "entity transaction" is a transaction transferring control of an organization, or substantially all assets of one, or subdividing an organization, or merging organizations. If propagation of a covered work results from an entity transaction, each party to that transaction who receives a copy of the work also receives whatever licenses to the work the party's predecessor in interest had or could give under the previous paragraph, plus a right to possession of the Corresponding Source of the work from the predecessor in interest, if the predecessor has it or can get it with reasonable efforts.

You may not impose any further restrictions on the exercise of the rights granted or affirmed under this License. For example, you may not impose a license fee, royalty, or other charge for exercise of rights granted under this License, and you may not initiate litigation (including a cross-claim or counterclaim in a lawsuit) alleging that any patent claim is infringed by making, using, selling, offering for sale, or importing the Program or any portion of it.

11. Patents.

A "contributor" is a copyright holder who authorizes use under this License of the Program or a work on which the Program is based. The work thus licensed is called the contributor's "contributor version".

A contributor's "essential patent claims" are all patent claims owned or controlled by the contributor, whether already acquired or hereafter acquired, that would be infringed by some manner, permitted by this License, of making, using, or selling its contributor version, but do not include claims that would be infringed only as a consequence of further modification of the contributor version. For purposes of this definition, "control" includes the right to grant patent sublicenses in a manner consistent with the requirements of this License.

Each contributor grants you a non-exclusive, worldwide, royalty-free patent license under the contributor's essential patent claims, to make, use, sell, offer for sale, import and otherwise run, modify and propagate the contents of its contributor version.

In the following three paragraphs, a "patent license" is any express agreement or commitment, however denominated, not to enforce a patent (such as an express permission to practice a patent or covenant not to sue for patent infringement). To "grant" such a patent license to a party means to make such an agreement or commitment not to enforce a patent against the party.

If you convey a covered work, knowingly relying on a patent license, and the Corresponding Source of the work is not available for anyone to copy, free of charge and under the terms of this License, through a publicly available network server or other readily accessible means, then you must either (1) cause the Corresponding Source to be so

available, or (2) arrange to deprive yourself of the benefit of the patent license for this particular work, or (3) arrange, in a manner consistent with the requirements of this License, to extend the patent license to downstream recipients. "Knowingly relying" means you have actual knowledge that, but for the patent license, your conveying the covered work in a country, or your recipient's use of the covered work in a country, would infringe one or more identifiable patents in that country that you have reason to believe are valid.

If, pursuant to or in connection with a single transaction or arrangement, you convey, or propagate by procuring conveyance of, a covered work, and grant a patent license to some of the parties receiving the covered work authorizing them to use, propagate, modify or convey a specific copy of the covered work, then the patent license you grant is automatically extended to all recipients of the covered work and works based on it.

A patent license is "discriminatory" if it does not include within the scope of its coverage, prohibits the exercise of, or is conditioned on the non-exercise of one or more of the rights that are specifically granted under this License. You may not convey a covered work if you are a party to an arrangement with a third party that is in the business of distributing software, under which you make payment to the third party based on the extent of your activity of conveying the work, and under which the third party grants, to any of the parties who would receive the covered work from you, a discriminatory patent license (a) in connection with copies of the covered work conveyed by you (or copies made from those copies), or (b) primarily for and in connection with specific products or compilations that contain the covered work, unless you entered into that arrangement, or that patent license was granted, prior to 28 March 2007.

Nothing in this License shall be construed as excluding or limiting any implied license or other defenses to infringement that may otherwise be available to you under applicable patent law.

12. No Surrender of Others' Freedom.

If conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot convey a covered work so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not convey it at all. For example, if you agree to terms that obligate you to collect a royalty for further conveying from those to whom you convey the Program, the only way you could satisfy both those terms and this License would be to refrain entirely from conveying the Program.

13. Use with the GNU Affero General Public License.

Notwithstanding any other provision of this License, you have permission to link or combine any covered work with a work licensed under version 3 of the GNU Affero General Public License into a single combined work, and to convey the resulting work. The terms of this License will continue to apply to the part which is the covered work, but the special requirements of the GNU Affero General Public License, section 13, concerning interaction through a network will apply to the combination as such.

14. Revised Versions of this License.

The Free Software Foundation may publish revised and/or new versions of the GNU General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies that a certain numbered version of the GNU General Public License "or any later version" applies to it, you have the option of following the terms and conditions either of that numbered version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of the GNU General Public License, you may choose any version ever published by the Free Software Foundation.

If the Program specifies that a proxy can decide which future versions of the GNU General Public License can be used, that proxy's public statement of acceptance of a version permanently authorizes you to choose that version for the Program.

Later license versions may give you additional or different permissions. However, no additional obligations are imposed on any author or copyright holder as a result of your choosing to follow a later version.

15. Disclaimer of Warranty.

THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. Limitation of Liability.

IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MODIFIES AND/OR CONVEYS THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

17. Interpretation of Sections 15 and 16.

If the disclaimer of warranty and limitation of liability provided above cannot be given local legal effect according to their terms, reviewing courts shall apply local law that most closely approximates an absolute waiver of all civil liability in connection with the Program, unless a warranty or assumption of liability accompanies a copy of the Program in return for a fee.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively state the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

one line to give the program's name and a brief idea of what it does. Copyright (C) year name of author

This program is free software: you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation, either version 3 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program. If not, see http://www.gnu.org/licenses/.

Also add information on how to contact you by electronic and paper mail.

If the program does terminal interaction, make it output a short notice like this when it starts in an interactive mode:

```
program Copyright (C) year name of author
This program comes with ABSOLUTELY NO WARRANTY; for details type 'show w'.
This is free software, and you are welcome to redistribute it
under certain conditions; type 'show c' for details.
```

The hypothetical commands 'show w' and 'show c' should show the appropriate parts of the General Public License. Of course, your program's commands might be different; for a GUI interface, you would use an "about box".

You should also get your employer (if you work as a programmer) or school, if any, to sign a "copyright disclaimer" for the program, if necessary. For more information on this, and how to apply and follow the GNU GPL, see http://www.gnu.org/licenses/.

The GNU General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Lesser General Public License instead of this License. But first, please read http://www.gnu.org/philosophy/why-not-lgpl.html.

Appendix J GNU Free Documentation License

Version 1.2, November 2002

Copyright © 2000,2001,2002 Free Software Foundation, Inc. 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301, USA.

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document free in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or non-commercially. Secondarily, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released

under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaT_EX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.

B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.

- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements."

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called

an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See http://www.gnu.org/copyleft/.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

J.1 ADDENDUM: How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

Copyright (C) year your name.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled ''GNU Free Documentation License''.

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the "with...Texts." line with this:

with the Invariant Sections being list their titles, with the Front-Cover Texts being list, and with the Back-Cover Texts being list.

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.

!	symbols	12
	tty	
'!' packet	tui	
	version	
11	windows	
#		
# (a comment) 19	'with-gdb-datadir'	
# in Modula-2	'with-relocated-sources'	
	'with-sysroot'	
	write	
\$	-b	
\$	-break-after	
	-break-commands	
\$\$	-break-condition	
\$_ and info breakpoints	-break-delete	302
\$_ and info line	-break-disable	302
\$_, \$, and value history 100	-break-enable	303
\$_, convenience variable	-break-info	303
\$, convenience variable	-break-insert	304
\$_exitcode, convenience variable 110	-break-list	305
\$_siginfo, convenience variable 110	-break-watch	306
\$bpnum, convenience variable	-c	12
\$cdir, convenience variable 87	-d	12
\$cwd, convenience variable 87	-data-disassemble	334
\$tpnum	-data-evaluate-expression	336
\$trace_file 139	-data-list-changed-registers	
\$trace_frame	-data-list-register-names	
\$trace_func 139	-data-list-register-values	
\$trace_line 139	-data-read-memory	
\$tracepoint 139	-e	
•	-enable-pretty-printing	
	-enable-timings	
-	-environment-cd.	
annotate	-environment-directory	
args	-environment-path	
batch	-environment-pwd	
batch-silent	-ex	
baud	-exec-arguments	
cd	-exec-arguments	
command	-exec-finish	
core		
directory	-exec-interrupt	
· ·	-exec-jump	
epoch	-exec-next	315 315
OVAL COMMANDA		010
exec	-exec-return	
fullname	-exec-run	
interpreter	-exec-step	
nowindows	-exec-step-instruction	
nx	-exec-until	
pid	-f	
quiet	-file-exec-and-symbols	
readnow	-file-exec-file	
return-child-result	-file-list-exec-source-file	
se	-file-list-exec-source-files	
silent	-file-symbol-file	
statistics	-gdb-exit	347

-gdb-set	'.build-id' directory 193
-gdb-show	'.debug' subdirectories
-gdb-version	'.gdbinit'
-inferior-tty-set	.gnu_debuglink sections 194
-inferior-tty-show	.note.gnu.build-id sections
-interpreter-exec	".o' files, reading symbols from
-1	
-list-features	1
-list-target-features 349	
-list-thread-groups	/proc
-n	, r
-nw	
-p	•
-q	::, context for variables/functions 95
-r	::, in Modula-2
-s	,
-stack-info-depth	
-stack-info-frame	<
-stack-list-arguments	<architecture></architecture>
-stack-list-frames	<pre><compatible>469</compatible></pre>
-stack-list-locals	<feature></feature>
-stack-list-variables	<pre><osabi>469</osabi></pre>
-stack-select-frame	<reg></reg>
-symbol-list-lines	<union></union>
-t	<pre><vector>469</vector></pre>
-target-attach	<xi:include></xi:include>
-target-detach	
-target-disconnect	?
-target-download	•
-target-file-delete	'?' packet 406
-target-file-get	
-target-file-put	
-target-select	<u>(0)</u>
-thread-info	©, referencing memory as an array 97
-thread-list-ids	
-thread-select	^
-var-assign	
-var-create	^connected
-var-delete	^done 295
-var-evaluate-expression	^error
-var-info-expression	^exit
-var-info-num-children	^running
-var-info-path-expression	
-var-info-type	
-var-list-children	-
-var-set-format	init on Command
-var-set-frozen	init on Function
-var-set-update-range	_NSPrintForDebugger, and printing Objective-C
-var-set-visualizer	objects
-var-show-attributes	
-var-show-format	4
-var-update	•
-w	"No symbol "foo" in current context" 96
-x	-
	٢
•	{
., Modula-2 scope operator	{type}94
,	

\mathbf{A}	assf	8
'A' packet	assignment	8
abbreviation	async output in GDB/MI	29.
abort (C-g)	async records in GDB/MI	290
accept-line (Newline or Return)	asynchronous execution	6
acknowledgment, for GDB remote	asynchronous execution, and process record and	
actions	replay	7
active targets	AT&T disassembly flavor	9
Ada	attach	3
Ada exception catching	attach to a program by name	20
Ada mode, general	attach&	
Ada task switching	auto-loading, Python	
Ada tasking and core file debugging	auto-retry, for remote TCP target 2	21
Ada, deviations from	automatic display	0
Ada, omissions from	automatic hardware breakpoints	4
Ada, problems	automatic overlay debugging 1	4
Ada, tasking	automatic thread selection	6
add new commands for external monitor 204	auxiliary vector	
add-inferior	AVR 2	23
add-shared-symbol-files 189	awatch	5
add-symbol-file		
add-symbol-file-from-memory	В	
address of a symbol	D	
address of LazyString	b (break)	4
address of Value	'b' packet	10
address size for remote targets	'B' packet	
ADP (Angel Debugger Protocol) logging 230	background execution	
advance location	backtrace	7
aggregates (Ada)	backtrace beyond main function	
AIX threads	backtrace limit	
alignment of remote memory accesses 408	backward-char (C-b) 3	
all-stop mode	backward-delete-char (Rubout) 3	
Alpha stack	backward-kill-line (C-x Rubout) 3	
ambiguous expressions94	backward-kill-word (M-DEL) 3	
AMD 29K register stack	backward-word (M-b) 3	
annotations	baud rate for remote targets 2	
annotations for errors, warnings and interrupts	'bc' packet	
357	bcache statistics	
annotations for invalidation messages 357	beginning-of-history (M-<) 3	
annotations for prompts	beginning-of-line (C-a)	
annotations for running programs	bell-style	
annotations for source display	bind-tty-special-chars	
append	bits in remote address	
append data to a file	bookmark	
apply command to several threads	break	
apropos	break task taskno (Ada) 1	
architecture debugging info	break thread threadno	
argument count in user-defined commands 255	break in overloaded functions	
arguments (to your program)	break on a system call	
arguments, to gdbserver	break on fork/exec	
arguments, to user-defined commands 255	BREAK signal instead of Ctrl-C	
ARM 32-bit mode	break, and Objective-C	
ARM RDI	breakpoint address adjusted	
array aggregates (Ada)	breakpoint annotation	
arrays 97	breakpoint commands for CDR/MI	
arrays in expressions	breakpoint commands for GDB/MI	
artificial array	breakpoint conditionsbreakpoint numbers	
absention in the first actions	DIEGEDOING HUMBELS	4

1 1 1 4	11: C 4:	104
breakpoint on events	calling functions	
breakpoint on memory address	calling make	
breakpoint on variable modification 43	capitalize-word (M-c)	
breakpoint ranges	case sensitivity in symbol names	
breakpoint subroutine, remote	case-insensitive symbol names	175
breakpointing Ada elaboration code 169	cast on Value	
breakpoints	casts, in expressions	93
breakpoints and tasks, in Ada 171	casts, to view memory	
breakpoints and threads	catch	
breakpoints in functions matching a regexp 45	catch Ada exceptions	
breakpoints in overlays	catch exceptions, list active handlers	
breakpoints, multiple locations	catchpoints	
breakpoints-invalid annotation	catchpoints, setting	
'bs' packet	cd	
bt (backtrace)	cdir	
bug criteria	Cell Broadband Engine	
bug reports	change working directory	30
bugs in GDB	character sets	
build ID sections	character-search (C-])	
build ID, and separate debugging files 193	character-search-backward (M-C-])	
building gdb, requirements for 391	charset	116
built-in simulator target	checkpoint	
	checkpoints and process id	41
	checks, range	150
\mathbf{C}	checks, type	
c (continue)	checksum, for GDB remote	
c (SingleKey TUI key)	children on pretty printer	
C and C++	choosing target byte order	
C and C++ checks	clear	
	clear, and Objective-C	
C and C++ constants	clear-screen (C-1)	
C and C++ defaults	clearing breakpoints, watchpoints, catchpoints	
C and C++ operators	clone-inferior	
'c' packet		
'C' packet	close, file-i/o system call	
C++	closest symbol and offset for an address	
C++ compilers	code address and its source line	
C++ exception handling 156	code of Type	
C++ overload debugging info	collect (tracepoints)	
C++ scope resolution	collected data discarded	
C++ symbol decoding style	colon, doubled as scope operator	
C++ symbol display 156	colon-colon, context for variables/functions	
C-L	command editing	365
C-x 1	command files	258
C-x 2	command history	245
C-x a	command hooks	257
C-x A	command interpreters	279
C-x C-a	command line editing	245
C-x o	command scripts, debugging	
C-x s	command tracing	
caching data of remote targets	COMMAND_BREAKPOINTS	
call	COMMAND_DATA	
call dummy stack unwinding	COMMAND_FILES	
call dummy stack unwinding on unhandled	COMMAND_MAINTENANCE	
· · · · · · · · · · · · · · · · · · ·	COMMAND_NONE	
exception	_	
call overloaded functions	COMMAND_OBSCURE	
call stack	COMMAND_RUNNING	
call stack traces	COMMAND_STACK	
call-last-kbd-macro (C-x e)	COMMAND STATUS	274

COMMAND_SUPPORT	core dump file
COMMAND_TRACEPOINTS	core dump file target
commands	core-file
commands annotation	crash of debugger
commands for C++	CRC algorithm definition
commands in python	CRC of memory block, remote request 416
commands to access python	CRIS
comment	CRIS mode
comment-begin	CRIS version
COMMON blocks, Fortran	Ctrl-BREAK, MS-Windows
common targets	ctrl-c message, in file-i/o protocol
${\tt compare-sections$	Ctrl-o (operate-and-get-next) 19
compatibility, gdb/mi and CLI	current Ada task ID
compilation directory	current directory 87
compiling, on Sparclet	current stack frame
complete	current thread
complete (TAB)	current thread, remote request
complete on Command	current_objfile
COMPLETE_COMMAND	cwd
COMPLETE_FILENAME	Cygwin DLL, debugging
COMPLETE_LOCATION	Cygwin-specific commands
COMPLETE_NONE	
COMPLETE_SYMBOL	D
completion	D
completion of Python commands	d (delete)
completion of quoted strings 20	d (SingleKey TUI key)
completion of structure field names 21	'd' packet 407
completion of union field names	'D' packet 407
completion-query-items	Darwin
compressed debug sections	data breakpoints
condition 57	data manipulation, in GDB/MI
conditional breakpoints	dead names, GNU Hurd
conditional tracepoints	debug expression parser
configuring GDB	debug formats and C++
confirmation	debug link sections
connection timeout, for remote TCP target 210	debug remote protocol
console i/o as part of file-i/o	debug_chaos
console interpreter	debugger crash
console output in GDB/MI	debugging C++ programs
const on Type	debugging information directory, global 193
constants, in file-i/o protocol	debugging information in separate files 193
continue	debugging multiple processes
continue&	debugging multithreaded programs (on HP-UX)
continuing	
continuing threads	debugging optimized code
control C, and remote debugging	debugging stub, example
controlling terminal	debugging target
convenience functions	debugging the Cygwin DLL
convenience functions in python	decimal floating point format
convenience variables	default collection action
convenience variables for tracepoints 139	default data directory
convenience variables, and trace state variables	default source path substitution
	default system root
convenience variables, initializing	define
convert-meta 369 copy-backward-word () 380	defining magnes interactively.
	defining macros interactively
copy-forward-word ()	delete
copy-region-as-kill ()	иетеле

delete breakpoints	document	
delete checkpoint checkpoint-id 41	documentation	389
delete display 101	don't repeat command	256
delete mem	don't repeat Python command	
delete tracepoint	dont-repeat	256
delete tvariable	dont_repeat on Command	272
delete-char (C-d)	DOS serial data link, remote debugging	220
delete-char-or-list ()	DOS serial port status	221
delete-horizontal-space () 380	down	. 80
deleting breakpoints, watchpoints, catchpoints	Down	283
55	down-silently	. 81
deliver a signal to a program	downcase-word (M-1)	
demangling C++ names	download server address (M32R)	
deprecated commands	download to Sparclet	
dereference on Value	download to VxWorks	
derived type of an object, printing 107	DPMI	
descriptor tables display	dump	
detach	dump all data collected at tracepoint	
detach (remote)	dump core from inferior	
detach from task, GNU Hurd	dump data to a file	
detach from thread, GNU Hurd	dump-functions ()	
detach inferior infno	dump-macros ()	
digit-argument (M-0, M-1, M) 381	dump-variables ()	
dir	dump/restore files	
direct memory access (DMA) on MS-DOS 220	DWARF 2 compilation units cache	
directories for source files	DWARF-2 CFI and CRIS	
	DWARF2 DIEs	
directory	dynamic linking	
directory, compilation	dynamic varobj	
directory, current	dynamic varobj	320
dis (disable)		
disable	\mathbf{E}	
disable display		
disable mem	e (edit)	
disable tracepoint	echo	
disable-completion	edit	
disassemble	editing	
disconnect	editing command lines	
disconnected tracing	editing source files	
displaced stepping debugging info	editing-mode	
displaced stepping support	eight-bit characters in strings	
displaced stepping, and process record and replay	elaboration phase	
	else	
display	Emacs	
display command history	emacs-editing-mode (C-e)	
display derived types		404
display displad out of some	empty response, for unsupported packets	
display disabled out of scope 102		
display GDB copyright	empty response, for unsupported packets	. 56
display GDB copyright 23 display of expressions 101	empty response, for unsupported packets	. 56 101
display GDB copyright 23 display of expressions 101 display remote monitor communications 201	empty response, for unsupported packets enable	. 56 101 113 131
display GDB copyright 23 display of expressions 101	empty response, for unsupported packets enable	. 56 101 113 131 369
display GDB copyright 23 display of expressions 101 display remote monitor communications 201	empty response, for unsupported packets enable	. 56 101 113 131 369
display GDB copyright 23 display of expressions 101 display remote monitor communications 201 display remote packets 252	empty response, for unsupported packets enable	. 56 101 113 131 369 . 55
display GDB copyright 23 display of expressions 101 display remote monitor communications 201 display remote packets 252 display_hint on pretty printer 270	empty response, for unsupported packets enable	. 56 101 113 131 369 . 55 278
display GDB copyright 23 display of expressions 101 display remote monitor communications 201 display remote packets 252 display_hint on pretty printer 270 DJGPP debugging 219	empty response, for unsupported packets enable	. 56 101 113 131 369 . 55 278
display GDB copyright23display of expressions101display remote monitor communications201display remote packets252display_hint on pretty printer270DJGPP debugging219dll-symbols221	empty response, for unsupported packets enable	. 56 101 113 131 369 . 55 278 . 58 259
display GDB copyright 23 display of expressions 101 display remote monitor communications 201 display remote packets 252 display_hint on pretty printer 270 DJGPP debugging 219 dll-symbols 221 DLLs with no debugging symbols 222 do (down) 80	empty response, for unsupported packets enable enable display enable mem enable-keypad enable/disable a breakpoint encoding of LazyString end (breakpoint commands) end (if/else/while commands)	. 56 101 113 131 369 . 55 278 . 58 259 256
display GDB copyright23display of expressions101display remote monitor communications201display remote packets252display_hint on pretty printer270DJGPP debugging219dll-symbols221DLLs with no debugging symbols222	empty response, for unsupported packets enable enable display enable mem enable tracepoint enable-keypad enable/disable a breakpoint encoding of LazyString end (breakpoint commands) end (if/else/while commands) end (user-defined commands)	. 56 101 113 131 369 . 55 278 . 58 259 256 382
display GDB copyright 23 display of expressions 101 display remote monitor communications 201 display remote packets 252 display_hint on pretty printer 270 DJGPP debugging 219 dll-symbols 221 DLLs with no debugging symbols 222 do (down) 80 do not print frame argument values 104	empty response, for unsupported packets enable enable display. enable mem. enable tracepoint. enable-keypad. enable/disable a breakpoint. encoding of LazyString. end (breakpoint commands) end (if/else/while commands). end (user-defined commands) end-kbd-macro (C-x))	. 56 101 113 131 369 . 55 278 259 256 382 378

entering numbers	file transfer	205
environment (of your program) 29	file transfer, remote protocol	432
errno values, in file-i/o protocol	file-i/o examples	448
error annotation	file-i/o overview	
error on valid input	File-I/O remote protocol extension	
error-begin annotation	file-i/o reply packet	
event debugging info	file-i/o request packet	
event designators	filename of Objfile	
event handling	fin (finish)	
examine process image	find	
examining data93	find downloadable SREC files (M32R)	
examining memory	find trace snapshot	
exception handlers 51	finish	
exception handlers, how to list 82	finish&	
exceptionHandler	flinching	
exceptions, python	float promotion	
exchange-point-and-mark (C-x C-x)	floating point	
exec-file	floating point registers	
executable file	floating point, MIPS remote	
executable file target	flush	
executable file, for remote target	flush_i_cache	
execute	flushregs	
execute commands from a file	focus	
execute forward or backward in time		
execute remote command, remote request 419	focus of debugging	
execution, foreground, background and	foo	
asynchronous	foreground execution	
· ·	fork, debugging programs which call	
exited annotation	format options	
exiting GDB	formatted output	
expand macro once	Fortran	
expand-tilde	Fortran Defaults	
expanding preprocessor macros	Fortran operators and expressions	
expression debugging info	Fortran-specific support in GDB	
expression parser, debugging info	forward-backward-delete-char ()	
expressions	forward-char (C-f)	
expressions in Ada	forward-search	
expressions in C or C++	forward-search-history (C-s)	
expressions in C++	forward-word (M-f)	
expressions in Modula-2	FR-V shared-library debugging	
extend GDB for remote targets	frame debugging info	
extending GDB	frame number	
extra signal information	frame pointer	. 77
	frame pointer register	111
F	frame, command	. 78
r	frame, definition	. 77
f (frame)	frame, selecting	. 80
f (SingleKey TUI key)	frame_stop_reason_string	277
'F' packet 407	frameless execution	
F reply packet	frames in python	
F request packet	frames-invalid annotation	357
fast tracepoints	free memory information (MS-DOS)	219
fatal signal	fstat, file-i/o system call	
fatal signals	ftrace	
FDL, GNU Free Documentation License 489	Fujitsu	
features of the remote protocol	full symbol tables, listing GDB's internal	
fg (resume foreground execution)	Function	
fields on Type	function call arguments, optimized out	
file	function entry/exit, wrong values of variables.	
	Tarrough Citor (CAIU, WICHE VALUED OF VALIABLES .	

functions without line info, and stepping \ldots	. 61	gdb.TYPE_CODE_INTERNAL_FUNCTION	
		gdb.TYPE_CODE_MEMBERPTR	269
		gdb.TYPE_CODE_METHOD	
G		gdb.TYPE_CODE_METHODPTR	269
'g' packet	407	gdb.TYPE_CODE_NAMESPACE	269
'G' packet		gdb.TYPE_CODE_PTR	268
g++, GNU C++ compiler		gdb.TYPE_CODE_RANGE	268
garbled pointers		gdb.TYPE_CODE_REF	269
GCC and C++		gdb.TYPE_CODE_SET	
gcore		gdb.TYPE_CODE_STRING	
		gdb.TYPE_CODE_STRUCT	
GDB bugs, reporting		gdb.TYPE_CODE_TYPEDEF	
GDB internal error		gdb.TYPE_CODE_UNION	
gdb module		gdb.TYPE_CODE_VOID	
GDB reference card		gdb.Value	
GDB startup		gdb.write	
GDB version number			
gdb.COMMAND_BREAKPOINTS		GDB/MI development	
gdb.COMMAND_DATA		GDB/MI General Design	
gdb.COMMAND_FILES	274	GDB/MI, async records	
gdb.COMMAND_MAINTENANCE	274	GDB/MI, breakpoint commands	
gdb.COMMAND_NONE	273	GDB/MI, compatibility with CLI	
gdb.COMMAND_OBSCURE		GDB/MI, data manipulation	
gdb.COMMAND_RUNNING		GDB/MI, input syntax	
gdb.COMMAND_STACK		GDB/MI, its purpose	289
gdb.COMMAND_STATUS		GDB/MI, output syntax	293
gdb.COMMAND_SUPPORT		GDB/MI, result records	295
gdb.COMMAND_TRACEPOINTS		GDB/MI, simple examples	
gdb.COMPLETE_COMMAND		GDB/MI, stream records	
gdb.COMPLETE_FILENAME		gdbarch debugging info	
		GDBHISTFILE, environment variable	
gdb.COMPLETE_LOCATION		gdbserver	
gdb.COMPLETE_NONE		gdbserver, multiple processes	
gdb.COMPLETE_SYMBOL		gdbserver, search path for libthread_db	
gdb.current_objfile		GDT	
gdb.execute			
gdb.flush		generate-core-file	
gdb.Function		get thread-local storage address, remote request	
gdb.history			
'gdb.ini'		getDebugChar	
gdb.LazyString	277	gettimeofday, file-i/o system call	
<pre>gdb.lookup_type</pre>	266	global debugging information directory	
gdb.Objfile	276	GNU C++	
gdb.objfiles		GNU Emacs.	
gdb.parameter		GNU Hurd debugging	
gdb.parse_and_eval		GNU/Hurd debug messages	
gdb.selected_frame		GNU/Linux LWP async debug messages	252
gdb.Type		GNU/Linux LWP debug messages	252
gdb.TYPE_CODE_ARRAY		<pre>gnu_debuglink_crc32</pre>	195
gdb.TYPE_CODE_BITSTRING			
gdb.TYPE_CODE_BOOL			
gdb.TYPE_CODE_CHAR		H	
gdb.TYPE_CODE_COMPLEX		h (holp)	91
		h (help)	
gdb.TYPE_CODE_DECFLOAT		'H' packet	
gdb.TYPE_CODE_ENUM		handle	
gdb.TYPE_CODE_ERROR		handle_exception	
gdb.TYPE_CODE_FLAGS		handling signals	
gdb.TYPE_CODE_FLT		hardware breakpoints	
gdb.TYPE_CODE_FUNC		hardware debug registers	
gdb.TYPE_CODE_INT	268	hardware watchpoints	. 49

hash mark while downloading	info all-registers
hbreak	info args
help 21	info auxv
help function	info breakpoints
help target	info catch 82
help user-defined	info checkpoints
heuristic-fence-post (Alpha, MIPS) 241	info classes
history	info common
history events	info copying
history expansion	info dcache
history expansion, turn on/off 246	info display
history file	info dl1
history number	info dos
history of values printed by GDB 108	info extensions
history size	info f (info frame)
history substitution	info files
history-preserve-point	info float
history-search-backward ()	info for known object files 400
history-search-forward ()	info frame
HISTSIZE, environment variable	info frame, show the source language 149
hook	info functions 177
hookpost	info handle
hooks, for commands	info inferiors
hooks, post-command	info io_registers, AVR
hooks, pre-command	info line
horizontal-scroll-mode	info line, and Objective-C 157
host character set	info locals
Host I/O, remote protocol	info macro
how many arguments (user-defined commands)	info mem
	info meminfo
HPPA support	info or1k spr
htrace	info os processes
hwatch	info pidlist
IIW 400II	info proc
	info program
I	info record
i (info)	info registers
'i' packet	info scope
'I' packet	info selectors
i/o30	info serial
I/O registers (Atmel AVR)	info set
i386	info share
'i386-stub.c'	info sharedlibrary
IDT	info signals
if	info source
ignore	info source, show the source language 149
	info sources
ignore count (of breakpoint)	
INCLUDE_RDB	info spu
incomplete type	info stack
indentation in structure display	info symbol
inferior	info target
inferior debugging info	info task taskno
inferior functions, calling	info tasks
inferior infno	
inferior tty	info threads
infinite recursion in user-defined commands 256	info threads (HP-UX)
info	info tp
info address	info tracepoints

info tvariables	K
info types	
info udot	n packet 100
info variables	Reffici classi damp
info vector	Kerner memory image
	Keymap
info w32	
info warranty	
info watchpoints $[n]$	
info win	
information about tracepoints	
inheritance	
init file	
init file name	
$\verb init-if-undefined $	kvm
initial frame	
initialization file, readline 368	3 $^{-}$ $^{-}$
inline functions, debugging	_
innermost frame	
input syntax for GDB/MI	languages
input-meta	last tracepoint number
insert-comment (M-#)	latest breakpoint 44
insert-completions (M-*)	layout
inspect	lazy strings in python
installation	lazy_string on Value
instructions, assembly	LDT 219
integral datatypes, in file-i/o protocol 445	leaving GDB 16
Intel	Left 283
Intel disassembly flavor	tength of LazyString
interaction, readline	11DKVIII
internal commands	, indiary list format, remote protocol 449
internal errors, control of GDB behavior 399	mint hardware breakpoints and watenpoints 209
	mint on number of printed array elements 104
internal GDB breakpoints	minus, in inc-i/o protocor
interpreter-exec	intespec
interrupt	Emax lightweight processes
interrupt	1100
interrupt debuggee on MS-Windows	inst decire chirecas, remote request
interrupt remote programs	
interrupting remote programs	, , , , , , , , , , , , , , , , , , ,
interrupting remote targets	· · ·) · · · · · · · · · · · · · · · ·
interrupts (remote protocol)	, , , , , , , , , , , , , , , , , , , ,
invalid input	v
invoke another interpreter	
invoke on Command	U II v
invoke on Function	,
${\tt is_optimized_out} \ \ of \ {\tt Value} \ \dots \dots \ \ \underline{ 264}$	load filename 201
is_valid on Frame	
isatty, file-i/o system call	load symbols from memory
isearch-terminators	
	locate address
т	lock scheduler
J	log output in GDB/MI
JIT compilation interface	logging file name
jump	logging GDB output
jump, and Objective-C	, 100kup_type
just-in-time compilation	loop_break
Jase in vinic compnation	loop_continue
	lseek flags, in file-i/o protocol

lseek, file-i/o system call	maint set python auto-load	263
	maint set python print-stack	261
Th. 4T	maint set show-debug-regs	401
\mathbf{M}	maint show dwarf2 max-cache-age	401
'm' packet	maint show internal-error	399
'M' packet	maint show internal-warning	399
M32-EVA target board address	maint show profile	401
M32R/Chaos debugging	maint show show-debug-regs	
m680x0	maint space	
'm68k-stub.c'	maint time	
machine instructions	maint translate-address	
macro define	maint undeprecate	
macro definition, showing	maintenance commands	
macro exp1	make	
macro expand	manual overlay debugging	
macro expansion, showing the results of	map an overlay	
preprocessor	mapinfo list, QNX Neutrino	
macro list	mapped address	
macro undef	mapped overlays	
macros, example of debugging with	mark-modified-lines	
macros, user-defined	mark-symlinked-directories	
mailing lists	match-hidden-files	
maint agent	maximum value for offset of closest symbol	
	mem	
maint agent-eval 397 maint check-symtabs 398	member functions	
maint check-symtabs	memory address space mappings	
maint cplus lilst_component	memory map format	
-	memory region attributes	
maint demangle	memory tracing	
maint deprecate 398 maint dump-me 398	memory transfer, in file-i/o protocol	
	memory used by commands	
maint info breakpoints	memory used for symbol tables	
maint info program-spaces	memory, alignment and size of remote accesses	101
maint info psymtabs	memory, angument and size of remote accesses	408
maint info sections	memory, viewing as typed object	
	memset	
maint info symtabs 179 maint internal-error 398	menu-complete ()	
	meta-flag	
maint internal-warning	mi interpreter	
maint packet	mil interpreter	
maint print architecture	mi2 interpreter	
maint print c-tdesc	minimal language	
	Minimal symbols and DLLs	
maint print dummy-frames	MIPS addresses, masking	
maint print objfiles	MIPS boards	
maint print psymbols	MIPS code compression	
maint print raw-registers	MIPS remote floating point	
maint print reggroups	MIPS stack	
maint print register-groups	miscellaneous settings	
maint print registers		
maint print statistics	MMX registers (x86)	
maint print symbols	mode_t values, in file-i/o protocol	
maint print target-stack	Modula-2	
maint print type		
maint print unwind, HPPA	Modula-2 checks	
maint set dwarf2 max-cache-age	Modula-2 constants	
maint set internal-error	Modula-2 defaults	
maint set internal-warning	Modula-2 operators	
maint set profile 401	Modula-2 types	162

Modula-2, deviations from	numbers for breakpoints 43
Modula-2, gdb support	
monitor	
monitor commands, for gdbserver 207	0
Motorola 680x0	object files, relocatable, reading symbols from
MS Windows debugging	189
MS-DOS system info	Objective-C
MS-DOS-specific commands	Objective-C, classes and selectors
multiple locations, breakpoints	Objective-C, print objects
multiple processes	Objfile
multiple processes with gdbserver 206	objfiles
multiple targets	objfiles in python
multiple threads	observer debugging info
multiple threads, backtrace	octal escapes in strings 105
multiple-symbols menu 94	older on Frame
multiprocess extensions, in remote protocol 423	online documentation
	opaque data types
TA T	open flags, in file-i/o protocol
\mathbf{N}	open, file-i/o system call
n (next)	OpenRISC 1000
n (SingleKey TUI key)	OpenRISC 1000 htrace
name on Frame	operating system information 475
names of symbols	operating system information, process list 475
namespace in C++	optimized code, debugging
native Cygwin debugging	optimized code, wrong values of variables 96
native DJGPP debugging	optimized out value in Python
negative breakpoint numbers	optional debugging messages
NetROM ROM emulator target	optional warnings 249
New systag message	or1k boards
New systag message, on HP-UX	or1ksim
newer on Frame	OS ABI
next	OS information
next&	out-of-line single-stepping
next-history (C-n)	outermost frame
nexti	output
nexti&	output formats
ni (nexti)	output syntax of GDB/MI
non-incremental-forward-search-history (M-n)	output-meta
378	overlay
non-incremental-reverse-search-history (M-p)	overlay area
378	overlay example program
non-member C++ functions, set breakpoint in 45	overlays
non-stop mode	overlays, setting breakpoints in 144
non-stop mode, and breakpoint always-inserted	overload-choice annotation
	overloaded functions, calling 154
non-stop mode, and process record and replay 73	overloaded functions, overload resolution 156
non-stop mode, and 'set displaced-stepping'	overloading in C++
	overwrite-mode ()
non-stop mode, remote request	
noninvasive task options	D
nosharedlibrary	P
notation, readline	'p' packet
notational conventions, for GDB/MI 289	'P' packet 408
notification packets	packet acknowledgment, for GDB remote 436
notify output in GDB/MI	packet size, remote protocol 422
NULL elements in arrays 105	packets, notification
number of array elements to print 104	packets, reporting on stdout
number representation	packets, tracepoint

page tables display (MS-DOS)	print messages on inferior start and exit	34
page-completions	print messages on thread start and exit	37
parameter	print settings	102
parse_and_eval	print structures in indented form	105
partial symbol dump	print-object	
partial symbol tables, listing GDB's internal 179	print/don't print memory addresses	
Pascal 1	printf	
Pascal objects, static members display 107	printing byte arrays	
Pascal support in GDB, limitations 159	printing data	
pass signals to inferior, remote request 418	printing frame argument values	
passcount	printing strings	
patching binaries	proc-trace-entry	
patching object files	proc-trace-exit	
path	proc-untrace-entry	
pause current task (GNU Hurd)	proc-untrace-exit	
pause current thread (GNU Hurd)	process detailed status information	
	process ID	
pauses in output	process info via '/proc'	
pc on Frame	process list, QNX Neutrino	
pending breakpoints		
PgDn	process record and replay	
PgUp	process status register	
physical address from linear address	processes, multiple	
pipe, target remote to	procfs API calls	
pipes	profiling GDB	
pmon, MIPS remote	program counter register	
po (print-object) 158	program entry point	
pointer on Type	programming in python	
pointer values, in file-i/o protocol 445	prompt	
pointer, finding referent	prompt annotation	
port rights, GNU Hurd	prompt-for-continue annotation	
port sets, GNU Hurd	protocol basics, file-i/o	
possible-completions (M-?) 381	protocol, GDB remote serial	
post-commands annotation	protocol-specific representation of datatypes, in	
post-overload-choice annotation 356	file-i/o protocol	
post-prompt annotation	ptrace system call	112
post-prompt-for-continue annotation 356	ptype	176
post-query annotation	<pre>putDebugChar</pre>	214
PowerPC architecture	pwd	30
pre-commands annotation	python	261
pre-overload-choice annotation 356	python api	262
pre-prompt annotation	python commands	272
pre-prompt-for-continue annotation 356	python convenience functions	
pre-query annotation	python exceptions	
prefix for data files	python functions	
prefix for shared library file names	python module	
prefix-meta (ESC)	python pagination	
premature return from system calls 69	python scripting	
preprocessor macro expansion, showing the results	python stdout	
of	Python, working with types	
	python, working with values from inferior	
pretty print arrays	pj mon, worming with various from interior	
pretty print C++ virtual function tables 108		
pretty_printers of Objfile	Q	
previous-history (C-p)	•	10
print	q (quit)	
print all frame argument values	q (SingleKey TUI key)	
print an Objective-C object description	'q' packet	
print array indexes	'Q' packet	
print frame argument values for scalars only 104	'qAttached' packet	427

'qAuth' packet	recent tracepoint number	
'qC' packet	record	
'qCRC' packet	record aggregates (Ada)	167
'qfThreadInfo' packet	record delete	. 75
'qGetTLSAddr' packet	record mode	. 73
'QNonStop' packet	record serial communications on file	209
QNX Neutrino	record stop	. 73
'qOffsets' packet	recording a session script	
'qP' packet	recording inferior's execution and replaying it.	
'QPassSignals' packet	redirection	
'qRcmd' packet	redraw-current-line ()	
'qSearch:memory' packet	reference card	
'QStartNoAckMode' packet	reference declarations	
'qsThreadInfo' packet	reference on Type	
'qSupported' packet	refresh	
'qSymbol' packet	register stack, AMD29K	
'QTDV' packet	registers	
'qThreadExtraInfo' packet		
	regs, Super-H	
'qTV' packet	regular expression	
query annotation	reloading symbols	
query attached, remote request	reloading the overlay table	
quit [expression]	relocatable object files, reading symbols from	
quit annotation	remote connection without stubs	
quoted-insert (C-q or C-v) 379	remote debugging	
quotes in commands	remote delete	
quoting Ada internal identifiers	remote get	
quoting names	remote memory comparison	
'qXfer' packet	remote monitor prompt	234
	remote packets, enabling and disabling	210
D	remote programs, interrupting	204
\mathbf{R}	remote programs, interrupting remote protocol debugging	252
	remote programs, interrupting	252
r (run)	remote programs, interrupting remote protocol debugging	$\frac{252}{403}$
r (run)	remote programs, interrupting remote protocol debugging remote protocol, binary data remote protocol, field separator	252 403 403
r (run)	remote programs, interrupting remote protocol debugging remote protocol, binary data	252 403 403 205
r (run)	remote programs, interrupting	252 403 403 205 415
r (run)	remote programs, interrupting	252 403 403 205 415 215
r (run)	remote programs, interrupting	252 403 403 205 415 215 212
r (run)	remote programs, interrupting	252 403 403 205 415 215 212 403
r (run) 25 r (SingleKey TUI key) 283 'r' packet 408 'R' packet 408 raise exceptions 54 range checking 150 range on Type 267 ranges of breakpoints 43	remote programs, interrupting	252 403 403 205 415 215 212 403 213
r (run) 25 r (SingleKey TUI key) 283 'r' packet 408 'R' packet 408 raise exceptions 54 range checking 150 range on Type 267 ranges of breakpoints 43 raw printing 98	remote programs, interrupting remote protocol debugging remote protocol, binary data. remote protocol, field separator remote put remote query requests remote serial debugging summary remote serial debugging, overview remote serial protocol remote serial stub remote serial stub list	252 403 403 205 415 215 212 403 213
r (run)	remote programs, interrupting remote protocol debugging remote protocol, binary data. remote protocol, field separator remote put remote query requests remote serial debugging summary remote serial debugging, overview remote serial protocol remote serial stub remote serial stub list remote serial stub, initialization	252 403 403 205 415 215 212 403 213 213 213
r (run)	remote programs, interrupting remote protocol debugging remote protocol, binary data. remote protocol, field separator remote put remote query requests remote serial debugging summary remote serial debugging, overview remote serial stub remote serial stub remote serial stub list remote serial stub, initialization remote serial stub, main routine	252 403 403 205 415 215 212 403 213 213 214
r (run) 25 r (SingleKey TUI key) 283 'r' packet 408 'R' packet 408 raise exceptions 54 range checking 150 range on Type 267 ranges of breakpoints 43 raw printing 98 rbreak 45 rc (reverse-continue) 71 RDI heartbeat 231	remote programs, interrupting remote protocol debugging remote protocol, binary data remote protocol, field separator remote put remote query requests remote serial debugging summary remote serial debugging, overview remote serial protocol remote serial stub remote serial stub list remote serial stub, initialization remote serial stub, main routine remote stub, example	252 403 403 205 415 212 403 213 213 213 214 212
r (run) 25 r (SingleKey TUI key) 283 'r' packet 408 'R' packet 408 raise exceptions 54 range checking 150 range on Type 267 ranges of breakpoints 43 raw printing 98 rbreak 45 rc (reverse-continue) 71 RDI heartbeat 231 rdilogenable 230	remote programs, interrupting remote protocol debugging remote protocol, binary data remote protocol, field separator remote put remote query requests remote serial debugging summary remote serial debugging, overview remote serial stub remote serial stub list remote serial stub, initialization remote serial stub, main routine remote stub, example remote stub, support routines	252 403 403 205 415 215 212 403 213 213 214 214 214
r (run) 25 r (SingleKey TUI key) 283 'r' packet 408 'R' packet 408 raise exceptions 54 range checking 150 range on Type 267 ranges of breakpoints 43 raw printing 98 rbreak 45 rc (reverse-continue) 71 RDI heartbeat 231 rdilogenable 230 rdilogfile 230	remote programs, interrupting remote protocol debugging remote protocol, binary data remote protocol, field separator remote put remote query requests remote serial debugging summary remote serial debugging, overview remote serial stub remote serial stub remote serial stub list remote serial stub, initialization remote serial stub, main routine remote stub, example remote stub, support routines remote target	252 403 403 205 415 215 212 403 213 213 214 212 214 200
r (run)	remote programs, interrupting remote protocol debugging remote protocol, binary data. remote protocol, field separator remote put remote query requests remote serial debugging summary remote serial debugging, overview remote serial protocol remote serial stub remote serial stub list remote serial stub, initialization remote serial stub, main routine remote stub, example remote stub, support routines remote target remote target, file transfer	252 403 403 205 415 212 403 213 213 214 212 214 200 205
r (run)	remote programs, interrupting remote protocol debugging remote protocol, binary data remote protocol, field separator remote put remote query requests remote serial debugging summary remote serial debugging, overview remote serial stub remote serial stub list remote serial stub, initialization remote serial stub, main routine remote stub, example remote target remote target, file transfer remote target, limit break- and watchpoints.	252 403 403 205 415 212 403 213 213 214 212 214 200 205 209
r (run)	remote programs, interrupting remote protocol debugging remote protocol, binary data remote protocol, field separator remote put remote query requests remote serial debugging summary remote serial debugging, overview remote serial stub remote serial stub remote serial stub list remote serial stub, initialization remote serial stub, main routine remote stub, example remote stub, support routines remote target remote target, file transfer remote target, limit break- and watchpoints remote timeout	252 403 403 205 415 212 403 213 213 214 212 214 200 205 209 209
r (run) 25 r (SingleKey TUI key) 283 'r' packet 408 'R' packet 408 raise exceptions 54 range checking 150 range on Type 267 ranges of breakpoints 43 raw printing 98 rbreak 45 rc (reverse-continue) 71 RDI heartbeat 231 rdilogenable 230 rdilogfile 230 re-read-init-file (C-x C-r) 382 read special object, remote request 425 read, file-i/o system call 441 read-only sections 191	remote programs, interrupting remote protocol debugging remote protocol, binary data remote protocol, field separator remote put remote query requests remote serial debugging summary remote serial debugging, overview remote serial stub remote serial stub list remote serial stub, initialization remote serial stub, main routine remote stub, example remote stub, support routines remote target remote target, file transfer remote target, limit break- and watchpoints remote timeout remotetimeout	252 403 403 205 415 212 403 213 213 214 212 214 200 205 209 237
r (run)	remote programs, interrupting remote protocol debugging remote protocol, binary data remote protocol, field separator remote put remote query requests remote serial debugging summary remote serial debugging, overview remote serial stub remote serial stub remote serial stub list remote serial stub, initialization remote serial stub, main routine remote stub, example remote stub, support routines remote target remote target, file transfer remote target, limit break- and watchpoints remote timeout remotetimeout remotetimeout remote actions from a tracepoint	252 403 403 205 415 215 212 403 213 213 214 212 205 209 209 237 133
r (run) 25 r (SingleKey TUI key) 283 'r' packet 408 'R' packet 408 raise exceptions 54 range checking 150 range on Type 267 ranges of breakpoints 43 raw printing 98 rbreak 45 rc (reverse-continue) 71 RDI heartbeat 231 rdilogenable 230 rdilogfile 230 re-read-init-file (C-x C-r) 382 read special object, remote request 425 read, file-i/o system call 441 read-only sections 191 read_var on Frame 277 reading symbols from relocatable object files 189	remote programs, interrupting remote protocol debugging remote protocol, binary data remote protocol, field separator remote put remote query requests remote serial debugging summary remote serial debugging, overview remote serial stub remote serial stub list remote serial stub, initialization remote serial stub, main routine remote stub, example remote stub, support routines remote target remote target, file transfer remote target, limit break- and watchpoints remote timeout remove-inferior	252 403 403 205 415 212 403 213 213 214 200 205 209 209 237 133 . 33
r (run) 25 r (SingleKey TUI key) 283 'r' packet 408 'R' packet 408 raise exceptions 54 range checking 150 range on Type 267 ranges of breakpoints 43 raw printing 98 rbreak 45 rc (reverse-continue) 71 RDI heartbeat 231 rdilogenable 230 rdilogfile 230 re-read-init-file (C-x C-r) 382 read special object, remote request 425 read, file-i/o system call 441 read-only sections 191 read_var on Frame 277 reading symbols from relocatable object files 189 reading symbols immediately 188	remote programs, interrupting remote protocol debugging remote protocol, binary data remote protocol, field separator remote put remote query requests remote serial debugging summary remote serial debugging, overview remote serial protocol remote serial stub remote serial stub bist remote serial stub, initialization remote serial stub, main routine remote stub, example remote stub, support routines remote target remote target, file transfer remote target, limit break- and watchpoints remote timeout remove-inferior rename, file-i/o system call	252 403 403 205 415 212 403 213 213 214 200 205 209 209 237 133 . 33 442
r (run)	remote programs, interrupting remote protocol debugging remote protocol, binary data remote protocol, field separator remote put remote query requests remote serial debugging summary remote serial debugging, overview remote serial protocol remote serial stub remote serial stub list remote serial stub, initialization remote serial stub, main routine remote stub, example remote stub, support routines remote target remote target, file transfer remote target, limit break- and watchpoints remote timeout remove-inferior rename, file-i/o system call Renesas	252 403 403 205 415 215 212 403 213 214 212 214 200 205 209 237 133 . 33 442 213
r (run)	remote programs, interrupting remote protocol debugging remote protocol, binary data remote protocol, field separator remote put remote query requests remote serial debugging summary remote serial debugging, overview remote serial protocol remote serial stub remote serial stub list remote serial stub, initialization remote serial stub, main routine remote stub, example remote stub, support routines remote target remote target, file transfer remote target, limit break- and watchpoints remote timeout remove-inferior rename, file-i/o system call Renesas repeated array elements	252 403 403 205 415 215 212 403 213 214 212 214 200 205 209 237 133 . 33 442 213 105
r (run)	remote programs, interrupting remote protocol debugging remote protocol, binary data remote protocol, field separator remote put remote query requests remote serial debugging summary remote serial debugging, overview remote serial protocol remote serial stub remote serial stub bist remote serial stub, initialization remote serial stub, main routine remote stub, example remote target remote target, file transfer remote target, limit break- and watchpoints remote timeout remove-inferior rename, file-i/o system call Renesas repeated array elements repeating command sequences	252 403 403 205 415 212 403 213 213 214 200 205 209 237 133 . 33 442 213 . 105 . 19
r (run)	remote programs, interrupting remote protocol debugging remote protocol, binary data remote protocol, field separator remote put remote query requests remote serial debugging summary remote serial debugging, overview remote serial protocol remote serial stub remote serial stub bist remote serial stub, initialization remote serial stub, main routine remote stub, example remote target remote target, file transfer remote target, limit break- and watchpoints remote timeout remove-inferior rename, file-i/o system call Renesas repeated array elements repeating commands	252 403 403 205 415 212 403 213 213 214 212 214 200 205 209 237 133 442 213 105
r (run)	remote programs, interrupting remote protocol debugging remote protocol, binary data remote protocol, field separator remote put remote query requests remote serial debugging summary remote serial debugging, overview remote serial protocol remote serial stub remote serial stub bist remote serial stub, initialization remote serial stub, main routine remote stub, example remote target remote target, file transfer remote target, limit break- and watchpoints remote timeout remove-inferior rename, file-i/o system call Renesas repeated array elements repeating command sequences	252 403 403 205 415 212 403 213 213 214 212 214 200 205 209 237 133 442 213 105

reporting bugs in GDB	sdireset	231
reprint the last value	sdistatus	
reset SDI connection, M32R	SDS protocol	
response time, MIPS debugging	sds, a command	
restart	search	
restart checkpoint-id	search path for libthread_db	
restore	searching memory	
restore data from a file	searching memory, in remote debugging	
result records in GDB/MI	searching source files	
resume threads of multiple processes	section	
simultaneously	section offsets, remote request	
resuming execution	segment descriptor tables	219
RET (repeat last command)	select Ctrl-C, BREAK or BREAK-g	
retransmit-timeout, MIPS protocol 234	select trace snapshot	
return	select-frame	
returning from a function	selected frame	
reverse execution	selected_frame	277
reverse-continue	selecting frame silently	. 78
reverse-finish	self-insert (a, b, A, 1, !,)	379
reverse-next	send command to remote monitor	204
reverse-nexti	send command to simulator	229
reverse-search	send interrupt-sequence on start	210
reverse-search-history (C-r)	send PMON command	234
reverse-step	send rights, GNU Hurd	225
reverse-stepi	sending files to remote systems	205
revert-line (M-r)	separate debugging information files	193
rewind program state	sequence-id, for GDB remote	
Right	serial connections, debugging	
rn (reverse-next)	serial line, target remote	
rni (reverse-nexti)	serial protocol, GDB remote	
ROM at zero address, RDI	server prefix	
rs (step)	server, command prefix	
rsi (reverse-stepi)	set	
run	set ABI for MIPS	241
run to main procedure	set ada trust-PAD-over-XVS	
run until specified location	set annotate	
run&	set architecture	199
running	set args	28
running and debugging Sparclet programs 238	set arm	
running programs backward	set auto-solib-add	191
running VxWorks tasks	set backtrace	
running, on Sparclet	set board-address	
rwatch 50	set breakpoint always-inserted	
	set breakpoint auto-hw	
	set breakpoint pending	
\mathbf{S}	set breakpoints in many functions	
s (SingleKey TUI key)	set breakpoints on all functions	
s (step)	set can-use-hw-watchpoints	
's' packet	set case-sensitive	
'S' packet	set charset	117
save command history	set check range	
save GDB output to a file	set check type	
save tracepoints for future sessions	set coerce-float-to-double	
save-tracepoints	set com1base	
scheduler locking mode	set com1irq	_
scope	set com2base	
scripting commands	set com2irq	
scripting with python 261	•	220

	com3irq		set mem inaccessible-by-default	
	com4base		set mips abi	241
set	com4irq	220	set mips compression	242
set	complaints	250	set mips mask-address	242
set	confirm	250	set mipsfpu	
set	cp-abi	249	set monitor-prompt, MIPS remote	
set	cygwin-exceptions	221	set monitor-warnings, MIPS remote	234
set	data-directory	198	set multiple-symbols	
set	debug	251	set new-console	221
	debug darwin		set new-group	222
	debug hppa		set non-stop	
	debug mach-o		set opaque-type-resolution	
	debug mips		set osabi	
set	debug monitor	201	set output-radix	
set	debug nto-debug	226	set overload-resolution	
	debug-file-directory		set pagination	
	debugevents		set powerpc	
	debugexceptions		set print	
	debugexec		set print frame-arguments	
	debugmemory		set print inferior-events	
			-	
	default-collect		set print thread-events	
	demangle-style		set processor	
	detach-on-fork		set procfs-file	218
	disable-randomization		set procfs-trace	
	disassemble-next-line		set prompt	
	disassembly-flavor		set radix	
	disconnected-tracing		set rdiheartbeat	
set	displaced-stepping	397	set rdiromatzero	
	download-path		set record insn-number-max	
	editing		set record stop-at-limit	. 74
set	endian	202	set remote	
set	environment	29	set remote system-call-allowed	445
set	exceptions, Hurd command	224	set remote-mips64-transfers-32bit-regs	242
set	exec-direction	. 72	set remotecache	119
set	exec-done-display	251	set remoteflow	208
set	exec-wrapper	. 27	set retransmit-timeout	234
	extension-language		set rstack_high_address	241
	fast tracepoint		set schedule-multiple	66
	follow-exec-mode		set script-extension	
set	follow-fork-mode	38	set sdstimeout	
	gnutarget		set server-address	
	hash, for remote monitors		set sh calling-convention	
	height		set shell	
	history expansion		set signal-thread	
	history filename		set signals, Hurd command	
	· · · · · · · · · · · · · · · · · · ·		set sigs, Hurd command	
	history save			
	history size		set sigthread	
	host-charset		set solib-absolute-prefix	
	inferior controlling terminal		set solib-search-path	
	inferior-tty		•	
	input-radix		set stack-cache	
	interactive-mode		set step-mode	
	language		set stop-on-solib-events	
set	libthread-db-search-path	. 37	set stopped, Hurd command	
	listsize		set struct-convention	
set	logging	. 17	$\verb set substitute-path $. 88
set	mach-exceptions	227	set symbol-reloading	178
	max-user-call-depth		set syn-garbage-limit, MIPS remote	

set sysroot	show com3base	221
·		
set target-async	show com3irq	
set target-charset	show com4base	
set target-wide-charset	show com4irq	
set task, Hurd commands	show commands	
set tcp	show complaints	250
set tdesc filename	show confirm	250
set thread, Hurd command 225	show convenience	109
set timeout	show copying	. 23
set trace-commands	show cp-abi	249
set tracepoint	show cygwin-exceptions	
set trust-readonly-sections	show data-directory	
set tui active-border-mode	show debug	
set tui border-kind	show debug darwin	
set tui border-mode	show debug mach-o	227
set unwind-on-terminating-exception 185	show debug mips	
	show debug monitor	
set unwindonsignal		
set variable	show debug nto-debug	
set verbose	show debug-file-directory	
set watchdog	show default-collect	
set width	show detach-on-fork	
set write	show directories	
set-mark (C-@)	show disassemble-next-line	
set_debug_traps	show disassembly-flavor	. 90
setting variables	show disconnected-tracing	135
setting watchpoints	show displaced-stepping	397
SH	show download-path	
'sh-stub.c'	show editing	
share	show environment	
shared libraries	show exceptions, Hurd command	
shared library events, remote reply 414	show exec-done-display	
sharedlibrary	show follow-fork-mode	
shell	show gnutarget	
shell escape	show hash, for remote monitors	
show	show height	
show ada trust-PAD-over-XVS	show history	
show all convenience functions	show history show host-charset show host-charset	
show all user variables	show inferior-tty	
show annotate	show input-radix	
show architecture	show interactive-mode	
show args	show language	
show arm	show last commands	
show auto-solib-add	show libthread-db-search-path	. 37
show backtrace 79	show listsize	
show board-address	show logging	
show breakpoint always-inserted	show mach-exceptions	
show breakpoint auto-hw	show max-user-call-depth	256
show breakpoint pending 48	show mem inaccessible-by-default	115
show can-use-hw-watchpoints 50	show mips abi	242
show case-sensitive	show mips compression	
show charset	show mips mask-address	
show check range	show mipsfpu	
show check type	show monitor-prompt, MIPS remote	
show coerce-float-to-double	show monitor-warnings, MIPS remote	
show com1base	show multiple-symbols	
show com1irq	show new-console	
show com2base	show new-group	
show com2irg	show non-stop	
	D110 # 11011 DUOP	. 01

show	opaque-type-resolution	show verbose	250
show	osabi	show version	. 23
show	output-radix	show warranty	. 23
show	overload-resolution	show width	247
show	pagination	show write	186
show	paths	show-all-if-ambiguous	371
show	print 102	show-all-if-unmodified	371
show	print inferior-events 34	si (stepi)	. 63
show	print thread-events	signal	183
show	processor 199	signal annotation	358
	procfs-file	signal-name annotation	357
	procfs-trace	signal-name-end annotation	357
show	prompt	signal-string annotation	357
	python auto-load	signal-string-end annotation	
show	radix	signalled annotation	
show	rdiheartbeat	signals	. 63
show	rdiromatzero	SIGQUIT signal, dump core of GDB	398
show	record insn-number-max	silent	. 58
show	record stop-at-limit	sim	239
	remote	sim, a command	229
show	remote system-call-allowed 445	simulator, Z8000	239
show	remote-mips64-transfers-32bit-regs	size of remote memory accesses	
	242	size of screen	247
show	remotecache	sizeof of Type	266
show	remoteflow	snapshot of a process	. 40
show	retransmit-timeout	software watchpoints	49
show	rstack_high_address 241	source	258
show	script-extension	source annotation	358
show	sdstimeout	source file and line of a symbol	103
show	server-address 231	source line and its code address	. 89
show	sh calling-convention	source path	. 86
show	shell	Sparc	
show	signal-thread	'sparc-stub.c'	213
show	signals, Hurd command	'sparcl-stub.c'	213
show	sigs, Hurd command 224	Sparclet	237
show	sigthread	SparcLite	
show	solib-search-path	Special Fortran commands	159
show	spu	specifying location	. 84
show	stack-cache	spr	
show	stop-on-solib-events	SPU	243
	stopped, Hurd command	SSE registers (x86)	111
	struct-convention	stack frame	
show	substitute-path 88	stack on Alpha	241
	symbol-reloading	stack on MIPS	
show	syn-garbage-limit, MIPS remote 234	stack pointer register	111
show	sysroot	stacking targets	199
show	target-async	standard registers	111
show	target-charset 117	start	26
show	target-wide-charset	start a new trace experiment	134
	task, Hurd commands	start-kbd-macro (C-x ()	381
show	$\verb tcp$	starting	
	tdesc filename	starting annotation	
show	thread, Hurd command	startup code, and backtrace	. 79
show	timeout	stat, file-i/o system call	
show	${\tt unwind-on-terminating-exception} \dots 185$	static members of C++ objects	
show	unwindonsignal	static members of Pascal objects	107
show	$\verb"user" \dots \dots$	status of trace data collection	
show	values	status output in GDB/MI	294

step 61	target character set	116
step&	target dbug	232
stepi	target ddb port	
stepi&	target debugging info	
stepping	target descriptions	
stepping into functions with no line info 61	target descriptions, ARM features	
stop a running trace experiment	target descriptions, i386 features	
stop on C++ exceptions	target descriptions, inclusion	
stop reply packets	target descriptions, M68K features	
stop, a pseudo-command	target descriptions, MIPS features	
stopped threads	target descriptions, PowerPC features	
stopping annotation	target descriptions, predefined types	
stream records in GDB/MI	target descriptions, standard features	
string on Value	target descriptions, XML format	
strip_typedefs on Type	target dink32	
struct return convention	target jtag	
struct stat, in file-i/o protocol	target lsi port	
struct timeval, in file-i/o protocol	target m32r	
	target m32rsdi	
struct/union returned in registers	target mips port	
structure field name completion	target on Type	
stub example, remote debugging	target op50n	
Super-H	target output in GDB/MI	
supported packets, remote query	target pmon port	
switching threads	target ppcbug	
switching threads automatically	target ppcbug1target r3900	
symbol decoding style, C++	_	
symbol dump	target rditarget rdp	
symbol from address	target record	
symbol lookup, remote request	target remote	
symbol names	target sds	
symbol table	target sim, with Z8000	
symbol tables, listing GDB's internal 179	target sparclite	
symbol, source file and line 103	target stack description	
symbol-file	target tfile	
symbols, reading from relocatable object files 189	target vxworks	
symbols, reading immediately 188	target w89k	
synchronize with remote MIPS target 234	task (Ada)	
syscall DSO	task attributes (GNU Hurd)	
sysinfo	task breakpoints, in Ada	
system calls and thread breakpoints	task exception port, GNU Hurd	
system root, alternate	task suspend count	
system, file-i/o system call	tbreak	
system-wide init file	TCP port, target remote	203
	tdump	
TT.	template_argument on Type	
${f T}$	terminal	
't' packet	teval (tracepoints)	133
'T' packet	Text User Interface	281
'T' packet reply	tfile	139
tab-insert (M-TAB)	tfind	136
tabset	thbreak	
tag of Type	this, inside C++ member functions	
target	thread apply	
target architecture	thread attributes info, remote request	
target array	thread breakpoints	69
target byte order	thread breakpoints and system calls	69

thread default settings, GNU Hurd 226	tvariable	32
thread identifier (GDB)	type casting memory	94
thread identifier (GDB), on HP-UX	type chain of a data type 4	00
thread identifier (system)	type checking 14	49
thread identifier (system), on HP-UX 36	type conversions in C++	54
thread info (Solaris)	type of LazyString	78
thread information, remote request 418	type of Value 2	64
thread list format	type on Frame	77
thread number	TYPE_CODE_ARRAY	68
thread properties, GNU Hurd	TYPE_CODE_BITSTRING 20	68
thread suspend count, GNU Hurd 226	TYPE_CODE_BOOL	69
thread threadno	TYPE_CODE_CHAR	69
thread-id, in remote protocol	TYPE_CODE_COMPLEX 20	69
threads and watchpoints	TYPE_CODE_DECFLOAT 2	69
threads of execution	TYPE_CODE_ENUM	68
threads, automatic switching	TYPE_CODE_ERROR	68
threads, continuing	TYPE_CODE_FLAGS	68
threads, stopped	TYPE_CODE_FLT	68
tilde-expand (M-~)	TYPE_CODE_FUNC	68
time of command execution	TYPE_CODE_INT	68
timeout for commands	TYPE_CODE_INTERNAL_FUNCTION 20	69
timeout for serial communications 209	TYPE_CODE_MEMBERPTR 20	69
timeout, for remote target connection 210	TYPE_CODE_METHOD	68
timeout, MIPS protocol	TYPE_CODE_METHODPTR 20	69
timestampping debugging info	TYPE_CODE_NAMESPACE 20	69
tload, M32R	TYPE_CODE_PTR	68
to_string on pretty printer	TYPE_CODE_RANGE	68
trace	TYPE_CODE_REF	69
trace experiment, status of	TYPE_CODE_SET	68
trace file format	TYPE_CODE_STRING	68
trace files	TYPE_CODE_STRUCT	68
trace state variable value, remote request 431	TYPE_CODE_TYPEDEF	69
trace state variables	TYPE_CODE_UNION	
traceback	TYPE_CODE_VOID	68
tracepoint actions	types in Python	66
tracepoint conditions		
tracepoint data, display	TT	
tracepoint deletion	\mathbf{U}	
tracepoint number	u (SingleKey TUI key)	83
tracepoint packets	u (until)	
tracepoint pass count	UDP port, target remote	
tracepoint variables	undisplay	
tracepoints	undo (C or C-x C-u)	
trailing underscore, in Fortran symbols 158	union field name completion	
translating between character sets	unions in structures, printing 10	
transpose-chars (C-t)	universal-argument ()	
transpose-words (M-t) 379	unix-filename-rubout ()	
tsave	unix-line-discard (C-u)	80
tstart	unix-word-rubout (C-w)	
tstatus	unknown address, locating	
tstop	unlink, file-i/o system call 4	
tty	unlinked object files	
TÜI	unload symbols from shared libraries 1	
TUI commands	unmap an overlay	
TUI configuration variables	unmapped overlays	
TUI key bindings	unqualified on Type	
tui reg	unset environment	
TUI single key mode	unset substitute-path	
<u> </u>	*	

unset tdesc filename	\mathbf{W}
unsupported languages	r (Single Very TIII Iron)
unsupported packets, empty response for 404	w (SingleKey TUI key) 285 watch 49
until 62	watchdog timer
until&	watchpoint annotation
unwind stack in called functions	watchpoints
unwind stack in called functions with unhandled	watchpoints and threads
exceptions	weak alias functions
unwind_stop_reason on Frame	whatis
up	where
Up	where to look for shared libraries
up-silently 81	while
upcase-word (M-u)	while-stepping (tracepoints) 133
update	wild pointer, interpreting
upload, M32R	winheight
use only software watchpoints 50	word completion
use_dbt_break	working directory
use_debug_dma	working directory (of your program) 30
use_ib_break	working language
use_mon_code	write
user-defined command	write data into object, remote request 427
user-defined macros	write, file-i/o system call
user-defined variables	writing convenience functions
	writing into corefiles
${f V}$	writing into executables 185
•	wrong values
v (SingleKey TUI key)	
value history	37
value on LazyString	\mathbf{X}
value optimized out, in backtrace	x (examine memory)
values from inferior, with Python	x command, default address 89
variable object debugging info	'X' packet
variable objects in GDB/MI	x(examine), and info line 89
variable values, wrong	Xilinx MicroBlaze
variables, readline	XInclude
variables, setting	XMD, Xilinx Microprocessor Debugger 232
'vAttach' packet	XML parser debugging 253
'vCont' packet	
'vCont?' packet	T 7
vector unit	Y
vector, auxiliary	yank (C-y)
verbose operation	yank-last-arg (M or M) 378
verify remote memory image 100	yank-nth-arg (M-C-y) 378
'vFile' packet	yank-pop (M-y)
'vFlashDone' packet	yanking text
'vFlashErase' packet	
'vFlashWrite' packet	
vi-editing-mode (M-C-j)	${f Z}$
virtual functions (C++) display 108	'z' packet
visible-stats	'Z' packets
'vKill' packet	'z0' packet
volatile on Type	'Z0' packet
'vRun' packet	'z1' packet
'vStopped' packet	'Z1' packet
VTBL display	'z2' packet
VxWorks	'Z2' packet
vxworks-timeout	'z3' packet

'Z3'	packet	413	Z8000	239
'z4'	packet	413		
'Z4'	packet	413	Zilog Z8000 simulator	239

The body of this manual is set in cmr10 at 10.95pt, with headings in cmb10 at 10.95pt and examples in cmtt10 at 10.95pt. cmti10 at 10.95pt, cmb10 at 10.95pt, and cmsl10 at 10.95pt are used for emphasis.