

Sourcery G++ Lite

ARM EABI

Sourcery G++ Lite 2009q3-68

Getting Started



Sourcery G++ Lite: ARM EABI: Sourcery G++ Lite 2009q3-68: Getting Started

CodeSourcery, Inc.

Copyright © 2005, 2006, 2007, 2008, 2009 CodeSourcery, Inc.

All rights reserved.

Abstract

This guide explains how to install and build applications with Sourcery G++ Lite, CodeSourcery's customized, validated, and supported version of the GNU Toolchain. Sourcery G++ Lite includes everything you need for application development, including C and C++ compilers, assemblers, linkers, and libraries.

When you have finished reading this guide, you will know how to use Sourcery G++ from the command line.

Table of Contents

Preface	v
1. Intended Audience	vi
2. Organization	vi
3. Typographical Conventions	vii
1. Quick Start	1
1.1. Installation and Set-Up	2
1.2. Configuring Sourcery G++ Lite for the Target System	2
1.3. Building Your Program	2
1.4. Running and Debugging Your Program	2
2. Installation and Configuration	4
2.1. Terminology	5
2.2. System Requirements	5
2.3. Downloading an Installer	6
2.4. Installing Sourcery G++ Lite	6
2.5. Installing Sourcery G++ Lite Updates	9
2.6. Setting up the Environment	9
2.7. Uninstalling Sourcery G++ Lite	11
3. Sourcery G++ Lite for ARM EABI	13
3.1. Included Components and Features	14
3.2. Library Configurations	14
3.3. Using Flash Memory	15
3.4. ARMv7-M Interrupt Handlers	16
3.5. ARM Interrupt Vectors	16
3.6. Using VFP Floating Point	16
3.7. ABI Compatibility	17
3.8. Object File Portability	18
3.9. ARM Profiling Implementation	18
4. Using Sourcery G++ from the Command Line	20
4.1. Building an Application	21
4.2. Running Applications on the Target System	21
4.3. Running Applications in the Simulator	21
4.4. Running Applications from GDB	22
5. CS3™: The CodeSourcery Common Startup Code Sequence	24
5.1. Startup Sequence	25
5.2. Exit and Embedded Systems	27
5.3. Memory Layout	27
5.4. Interrupt Vectors and Handlers	29
5.5. Linker Scripts	31
5.6. Supported Boards for ARM EABI	33
5.7. Interrupt Vector Tables	34
5.8. Regions and Memory Sections	34
6. Sourcery G++ Debug Sprite	36
6.1. Probing for Debug Devices	37
6.2. Debug Sprite Example	37
6.3. Invoking Sourcery G++ Debug Sprite	38
6.4. Sourcery G++ Debug Sprite Options	39
6.5. Remote Debug Interface Devices	39
6.6. Actel FlashPro Devices	40
6.7. Altera Devices	40
6.8. Debugging a Remote Board	41
6.9. Supported Board Files	42

6.10. Board File Syntax	42
7. Next Steps with Sourcery G++	46
7.1. Sourcery G++ Subscriptions	47
7.2. Sourcery G++ Knowledge Base	48
7.3. Manuals for GNU Toolchain Components	48
A. Sourcery G++ Lite Release Notes	50
A.1. Changes in Sourcery G++ Lite for ARM EABI	51
B. Sourcery G++ Lite Licenses	65
B.1. Licenses for Sourcery G++ Lite Components	66
B.2. Sourcery G++ Software License Agreement	67

Preface

This preface introduces the Sourcery G++ Lite Getting Started guide. It explains the structure of this guide and describes the documentation conventions used.

1. Intended Audience

This guide is written for people who will install and/or use Sourcery G++ Lite. This guide provides a step-by-step guide to installing Sourcery G++ Lite and to building simple applications. Parts of this document assume that you have some familiarity with using the command-line interface.

2. Organization

This document is organized into the following chapters and appendices:

Chapter 1, “Quick Start”	This chapter includes a brief checklist to follow when installing and using Sourcery G++ Lite for the first time. You may use this chapter as an abbreviated guide to the rest of this manual.
Chapter 2, “Installation and Configuration”	This chapter describes how to download, install and configure Sourcery G++ Lite. This section describes the available installation options and explains how to set up your environment so that you can build applications.
Chapter 3, “Sourcery G++ Lite for ARM EABI”	This chapter contains information about using Sourcery G++ Lite that is specific to ARM EABI targets. You should read this chapter to learn how to best use Sourcery G++ Lite on your target system.
Chapter 4, “Using Sourcery G++ from the Command Line”	This chapter explains how to build applications with Sourcery G++ Lite using the command line. In the process of reading this chapter, you will build a simple application that you can use as a model for your own programs.
Chapter 5, “CS3™: The CodeSourcery Common Startup Code Sequence”	CS3 is CodeSourcery's low-level board support library. This chapter describes the organization of the system startup code and tells you how you can customize it, such as by defining your own interrupt handlers. This chapter also documents the boards supported by Sourcery G++ Lite and the compiler and linker options you need to use with them.
Chapter 6, “Sourcery G++ Debug Sprite”	This chapter describes the use of the Sourcery G++ Debug Sprite for remote debugging. The Sprite allows you to debug programs running on a bare board without an operating system. This chapter includes information about the debugging devices and boards supported by the Sprite for ARM EABI.
Chapter 7, “Next Steps with Sourcery G++”	This chapter describes where you can find additional documentation and information about using Sourcery G++ Lite and its components. It also provides information about Sourcery G++ subscriptions. CodeSourcery customers with Sourcery G++ subscriptions receive comprehensive support for Sourcery G++.
Appendix A, “Sourcery G++ Lite Release Notes”	This appendix contains information about changes in this release of Sourcery G++ Lite for ARM EABI. You should read through these notes to learn about new features and bug fixes.

Appendix B, “Sourcery G++ Lite Licenses” This appendix provides information about the software licenses that apply to Sourcery G++ Lite. Read this appendix to understand your legal rights and obligations as a user of Sourcery G++ Lite.

3. Typographical Conventions

The following typographical conventions are used in this guide:

<code>> command arg ...</code>	A command, typed by the user, and its output. The “>” character is the command prompt.
command	The name of a program, when used in a sentence, rather than in literal input or output.
<code>literal</code>	Text provided to or received from a computer program.
<code>placeholder</code>	Text that should be replaced with an appropriate value when typing a command.
<code>\</code>	At the end of a line in command or program examples, indicates that a long line of literal input or output continues onto the next line in the document.

Chapter 1

Quick Start

This chapter includes a brief checklist to follow when installing and using Sourcery G++ Lite for the first time. You may use this chapter as an abbreviated guide to the rest of this manual.

Sourcery G++ Lite for ARM EABI is intended for developers working on embedded applications or firmware for boards without an operating system, or that run an RTOS or boot loader. This Sourcery G++ configuration is not intended for Linux or uClinux kernel or application development.

Follow the steps given in this chapter to install Sourcery G++ Lite and build and run your first application program. The checklist given here is not a tutorial and does not include detailed instructions for each step; however, it will help guide you to find the instructions and reference information you need to accomplish each step.

You can find additional details about the components, libraries, and other features included in this version of Sourcery G++ Lite in Chapter 3, “Sourcery G++ Lite for ARM EABI”.

1.1. Installation and Set-Up

Install Sourcery G++ Lite on your host computer. You may download an installer package from the Sourcery G++ web site¹, or you may have received an installer on CD. The installer is an executable program that pops up a window on your computer and leads you through a series of dialogs to configure your installation. If the installation is successful, it will offer to launch the Getting Started guide. For more information about installing Sourcery G++ Lite, including host system requirements and tips to set up your environment after installation, refer to Chapter 2, “Installation and Configuration”.

Install drivers for your debug device. If you plan to use the Sourcery G++ Debug Sprite, you may need to install drivers, libraries, or other software on your host system. Refer to Chapter 6, “Sourcery G++ Debug Sprite” for a list of supported devices and information about installing drivers and other device set-up. Sourcery G++ Lite also supports third-party debug devices that communicate via the GDB remote serial protocol. If you plan to use one of these devices, follow the manufacturer's directions to connect the device and install any required drivers or software.

1.2. Configuring Sourcery G++ Lite for the Target System

Identify your target board. On bare-metal targets, you must explicitly specify a linker script for your target board on your link command line. Supported boards are listed in Chapter 5, “CS3™: The CodeSourcery Common Startup Code Sequence”. You can also choose a simulator as your target board.

1.3. Building Your Program

Build your program with Sourcery G++ command-line tools. Create a simple test program, and follow the directions in Chapter 4, “Using Sourcery G++ from the Command Line” to compile and link it using Sourcery G++ Lite. On bare-metal targets, you must specify a linker script using the `-T` option on your link command line. Supported boards and linker scripts are listed in Chapter 5, “CS3™: The CodeSourcery Common Startup Code Sequence”.

1.4. Running and Debugging Your Program

The steps to run or debug your program depend on your target system and how it is configured. Choose the appropriate method for your target.

¹ http://www.codesourcery.com/gnu_toolchains/

Run or debug your program in the simulator. Sourcery G++ Lite includes an instruction-set simulator, which provides an easy way to run or debug your program without requiring target hardware. The simulator can be run directly from the command line (see Section 4.3, “Running Applications in the Simulator”) or via the debugger (see Section 4.4, “Running Applications from GDB”).

Debug your program on the target using the Debug Sprite. You can use the Sourcery G++ Debug Sprite to load and execute your program on the target from the debugger. Refer to Section 4.4, “Running Applications from GDB” for instructions on using the Sprite from the GDB command line. Detailed reference material for the Sourcery G++ Debug Sprite, including information about supported debug devices, can be found in Chapter 6, “Sourcery G++ Debug Sprite”.

Debug your program on the target using a third-party debug device. Sourcery G++ supports debugging programs on the remote target using third-party debug devices that can communicate via the GDB remote serial protocol. For command-line GDB instructions, see Section 4.4, “Running Applications from GDB”.

Chapter 2

Installation and Configuration

This chapter explains how to install Sourcery G++ Lite. You will learn how to:

1. Verify that you can install Sourcery G++ Lite on your system.
2. Download the appropriate Sourcery G++ Lite installer.
3. Install Sourcery G++ Lite.
4. Configure your environment so that you can use Sourcery G++ Lite.

2.1. Terminology

Throughout this document, the term *host system* refers to the system on which you run Sourcery G++ while the term *target system* refers to the system on which the code produced by Sourcery G++ runs. The target system for this version of Sourcery G++ is `arm-none-eabi`.

If you are developing a workstation or server application to run on the same system that you are using to run Sourcery G++, then the host and target systems are the same. On the other hand, if you are developing an application for an embedded system, then the host and target systems are probably different.

2.2. System Requirements

2.2.1. Host Operating System Requirements

This version of Sourcery G++ supports the following host operating systems and architectures:

- Microsoft Windows NT 4, Windows 2000, Windows XP, and Windows Vista systems using IA32, AMD64, and EM64T processors.
- GNU/Linux systems using IA32, AMD64, or EM64T processors, including Debian 3.1 (and later), Red Hat Enterprise Linux 3 (and later), and SuSE Enterprise Linux 8 (and later).

Sourcery G++ is built as a 32-bit application. Therefore, even when running on a 64-bit host system, Sourcery G++ requires 32-bit host libraries. If these libraries are not already installed on your system, you must install them before installing and using Sourcery G++ Lite. Consult your operating system documentation for more information about obtaining these libraries.

Installing on Ubuntu and Debian GNU/Linux Hosts

The Sourcery G++ graphical installer is incompatible with the **dash** shell, which is the default `/bin/sh` for recent releases of the Ubuntu and Debian GNU/Linux distributions. To install Sourcery G++ Lite on these systems, you must make `/bin/sh` a symbolic link to one of the supported shells: **bash**, **csh**, **tcsh**, **zsh**, or **ksh**.

For example, on Ubuntu systems, the recommended way to do this is:

```
> sudo dpkg-reconfigure -pflow dash
Install as /bin/sh? No
```

This is a limitation of the installer and uninstaller only, not of the installed Sourcery G++ Lite toolchain.

2.2.2. Host Hardware Requirements

In order to install and use Sourcery G++ Lite, you must have at least 128MB of available memory.

The amount of disk space required for a complete Sourcery G++ Lite installation directory depends on the host operating system and the number of target libraries included. Typically, you should plan on at least 400MB.

In addition, the graphical installer requires a similar amount of temporary space during the installation process. On Microsoft Windows hosts, the installer uses the location specified by the `TEMP` environment variable for these temporary files. If there is not enough free space on that volume, the installer

prompts for an alternate location. On Linux hosts, the installer puts temporary files in the directory specified by the `IATEMPDIR` environment variable, or `/tmp` if that is not set.

2.2.3. Target System Requirements

See Chapter 3, “Sourcery G++ Lite for ARM EABI” for requirements that apply to the target system.

2.3. Downloading an Installer

If you have received Sourcery G++ Lite on a CD, or other physical media, then you do not need to download an installer. You may skip ahead to Section 2.4, “Installing Sourcery G++ Lite”.

You can download Sourcery G++ Lite from the Sourcery G++ web site¹. This free version of Sourcery G++, which is made available to the general public, does not include all the functionality of CodeSourcery's product releases. If you prefer, you may instead purchase or register for an evaluation of CodeSourcery's product toolchains at the Sourcery G++ Portal². For more information about subscriptions for Sourcery G++ product releases, see Section 7.1, “Sourcery G++ Subscriptions”.

Once you have navigated to the appropriate web site, download the installer that corresponds to your host operating system. For Microsoft Windows systems, the Sourcery G++ installer is provided as an executable with the `.exe` extension. For GNU/Linux systems Sourcery G++ Lite is provided as an executable installer package with the `.bin` extension. You may also install from a compressed archive with the `.tar.bz2` extension.

On Microsoft Windows systems, save the installer to the desktop. On GNU/Linux systems, save the download package in your home directory.

2.4. Installing Sourcery G++ Lite

The method used to install Sourcery G++ Lite depends on your host system and the kind of installation package you have downloaded.


2.4.1. Using the Sourcery G++ Lite Installer on Microsoft Windows

If you have received Sourcery G++ Lite on CD, insert the CD in your computer. On most computers, the installer then starts automatically. If your computer has been configured not to automatically run CDs, open *My Computer*, and double click on the CD. If you downloaded Sourcery G++ Lite, double-click on the installer.

After the installer starts, follow the on-screen dialogs to install Sourcery G++ Lite. The installer is intended to be self-explanatory and on most pages the defaults are appropriate.


¹ http://www.codesourcery.com/gnu_toolchains/

² <https://support.codesourcery.com/GNUToolchain/>




Running the Installer. The graphical installer guides you through the steps to install Sourcery G++ Lite.

You may want to change the install directory pathname and customize the shortcut installation.




Choose Install Folder. Select the pathname to your install directory.



Choose Shortcut Folder. You can customize where the installer creates shortcuts for quick access to Sourcery G++ Lite.

When the installer has finished, it asks if you want to launch a viewer for the Getting Started guide. Finally, the installer displays a summary screen to confirm a successful install before it exits.



Install Complete. You should see a screen similar to this after a successful install.

If you prefer, you can run the installer in console mode rather than using the graphical interface. To do this, invoke the installer with the `-i console` command-line option. For example:

```
> /path/to/package.exe -i console
```

2.4.2. Using the Sourcery G++ Lite Installer on GNU/Linux Hosts

Start the graphical installer by invoking the executable shell script:

```
> /bin/sh ./path/to/package.bin
```

After the installer starts, follow the on-screen dialogs to install Sourcery G++ Lite. For additional details on running the installer, see the discussion and screen shots in the Microsoft Windows section above.

If you prefer, or if your host system does not run the X Window System, you can run the installer in console mode rather than using the graphical interface. To do this, invoke the installer with the `-i console` command-line option. For example:

```
> /bin/sh ./path/to/package.bin -i console
```

2.4.3. Installing Sourcery G++ Lite from a Compressed Archive

You do not need to be a system administrator to install Sourcery G++ Lite from a compressed archive. You may install Sourcery G++ Lite using any user account and in any directory to which you have write access. This guide assumes that you have decided to install Sourcery G++ Lite in the `$HOME/CodeSourcery` subdirectory of your home directory and that the filename of the package you have downloaded is `/path/to/package.tar.bz2`. After installation the toolchain will be in `$HOME/CodeSourcery/sourceryg++-2009q3`.

First, uncompress the package file:

```
> bunzip2 /path/to/package.tar.bz2
```

Next, create the directory in which you wish to install the package:

```
> mkdir -p $HOME/CodeSourcery
```

Change to the installation directory:

```
> cd $HOME/CodeSourcery
```

Unpack the package:

```
> tar xf /path/to/package.tar
```

2.5. Installing Sourcery G++ Lite Updates

If you have already installed an earlier version of Sourcery G++ Lite for ARM EABI on your system, it is not necessary to uninstall it before using the installer to unpack a new version in the same location. The installer detects that it is performing an update in that case.

If you are installing an update from a compressed archive, it is recommended that you remove any previous installation in the same location, or install in a different directory.

Note that the names of the Sourcery G++ commands for the ARM EABI target all begin with **arm-none-eabi**. This means that you can install Sourcery G++ for multiple target systems in the same directory without conflicts.

2.6. Setting up the Environment

As with the installation process itself, the steps required to set up your environment depend on your host operating system.

2.6.1. Setting up the Environment on Microsoft Windows Hosts

2.6.1.1. Setting the PATH

In order to use the Sourcery G++ tools from the command line, you should add them to your PATH. You may skip this step if you used the graphical installer, since the installer automatically adds Sourcery G++ to your PATH.

To set the PATH on a Microsoft Windows Vista system, use the following command in a `cmd.exe` shell:

```
> setx PATH "%PATH%;C:\Program Files\Sourcery G++\bin"
```

where `C:\Program Files\Sourcery G++` should be changed to the path of your Sourcery G++ Lite installation.

To set the PATH on a system running a Microsoft Windows version other than Vista, from the desktop bring up the Start menu and right click on My Computer. Select Properties, go to the Advanced tab, then click on the Environment Variables button. Select the PATH variable and click the Edit. Add the string `;C:\Program Files\Sourcery G++\bin` to the end, and click OK. Again, you must adjust the pathname to reflect your installation directory.

You can verify that your PATH is set up correctly by starting a new `cmd.exe` shell and running:

```
> arm-none-eabi-g++ -v
```

Verify that the last line of the output contains: `Sourcery G++ Lite 2009q3-68`.

2.6.1.2. Working with Cygwin

Sourcery G++ Lite does not require Cygwin or any other UNIX emulation environment. You can use Sourcery G++ directly from the Windows command shell. You can also use Sourcery G++ from within the Cygwin environment, if you prefer.

The Cygwin emulation environment translates Windows path names into UNIX path names. For example, the Cygwin path `/home/user/hello.c` corresponds to the Windows path `c:\cygwin\home\user\hello.c`. Because Sourcery G++ is not a Cygwin application, it does not, by default, recognize Cygwin paths.

If you are using Sourcery G++ from Cygwin, you should set the `CYGPATH` environment variable. If this environment variable is set, Sourcery G++ Lite automatically translates Cygwin path names into Windows path names. To set this environment variable, type the following command in a Cygwin shell:

```
> export CYGPATH=cygpath
```

To resolve Cygwin path names, Sourcery G++ relies on the `cygpath` utility provided with Cygwin. You must provide Sourcery G++ with the full path to `cygpath` if `cygpath` is not in your PATH. For example:

```
> export CYGPATH=c:/cygwin/bin/cygpath
```

directs Sourcery G++ Lite to use `c:/cygwin/bin/cygpath` as the path conversion utility. The value of `CYGPATH` must be an ordinary Windows path, not a Cygwin path.

2.6.2. Setting up the Environment on GNU/Linux Hosts

If you installed Sourcery G++ Lite using the graphical installer then you may skip this step. The installer does this setup for you.

Before using Sourcery G++ Lite you should add it to your `PATH`. The command you must use varies with the particular command shell that you are using. If you are using the C Shell (**cs**h or **tc**sh), use the command:

```
> setenv PATH $HOME/CodeSourcery/Sourcery_G++/bin:$PATH
```

If you are using Bourne Shell (**sh**), the Korn Shell (**ksh**), or another shell, use:

```
> PATH=$HOME/CodeSourcery/Sourcery_G++/bin:$PATH
> export PATH
```

If you are not sure which shell you are using, try both commands. In both cases, if you have installed Sourcery G++ Lite in an alternate location, you must replace the directory above with `bin` subdirectory of the directory in which you installed Sourcery G++ Lite.

You may also wish to set the `MANPATH` environment variable so that you can access the Sourcery G++ manual pages, which provide additional information about using Sourcery G++. To set the `MANPATH` environment variable, follow the same steps shown above, replacing `PATH` with `MANPATH`, and `bin` with `share/doc/sourceryg++-arm-none-eabi/man`.

You can test that your `PATH` is set up correctly by running the following command:

```
> arm-none-eabi-g++ -v
```

Verify that the last line of the output contains: `Sourcery G++ Lite 2009q3-68`.

2.7. Uninstalling Sourcery G++ Lite

The method used to uninstall Sourcery G++ Lite depends on the method you originally used to install it. If you have modified any files in the installation it is recommended that you back up these changes. The uninstall procedure may remove the files you have altered.

2.7.1. Using the Sourcery G++ Lite Uninstaller on Microsoft Windows

For Windows hosts other than Microsoft Windows Vista, select `Start`, then `Control Panel`. Select `Add or Remove Programs`. Scroll down and click on `Sourcery G++ for ARM EABI`. Select `Change/Remove` and follow the on-screen dialogs to uninstall Sourcery G++ Lite.

On Microsoft Windows Vista hosts, select `Start`, then `Settings` and finally `Control Panel`. Select the `Uninstall a program` task. Scroll down and double click on `Sourcery G++ for ARM EABI`. Follow the on-screen dialogs to uninstall Sourcery G++ Lite.

You can run the uninstaller in console mode, rather than using the graphical interface, by invoking the `Uninstall` executable found in your Sourcery G++ Lite installation directory with the `-i console` command-line option.

To uninstall third-party drivers bundled with Sourcery G++ Lite, first disconnect the associated hardware device. Then use `Add or Remove Programs` (non-Vista) or `Uninstall a program` (Vista) to remove the drivers separately. Depending on the device, you may need to reboot your computer to complete the driver uninstall.

2.7.2. Using the Sourcery G++ Lite Uninstaller on GNU/Linux

You should use the provided uninstaller to remove a Sourcery G++ Lite installation originally created by the executable installer script. The `arm-none-eabi` directory located in the `install` directory will be removed entirely by the uninstaller. Please back up any changes you have made to this directory, such as modified linker scripts.

Start the graphical uninstaller by invoking the executable `Uninstall` shell script located in your installation directory. After the uninstaller starts, follow the on-screen dialogs to uninstall Sourcery G++ Lite.

You can run the uninstaller in console mode, rather than using the graphical interface, by invoking the `Uninstall` script with the `-i console` command-line option.

2.7.3. Uninstalling a Compressed Archive Installation

If you installed Sourcery G++ Lite from a `.tar.bz2` file, you can uninstall it by manually deleting the installation directory created in the `install` procedure.

Chapter 3

Sourcery G++ Lite for ARM EABI

This chapter contains information about features of Sourcery G++ Lite that are specific to ARM EABI targets. You should read this chapter to learn how to best use Sourcery G++ Lite on your target system.

3.1. Included Components and Features

This section briefly lists the important components and features included in Sourcery G++ Lite for ARM EABI, and tells you where you may find further information about these features.

Component	Version	Notes
GNU programming tools		
GNU Compiler Collection	4.4.1	Separate manual included.
GNU Binary Utilities	2.19.51	Includes assembler, linker, and other utilities. Separate manuals included.
Debugging support and simulators		
GNU Debugger	6.8.50	Separate manual included.
Sourcery G++ Debug Sprite for ARM	2009q3-68	See Chapter 6, “Sourcery G++ Debug Sprite”.
GDB Simulator	N/A	See Section 4.3, “Running Applications in the Simulator”.
Target libraries		
CodeSourcery Common Startup Code Sequence	2009q3-68	See Chapter 5, “CS3™: The CodeSourcery Common Startup Code Sequence”.
Newlib C Library	1.17.0	Separate manuals included.
Other utilities		
GNU Make	N/A	Build support on Windows hosts.
GNU Core Utilities	N/A	Build support on Windows hosts.

3.2. Library Configurations

Sourcery G++ includes copies of run-time libraries that have been built with optimizations for different target architecture variants or other sets of build options. Each such set of libraries is referred to as a *multilib*. When you link a target application, Sourcery G++ selects the multilib matching the build options you have selected.

Sourcery G++ Lite includes linker scripts as well as runtime libraries for each multilib. You can find these files in multilib-specific subdirectories of the `arm-none-eabi/lib` directory of your Sourcery G++ install.

3.2.1. Included Libraries

The following library configurations are available in Sourcery G++ Lite for ARM EABI.

ARMv4 - Little-Endian, Soft-Float	
Command-line option(s):	default
Library subdirectory:	./

ARMv4 Thumb - Little-Endian, Soft-Float	
Command-line option(s):	-mthumb
Library subdirectory:	tthumb/

ARMv7 Thumb-2 - Little-Endian, Soft-Float	
Command-line option(s):	<code>-mthumb -march=armv7 -mfix-cortex-m3-ldrd</code>
Library subdirectory:	<code>thumb2/</code>

ARMv6-M Thumb - Little-Endian, Soft-Float	
Command-line option(s):	<code>-mthumb -march=armv6-m</code>
Library subdirectory:	<code>armv6-m/</code>

3.2.2. Library Selection

A given multilib may be compatible with additional processors and build options beyond those listed above. However, even if a particular set of command-line options produces code compatible with one of the provided multilibs, those options may not be sufficient to identify the intended library to the linker. For example, on some targets, specifying only a processor option on the command line may imply architecture features or floating-point support for compilation, but not for library selection. The details of the mapping from command-line options to multilibs are target-specific and quite complex. Therefore, it is recommended that your link command line include exactly the options listed in the tables above for your intended target multilib. In some cases, you may need to supply different options for linking than for compilation.

If you are uncertain which multilib is selected by a particular set of command-line options, GCC can tell you if you invoke it with the `-print-multi-directory` option in addition to your other build options. For example:

```
> arm-none-eabi-gcc -print-multi-directory options...
```

The output of this command is a directory name for the multilib, which you can look up in the tables given previously.

3.3. Using Flash Memory

Sourcery G++ Lite supports development and debugging of applications loaded into flash memory on ARM EABI targets. There are three steps involved:

1. You must use an appropriate linker script that identifies the ROM memory region on your target board, and locates the program text within that region. Refer to Chapter 5, “CS3™: The Code-Sourcery Common Startup Code Sequence” for information about the boards supported by Sourcery G++.
2. Next, load your program into the flash memory on your target board. You must use third-party tools to program the flash memory.
3. Finally, when debugging a program in flash memory, GDB must be told about the ROM region so that it knows where it must use hardware breakpoints to control program execution. If you are using the Sourcery G++ Debug Sprite to debug your program, the Sprite does this automatically, using the memory map provided in the board configuration file. Otherwise, you must provide this information explicitly.

When using GDB from the command line, you can mark the flash memory as read-only by using the command:

```
(gdb) mem start end ro
```

where *start* and *end* define the address range of the read-only memory region.

Although GDB automatically attempts to use hardware breakpoints on code locations in the read-only memory region, on many targets the number of available hardware breakpoints is very small. Furthermore, GDB also uses hardware breakpoints internally to implement commands such as **step**, **next**, and **finish**. Thus the number of breakpoints you can explicitly set in ROM may be fewer than the number supported by the target system.

For example, ARM7TDMI cores support only one hardware breakpoint, which must also be used internally by the debugger if you set any software breakpoints in RAM. On ARM9 cores, there are two hardware breakpoints supported and one is consumed by the debugger if you set any software breakpoints.

3.4. ARMv7-M Interrupt Handlers

Because of a discrepancy between the ARMv7-M Architecture and the ARM EABI, it is not safe to use normal C functions directly as interrupt handlers. The EABI requires the stack be 8-byte aligned, whereas ARMv7-M only guarantees 4-byte alignment when calling an interrupt vector. This can cause subtle runtime failures, usually when 8-byte types are used.

Functions that are used directly as interrupt handlers should be annotated with `__attribute__((__interrupt__))`. This tells the compiler to add special stack alignment code to the function prologue.

3.5. ARM Interrupt Vectors

The ARM CS3 implementation treats all systems as address vector processors. On traditional ARM systems (everything except M profile devices) an exception causes the processor to jump to a fixed address. However the arrangement of these addresses is such that the CS3 code vector model is not feasible.

The default CS3 vector table emulates an address vector system by placing an indirect branch at the real exception vector. If you override the entire vector table (rather than individual vectors) the indirect stubs are not included, and your replacement table is placed at address zero.

3.6. Using VFP Floating Point

3.6.1. Enabling Hardware Floating Point

GCC provides three basic options for compiling floating-point code:

- Software floating point emulation, which is the default. In this case, the compiler implements floating-point arithmetic by means of library calls.
- VFP hardware floating-point support using the soft-float ABI. This is selected by the `-mfloat-abi=softfp` option. When you select this variant, the compiler generates VFP floating-point instructions, but the resulting code uses the same call and return conventions as code compiled with software floating point.
- VFP hardware floating-point support using the VFP ABI, which is the VFP variant of the Procedure Call Standard for the ARM® Architecture (AAPCS). This ABI uses VFP registers to pass function arguments and return values, resulting in faster floating-point code. To use this variant, compile with `-mfloat-abi=hard`.

You can freely mix code compiled with either of the first two variants in the same program, as they both use the same soft-float ABI. However, code compiled with the VFP ABI is not link-compatible with either of the other two options. If you use the VFP ABI, you must use this option to compile your entire program, and link with libraries that have also been compiled with the VFP ABI. For example, you may need to use the VFP ABI in order to link your program with other code compiled by the ARM RealView® compiler, which uses this ABI.

Sourcery G++ Lite for ARM EABI includes libraries built with software floating point, which are compatible with VFP code compiled using the soft-float ABI. While the compiler is capable of generating code using the VFP ABI, no compatible runtime libraries are provided in Sourcery G++ Lite. However, VFP hard-float libraries built with both ABIs are available to Sourcery G++ Professional Edition subscribers.

Note that, in addition to selecting hard/soft float and the ABI via the `-mfloat-abi` option, you can also compile for a particular FPU using the `-mfpu` option. For example, `-mfpu=neon` selects VFPv3 with NEON coprocessor extensions.

3.6.2. NEON SIMD Code

Sourcery G++ includes support for automatic generation of NEON SIMD vector code. Autovectorization is a compiler optimization in which loops involving normal integer or floating-point code are transformed to use NEON SIMD instructions to process several data elements at once.

To enable generation of NEON vector code, use the command-line options `-ftree-vectorize -mfpu=neon -mfloat-abi=softfp`. The `-mfpu=neon` option also enables generation of VFPv3 scalar floating-point code.

Sourcery G++ also includes support for manual generation of NEON SIMD code using C intrinsic functions. These intrinsics, the same as those supported by the ARM RealView® compiler, are defined in the `arm_neon.h` header and are documented in the 'ARM NEON Intrinsics' section of the GCC manual. The command-line options `-mfpu=neon -mfloat-abi=softfp` must be specified to use these intrinsics; `-ftree-vectorize` is not required.

3.6.3. Half-Precision Floating Point

Sourcery G++ for ARM EABI includes support for half-precision (16-bit) floating point, including the new `__fp16` data type in C and C++, support for generating conversion instructions when compiling for processors that support them, and library functions for use in other cases.

To use half-precision floating point, you must explicitly enable it via the `-mfp16-format` command-line option to the compiler. For more information about `__fp16` representations and usage from C and C++, refer to the GCC manual.

3.7. ABI Compatibility

The Application Binary Interface (ABI) for the ARM Architecture is a collection of standards, published by ARM Ltd. and other organizations. The ABI makes it possible to combine tools from different vendors, including Sourcery G++ and ARM RealView®.

Sourcery G++ implements the ABI as described in these documents, which are available from the ARM Information Center¹:

- BSABI - ARM IHI 0036B (10 October 2008)

¹ <http://infocenter.arm.com>

- BPABI - ARM IHI 0037B (10 October 2008)
- EHABI - ARM IHI 0038A (10 October 2008)
- CLIBABI - ARM IHI 0039A (10 October 2008)
- AADWARF - ARM IHI 0040A (10 October 2008)
- CPPABI - ARM IHI 0041B (10 October 2008)
- AAPCS - ARM IHI 0042C (10 October 2008)
- RTABI - ARM IHI 0043B (10 October 2008)
- AAELF - ARM IHI 0044C (10 October 2008)
- ABI Addenda - ARM IHI 0045B (10 October 2008)

Sourcery G++ currently produces DWARF version 2, rather than DWARF version 3 as specified in AADWARF.

3.8. Object File Portability

It is possible to create object files using Sourcery G++ for ARM EABI that are link-compatible with the GNU C library provided with Sourcery G++ for ARM GNU/Linux as well as with the CodeSourcery C Library or Newlib C Library provided with ARM EABI toolchains. These object files are additionally link-compatible with other ARM C Library ABI-compliant static linking environments and toolchains.

To use this feature, when compiling your files with the bare-metal ARM EABI toolchain define the preprocessor constant `_AEABI_PORTABILITY_LEVEL` to 1 before including any system header files. For example, pass the option `-D_AEABI_PORTABILITY_LEVEL=1` on your compilation command line. No special options are required when linking the resulting object files. When building applications for ARM EABI, files compiled with this definition may be linked freely with those compiled without it.

Files compiled in this manner may not use the functions `fgetpos` or `fsetpos`, or reference the type `fpos_t`. This is because Newlib assumes a representation for `fpos_t` that is not AEABI-compliant.

Note that object files are only portable from EABI to GNU/Linux, and not vice versa; object files compiled for ARM GNU/Linux targets cannot be linked into ARM EABI executables.

3.9. ARM Profiling Implementation

Profiling is enabled by means of the `-pg` compiler option. In this mode, the compiler inserts a call to `__gnu_mcount_nc` into every function prologue. However, no implementation of `__gnu_mcount_nc` is provided (to do so would be impossible without knowledge of the execution environment).

You must provide your own implementation of `__gnu_mcount_nc`. Here are the requirements:

- On exit, pop the top value from the stack, and place it in the `lr` register. The `sp` register should be adjusted accordingly. For example, this is how to write it as a stub function:

```
.globl __gnu_mcount_nc
.type __gnu_mcount_nc, %function
__gnu_mcount_nc:
    mov    ip, lr
    pop    { lr }
    bx    ip
```

- Preserve all other register state except for `r12` and the CPSR condition code bits. In particular all coprocessor state and registers `r0-r3` must be preserved.
- Record and count all occurrences of the function calls in the program. The caller can be determined from the `lr` value stored on the top of the stack (on entry to `__gnu_mcount_nc`), and the callee can be determined from the current value of the `lr` register (i.e. the caller of this function).
- Arrange for the data to be saved to a file named `gmon.out` when the program exits (via `atexit`). Refer to the **gprof** profiler manual for more information.

Chapter 4

Using Sourcery G++ from the Command Line

This chapter demonstrates the use of Sourcery G++ Lite from the command line.

4.1. Building an Application

This chapter explains how to build an application with Sourcery G++ Lite using the command line. As elsewhere in this manual, this section assumes that your target system is arm-none-eabi, as indicated by the **arm-none-eabi** command prefix.

Using an editor (such as **notepad** on Microsoft Windows or **vi** on UNIX-like systems), create a file named `main.c` containing the following simple factorial program:

```
#include <stdio.h>

int factorial(int n) {
    if (n == 0)
        return 1;
    return n * factorial (n - 1);
}

int main () {
    int i;
    int n;
    for (i = 0; i < 10; ++i) {
        n = factorial (i);
        printf ("factorial(%d) = %d\n", i, n);
    }
    return 0;
}
```

Compile and link this program using the command:

```
> arm-none-eabi-gcc -o factorial main.c -T script
```

Sourcery G++ requires that you specify a linker script with the `-T` option to build applications for bare-board targets. Linker errors like undefined reference to ``read'` are a symptom of failing to use an appropriate linker script. Default linker scripts are provided in `arm-none-eabi/lib`. Refer to Chapter 5, “CS3™: The CodeSourcery Common Startup Code Sequence” for information about the boards and linker scripts supported by Sourcery G++ Lite.

There should be no output from the compiler. (If you are building a C++ application, instead of a C application, replace **arm-none-eabi-gcc** with **arm-none-eabi-g++**.)

4.2. Running Applications on the Target System

Consult your target board documentation for instructions on loading programs onto the target, and running them. Alternatively, you can use the Sourcery G++ Debug Sprite from within GDB to download and run programs on the target via a supported hardware debugging device.

4.3. Running Applications in the Simulator

Sourcery G++ Lite includes a simulator that you can use on the host system to run programs compiled for the target system. Since you do not need target hardware, this is the easiest way to try out Sourcery G++.

To use the simulator run:

```
> arm-none-eabi-run factorial
```

You should see the expected output:

```
factorial(0) = 1
factorial(1) = 1
factorial(2) = 2
factorial(3) = 6
factorial(4) = 24
factorial(5) = 120
factorial(6) = 720
factorial(7) = 5040
factorial(8) = 40320
factorial(9) = 362880
```

You can also use the simulator to execute target programs when debugging with GDB. See Section 4.4, “Running Applications from GDB” for more information.

The simulator supports the ARMv4 (StrongARM), ARMv4T (ARM7TDMI, ARM920, ARM9TDMI), ARMv5, and ARMv5TE (ARM926, Xscale) instruction sets. The arm-none-eabi-run simulator also includes support for Thumb instructions.

4.4. Running Applications from GDB

You can run GDB, the GNU Debugger, on your host system to debug programs running remotely on a target board or system. You can also run and debug programs using the GDB simulator.

While this section explains the alternatives for using GDB to run and debug application programs, explaining the use of the GDB command-line interface is beyond the scope of this document. Please refer to the GDB manual for further instructions.

4.4.1. Connecting to the GDB Simulator

GDB includes a simulator that allows you to debug ARM EABI applications without target hardware. To start and connect to the simulator from within GDB, use this command:

```
(gdb) target sim
```

4.4.2. Connecting to the Sourcery G++ Debug Sprite

The Sourcery G++ Debug Sprite is a program that runs on the host system to support hardware debugging devices. You can use the Debug Sprite to run and debug programs on a target board without an operating system, or to debug an operating system kernel. See Chapter 6, “Sourcery G++ Debug Sprite” for detailed information about the supported devices.

You can start the Sprite directly from within GDB:

```
(gdb) target remote | arm-none-eabi-sprite arguments
```

Refer to Section 6.3, “Invoking Sourcery G++ Debug Sprite” for a full description of the Sprite arguments.

4.4.3. Connecting to an External GDB Server

From within GDB, you can connect to a running **gdbserver** or other debugging stub that uses the GDB remote protocol using:

```
(gdb) target remote host:port
```

where *host* is the host name or IP address of the machine the stub is running on, and *port* is the port number it is listening on for TCP connections.

4.4.4. Loading and Running Applications

Connecting to a bare-metal target or simulator from GDB does not cause your program to be loaded into target memory. You must do this explicitly from GDB after you connect:

```
(gdb) load
```

Alternatively, you can use third-party tools to load your application into flash memory before starting GDB.

To begin execution of your application, you should generally use the **continue** command:

```
(gdb) continue
```

However, you should use **run** instead of **continue** to start your program if you used **target sim** to connect:

```
(gdb) run
```

Chapter 5

CS3™: The CodeSourcery Common Startup Code Sequence

CS3 is CodeSourcery's low-level board support library. This chapter describes the organization of the system startup code and tells you how you can customize it, such as by defining your own interrupt handlers. This chapter also documents the boards supported by Sourcery G++ Lite and the compiler and linker options you need to use with them.

Many developers turn to the GNU toolchain for its cross-platform consistency: having a single system support so many different processors and boards helps to limit risk and keep learning curves gentle. Historically, however, the GNU toolchain has lacked a consistent set of conventions for processor- and board-level initialization, language run-time setup, and interrupt and trap handler definition.

The CodeSourcery Common Startup Code Sequence (CS3) addresses this problem. For each supported system, CS3 provides a set of linker scripts describing the system's memory map, and a board support library providing generic reset, startup, and interrupt handlers. These scripts and libraries all follow a standard set of conventions across a range of processors and boards.

In addition to providing linker support, CS3's functionality is fully integrated with the Sourcery G++ Debug Sprite. For each supported board, CS3 provides the board file containing the memory map and initialization sequence required for debugging applications on the board via the Sprite, as documented in Section 6.9, "Supported Board Files".

5.1. Startup Sequence

CS3 divides the startup sequence into three phases:

- In the *hard reset phase*, we initialize the memory controller and set up the memory map.
- In the *assembly initialization phase*, we prepare the stack to run C code, and jump to the C initialization function.
- In the *C initialization phase*, we initialize the data areas, run constructors for statically-allocated objects, and call `main`.

The hard reset and assembly initialization phases are necessarily written in assembly language; at reset, there may not yet be stack to hold compiler temporaries, or perhaps even any RAM accessible to hold the stack. These phases do the minimum necessary to prepare the environment for running simple C code. Then, the code for the final phase may be written in C; CS3 leaves as much as possible to be done at this point.

The CodeSourcery board support library provides default code for all three phases. The hard reset phase is implemented by board-specific code. The assembly initialization phase is implemented by code specific to the processor family. The C initialization phase is implemented by generic code.

5.1.1. The Hard Reset Phase

This phase is responsible for initializing board-specific registers, such as memory base registers and DRAM controllers, or scanning memory to check the available size. It is written in assembler and ends with a jump to `_start`, which is where the assembly initialization phase begins.

The hard reset code is in a section named `.cs3.reset`. The section must define a symbol named `__cs3_reset_sys`, where `sys` is a name for the board being initialized; for example, the reset code for Altera Cyclone III Cortex-M1 boards would be named `__cs3_reset_cycloneiii_cm1`. The linker script defines the symbol `__cs3_reset` to refer to this reset address. If you need to refer to the reset address from generic code, you can use this non-specific `__cs3_reset` name.

When using the Sourcery G++ Debug Sprite, the Sprite is responsible for carrying out the initialization done in this phase. In this case execution begins at the start of the assembly initialization phase, except for simulators as described below.

Some simulators provide a supervisory operation to determine the amount of available memory. This operation is performed in the hard reset phase. Thus for simulators, execution always begins at `__cs3_reset_sys`.

The CodeSourcery board support library provides reasonable default reset code, but you may provide your own reset code by defining `__cs3_reset_sys` in an object file or library, in a `.cs3.reset` section.

5.1.2. The Assembly Initialization Phase

This phase is responsible for initializing the stack pointer and creating an initial stack frame. The symbol `_start` marks the entry point of the assembly initialization code; this name lacks the `__cs3` prefix because it is the symbol traditionally used by debuggers and other integrated development environments for the address where program execution begins. The assembly initialization phase ends with a call or jump to `__cs3_start_c`.

Simulators typically initialize the stack pointer and initial stack frame automatically on startup. CS3 can also support targets running a boot monitor that likewise initializes the stack before starting user code. On these targets, CS3 does not perform the assembly initialization phase at all; instead, `_start` is aliased to `__cs3_reset_sys`, so that execution always starts with the hard reset phase. The hard reset phase then ends with a jump directly to `__cs3_start_c`.

On the other hand, on bare-board targets setting the stack pointer explicitly in the assembly initialization phase is required even if the processor itself initializes the stack pointer automatically on reset. This is to support restarting programs from `_start` in the debugger.

The value of the symbol `__cs3_stack` provides the initial value of the stack pointer. The CodeSourcery linker scripts provide a default value for this symbol, which you may override by defining `__cs3_stack` yourself.

The initial stack frame is created for the use of ordinary C and C++ calling conventions. The stack should be initialized so that backtraces stop cleanly at this point; this might entail zeroing a dynamic link pointer, or providing hand-written DWARF call frame information.

Finally, we call the C function `__cs3_start_c`. This function never returns, and `_start` need not be prepared to handle a return from it.

As with the hard reset code, the CodeSourcery board support library provides reasonable default assembly initialization code. However, you may provide your own code by providing a definition for `_start`, either in an object file or a library.

5.1.3. The C Initialization Phase

Finally, C code can be executed. The C startup function is declared as follows:

```
void __cs3_start_c (void) __attribute__((noreturn));
```

In this function we take the following steps:

- Initialize all `.data`-like sections by copying their contents.
- Clear all `.bss`-like sections.
- Run constructors for statically-allocated objects, recorded using whatever conventions are usual for C++ on the target architecture.

CS3 reserves priorities from 0 to 100 for use by initialization code. You can handle tasks like enabling interrupts, initializing coprocessors, pointing control registers at interrupt vectors, and so on by defining constructors with appropriate priorities.

- Call `main` as appropriate.
- Call `exit`, if it is available.

As with the hard reset and assembly initialization code, the CodeSourcery board support library provides a reasonable definition for the `__cs3_start_c` function. You may override this by providing a definition for `__cs3_start_c`, either in an object file or in a library.

The CodeSourcery-provided definition of `__cs3_start_c` can pass command-line arguments to `main` using the normal C `argc` and `argv` mechanism if the board support package provides corresponding definitions for `__cs3_argc` and `__cs3_argv`. For example:

```
int __cs3_argc;  
char **__cs3_argv;
```

These variables should be initialized using a constructor function, which is run by `__cs3_start_c` after it initializes the data segment. Use the `constructor` attribute on the function definition:

```
__attribute__((constructor))  
static void __cs3_init_args (void) {  
    __cs3_argc = ...;  
    __cs3_argv = ...;  
}
```

The constructor function may have an arbitrary name; `__cs3_init_args` is used only for illustrative purposes here.

If definitions of `__cs3_argc` and `__cs3_argv` are not provided, then the default `__cs3_start_c` function invokes `main` with zero as the `argc` argument and a null pointer as `argv`.

5.2. Exit and Embedded Systems

A program running on an embedded system is usually designed never to exit — it runs until the system is powered down. The C and C++ standards leave it unspecified as to whether `exit` is called at program termination. If the program never exits, then there is no reason to include `exit`, facilities to run functions registered with `atexit`, or global destructors. This code would never be run and would therefore just waste space in the application.

The CS3 startup code, by itself, does not cause `exit` to be present in the application. It dynamically checks whether `exit` is present, and only calls it if it is. If you require `exit` to be present, either refer to it within your application, or add `-Wl,-u,exit` to the linking command line.

Similarly, code to register global destructors is only invoked when `atexit` is already in the executable; CS3, by itself, does not cause `atexit` to be present. If you require `atexit`, either refer to it within your application, or add `-Wl,-u,atexit` to the linking command line.

5.3. Memory Layout

The header file `cs3.h` declares variables and types that describe the layout of memory on the system to C code. The variables are defined by the CS3 linker script or in the board support library.

The following variables describe the regions of memory to be initialized at startup:

```

/* The number of elements in __cs3_regions. */
const size_t __cs3_region_num;

/* An untyped object, aligned on an eight-byte boundary. */
typedef unsigned char __cs3_byte_align8
    __attribute__((aligned (8)));

struct __cs3_region
{
    /* Flags for this region. None defined yet. */
    unsigned flags;

    __cs3_byte_align8 *init; /* Region's initial contents. */
    __cs3_byte_align8 *data; /* Region's start address. */

    /* These sizes are always a multiple of eight. */
    size_t init_size; /* Size of initial data. */
    size_t zero_size; /* Additional size to be zeroed. */
};

/* An array of memory regions. __cs3_regions[0] describes
   the region holding .data and .bss. */
extern const struct __cs3_region __cs3_regions[];

```

The following variables describe the area of memory to be used for the dynamically-allocated heap:

```

/* The addresses of these objects are the start and end of
   free space for the heap, typically following .data and .bss.
   However, &__cs3_heap_end may be zero, meaning that we must
   determine the heap limit in some other way --- perhaps via a
   supervisory operation on a simulator, or simply by treating
   the stack pointer as the limit. */
extern __cs3_byte_align8 __cs3_heap_start[];
extern __cs3_byte_align8 __cs3_heap_end[];

/* The end of free space for the heap, or zero if we haven't been
   able to determine it. It usually points to __cs3_heap_end,
   but in some configurations, may be overridden by a supervisory
   call in the reset code. */
extern void *__cs3_heap_limit;

```

For each region named *R*, *cs3.h* declares the following variables:

```

/* The start of the region. */
extern unsigned char __cs3_region_start_R[]
    __attribute__((aligned (8)));

/* The region's size, in bytes. */
extern const size_t __cs3_region_size_R;

```

At the assembly level, the linker script also defines symbols with the same names and values.

If the region is initialized, then `cs3.h` also declares the following variables, corresponding to the region's element in `__cs3_regions`:

```
/* The data we initialize the region with. */
extern const unsigned char __cs3_region_init_R[]
    __attribute__((aligned (8)));

/* The size of the initialized portion of the region. */
extern const size_t __cs3_region_init_size_R;

/* The size of the additional portion to be zeroed. */
extern const size_t __cs3_region_zero_size_R;
```

Any of these identifiers may actually be defined as preprocessor macros that expand to expressions of the appropriate type and value.

Like the `struct __cs3_region` members, these regions are all aligned on eight-byte boundaries, and their sizes are multiples of eight bytes.

CS3 linker scripts place the contents of sections named `.cs3.region-head.R` at the start of each memory region. Note that CS3 itself may want to place items (like interrupt vector tables) at these locations; if there is a conflict, CS3 raises an error at link time.

5.4. Interrupt Vectors and Handlers

CS3 provides standard handlers for interrupts, exceptions and traps, but also allows you to easily define your own handlers as needed. In this section, we use the term *interrupt* as a generic term for this entire class of events.

Different processors handle interrupts in various ways, but there are two general approaches:

- Some processors fetch an address from an array indexed by the interrupt number, and jump to that address. We call these *address vector* processors; Cortex-M1 systems are a typical example.
- Others multiply the interrupt number by some constant factor, add a base address, and jump directly to that address. Here, the interrupt vector consists of blocks of code, so we call these *code vector* processors; PowerPC systems are a typical example.

On address vector processors, the CS3 library provides an array of pointers to interrupt handlers named `__cs3_interrupt_vector_form`, occupying a section named `.cs3.interrupt_vector`, where *form* identifies the particular processor variant the vector is appropriate for. If the processor supports more than one variety of interrupt vector (for example, a full-length form and a shortened form), then *form* identifies the variety as well. Each entry in the vector holds a reference to a symbol named `__cs3_isr_int`, where *int* is the customary name of that interrupt on the processor, or a number if there is no consistently used name. The library further provides a reasonable default definition for each `__cs3_isr_int` handler routine.

To override an individual handler, provide your own definition for the appropriate `__cs3_isr_int` symbol. The definition need not be placed in any particular object file section.

Interrupt handlers typically require special call/return and register usage conventions that are target-specific and beyond the scope of this document. As an alternative to writing interrupt handlers in assembly language, on some targets they may be written in C using the `interrupt` attribute. For example, to override the `__cs3_isr_nmi` handler, use the following definition:

```
void __attribute__((interrupt)) __cs3_isr_nmi (void)
{
    ... custom handler code ...
}
```

To override the entire interrupt vector, you can define `__cs3_interrupt_vector_form`, placing the definition in a section named `.cs3.interrupt_vector`. The linker script reports an error if the `.cs3.interrupt_vector` section is empty, to ensure that the definition of `__cs3_interrupt_vector_form` occupies the proper section.

You may define the vector in C with an array of pointers using the `section` attribute to place it in the appropriate section. For example, to override the interrupt vector on Altera Cyclone III Cortex-M1 boards, make the following definition:

```
typedef void handler(void);
handler *__attribute__((section(".cs3.interrupt_vector")))
__cs3_interrupt_vector_micro[] =
{ ... };
```

On code vector processors, we follow the same conventions, with the following exceptions:

- In addition to being named `__cs3_isr_int`, each interrupt handler must also occupy a section named `.cs3.interrupt_int`. Naturally, each handler must fit within a single interrupt vector entry.
- Instead of providing a default definition for `__cs3_interrupt_vector_form` in the library, the linker script gathers the `.cs3.interrupt_int` sections together, in the proper order and on the necessary address boundaries, and defines the `__cs3_interrupt_vector_form` symbol to refer to its start.

To override an individual handler on a code vector processor, you provide your own definition for `__cs3_isr_int`, placed in an appropriate section. The linker script ensures that each `.cs3.interrupt_int` section is non-empty, so that placing a handler in the wrong section elicits an error at link time.

CS3 does not allow you to override the entire interrupt vector on code vector processors, because the code vector must be constructed by the linker script, and thus cannot come from a library or object file. However, the portion of the linker script that constructs the interrupt vector occupies its own file, which other linker scripts can incorporate using the **INCLUDE** linker script command, making it easier to replace the linker script entirely and still take advantage of CS3's other features.

Some processors, like the Innovasic fido, use more than one interrupt vector: the processor provides several interrupt vector pointer registers, each used in different circumstances. Each register may point to a different vector, or some or all may share vectors.

On these processors, CS3 provides only a single pre-constructed interrupt vector, but defines a separate symbol for each interrupt vector pointer register; all the symbols point to the pre-constructed vector by default. The CS3 startup code initializes each register from the corresponding symbol. You can provide your own vectors by defining the appropriate symbols.

For example, the fido processor has five contexts, each of which can use its own interrupt vector; on this architecture, CS3 defines the standard `__cs3_interrupt_vector_fido` symbol referring to the pre-constructed vector, and then goes on to define per-context symbols `__cs3_interrupt_vector_fido_ctx0`, `__cs3_interrupt_vector_fido_ctx1`, and so on, all referring to `__cs3_interrupt_vector_fido`. The CS3 startup code sets each context's vector register to

the value of the corresponding symbol. By default, all the contexts share an interrupt vector, but if your code provides its own definition for `__cs3_interrupt_vector_fido_ctx1`, then the startup code initializes context one's register to point to that vector instead.

This arrangement requires you to use a different approach to specify a handler for a secondary context that differs from the corresponding handler in the primary context. For example, to handle division-by-zero exceptions in context 1 with the function `ctx1_divide_by_zero`, you should write the following:

```
typedef void (*handler_type) (void);
handler_type __cs3_interrupt_vector_fido_ctx1[256];
extern handler_type __cs3_interrupt_vector_fido[256];

__attribute__((interrupt))
void
ctx1_divide_by_zero (void)
{
    /* Your code here.  */
}

__attribute__((constructor))
void
initialize_vector_ctx1 (void)
{
    /* Initialize our custom vector from the
       pre-constructed CS3 vector.  */
    memcpy (__cs3_interrupt_vector_fido_ctx1,
           __cs3_interrupt_vector_fido,
           sizeof (__cs3_interrupt_vector_fido));

    /* Initialize custom interrupt handlers.  */
    __cs3_interrupt_vector_fido_ctx1[5] = ctx1_divide_by_zero;
}
```

With this code in place, when a division-by-zero exception occurs in context 1, the processor calls `ctx1_divide_by_zero` to handle it. Defining `initialize_vector_ctx1` with the constructor attribute arranges for CS3 to call it before calling your program's main function.

5.5. Linker Scripts

CS3 provides linker scripts for each supported board. Each board may be used in a number of different configurations, and these are reflected in the linker script names. The linker scripts are named `board-profile-hosted.ld`, where *board* is the name of the board, *profile* describes the memory arrangement used and *-hosted* indicates whether hosting or semihosting is provided.

Caution

Linker scripts are required to create executable programs for ARM EABI targets. When invoking the Sourcery G++ linker from the command line, you must explicitly supply a linker script using the `-T` option; otherwise a link error results.

5.5.1. Program and Data Placement

Many boards have both RAM and ROM (flash) memory devices. CS3 provides distinct linker scripts to place the application either entirely in RAM, or in ROM where data is initialized during the C initialization phase.

Some boards have very small amounts of RAM memory. If you use large library functions (such as `printf` and `malloc`), you may overflow the available memory. You may need to use the ROM-based linker scripts for such programs, so that the program itself is stored in ROM. You may be able to reduce the total amount of memory used by your program by replacing portions of the Sourcery G++ runtime library and/or startup code.

5.5.2. Hosting and Semihosting

CS3 is designed to support boards without an operating system. To allow functions like `open` and `write` to work without operating system support, a *semihosting* feature is supported, in conjunction with the debugger.

With semihosting enabled, these system calls are translated into equivalent function calls on your host system. You can only use these function calls while connected to the debugger; if you try to use them when disconnected from the debugger, you will get a hardware exception.

Semihosting requires support from the remote GDB debugging stub or agent, as well as the debugger itself. The Sourcery G++ Debug Sprite implements semihosting for all supported devices. Semihosting is also supported by the GDB Simulator included with Sourcery G++ Lite. However, semihosting may not be supported by debugging stubs provided by third parties. If you are using a debug device that communicates with GDB using the GDB remote protocol, check the documentation for your device to see whether semihosting is supported.

A good use of semihosting is to display debugging messages. For example, this program prints a message on the standard error stream on the host:

```
#include <unistd.h>

int main () {
    write (STDERR_FILENO, "Hello, world!\n", 14);
    return 0;
}
```

The hosted CS3 linker scripts provide the semihosting support, and as such programs linked with them may only be run with the debugger. The unhosted CS3 linker scripts provide stub versions of the system calls, which return an appropriate error value in `errno`. If such a stub system call is required in the executable, the linker also produces a warning. Such a warning may indicate that you have left debugging code active, and that your executable is larger than it might need to be.

Some targets supported by CS3 can run a boot monitor that provides console I/O services and other basic system calls. CS3 can also provide hosting via these facilities; where a boot monitor is supported, this is noted in the board tables below. Unlike semihosting, hosting via the boot monitor can be used when running programs outside of the debugger.

5.5.3. Choosing a Linker Script

When using Sourcery G++ from the command line, you must add `-T script` to your linking command, where *script* is the appropriate linker script. For example, to target Altera Cyclone III Cortex-M1 boards, you could link with `-T cycloneiii-cml-ram-hosted.ld`.

5.6. Supported Boards for ARM EABI

CS3 provides support for the following boards on ARM EABI targets.

Altera Cyclone III Cortex-M1		
Processor name:	Cortex-M1	
Processor options:	-mcpu=cortex-m1 -mthumb	
Memory regions:	ram, rom, itcm	
Linker scripts:	RAM Hosted	cycloneiii-cm1-ram-hosted.ld
	RAM Unhosted	cycloneiii-cm1-ram.ld
	ROM Hosted	cycloneiii-cm1-rom-hosted.ld
	ROM Unhosted	cycloneiii-cm1-rom.ld

ARM M-profile Simulator		
Processor name:	Cortex-M3	
Processor options:	-mcpu=cortex-m3 -mthumb	
Memory regions:	ram	
Linker scripts:	Hosted	generic-m-hosted.ld
	Unhosted	generic-m.ld

ARM Simulator (VFP)		
Processor name:	unspecified	
Processor options:	none	
Memory regions:	ram	
Linker scripts:	Hosted	generic-vfp-hosted.ld
	Unhosted	generic-vfp.ld

ARM Simulator		
Processor name:	unspecified	
Processor options:	none	
Memory regions:	ram	
Linker scripts:	Hosted	generic-hosted.ld
	Unhosted	generic.ld

ARMulator (RDI)		
Processor name:	unspecified	
Processor options:	none	
Memory regions:	ram	
Linker scripts:	RAM Hosted	armulator-ram-hosted.ld
	RAM Unhosted	armulator-ram.ld

5.7. Interrupt Vector Tables

The ARM interrupt vector table (`__cs3_interrupt_vector_arm`) contents are:

Number	Name	Meaning
0	<code>__cs3_reset</code>	Reset entry point
1	<code>__cs3_isr_undef</code>	Undefined Instruction
2	<code>__cs3_isr_swi</code>	Software Interrupt/Supervisor Call
3	<code>__cs3_isr_pabort</code>	Prefetch Abort
4	<code>__cs3_isr_dabort</code>	Data Abort
5	<code>__cs3_isr_reserved</code>	
6	<code>__cs3_isr_irq</code>	External Interrupt (IRQ)
7	<code>__cs3_isr_fiq</code>	Fast Interrupt (FIQ)

The Microcontroller Profile interrupt vector table (`__cs3_interrupt_vector_micro`) contents are:

Number	Name	Meaning
0	<code>__cs3_stack</code>	Initial stack pointer
1	<code>__cs3_reset</code>	Reset entry point
2	<code>__cs3_isr_nmi</code>	Non Maskable Interrupt
3	<code>__cs3_isr_hard_fault</code>	Hardware fault
4	<code>__cs3_isr_mpu_fault</code>	MPU fault
5	<code>__cs3_isr_bus_fault</code>	Bus fault
6	<code>__cs3_isr_usage_fault</code>	Usage fault
7..10	<code>__cs3_isr_reserved_7..10</code>	Reserved for future use
11	<code>__cs3_isr_svcall</code>	System Vector Call
12	<code>__cs3_isr_debug</code>	Debug interrupt
13	<code>__cs3_isr_reserved_13</code>	Reserved for future use
14	<code>__cs3_isr_pendsv</code>	
15	<code>__cs3_isr_systick</code>	System Ticker
16..47	<code>__cs3_isr_external_0..31</code>	External interrupt

5.8. Regions and Memory Sections

The following regions are defined for ARM EABI.

Region	Contents
<code>itcm</code>	ITCM
<code>ram</code>	<code>.data</code> and <code>.bss</code> sections. In ram-based profiles, also contains <code>.text</code> and other program-like sections.
<code>rom</code>	<code>.text</code> and other program-like sections. Initialization values for data-like regions.

Note that not all regions are provided in every linker script or profile; see the documentation of the individual linker scripts in Section 5.6, “Supported Boards for ARM EABI”, above.

Regions documented as "Memory regions" correspond to similarly-named program sections. For example, the linker script assigns the `.ram` section to the `ram` region. You can explicitly locate data or code in these sections using section attributes in your source C or C++ code. Section attributes are especially useful on code compiled for boards that include special memory banks, such as a fast on-chip cache memory, in addition to the default `ram` and/or `rom` regions. CS3 arranges for additional data-like sections to be initialized in the same way as the default `.data` section.

As an example to illustrate the attribute syntax, you can put a variable `v` in the `.ram` section using:

```
int v __attribute__ ((section (".ram")));
```

To declare a function `f` in this section, use:

```
int f (void) __attribute__ ((section (".ram"))) {...}
```

For more information about attribute syntax, see the GCC manual.

Regions documented as "Other regions" do not have a corresponding program section. Typically, these regions correspond to memory-mapped control and I/O registers that cannot be used for general data or program storage. If you need to manipulate data in these regions, you can use the CS3 memory layout facilities declared in `cs3.h`, as described in Section 5.3, “Memory Layout”.

Memory maps for boards supported by Sourcery G++ Lite for ARM EABI are documented in XML files in the `arm-none-eabi/lib/boards/` subdirectory of your Sourcery G++ installation directory.

Chapter 6

Sourcery G++ Debug Sprite

This chapter describes the use of the Sourcery G++ Debug Sprite for remote debugging. The Sprite allows you to debug programs running on a bare board without an operating system. This chapter includes information about the debugging devices and boards supported by the Sprite for ARM EABI.

Sourcery G++ Lite contains the Sourcery G++ Debug Sprite for ARM EABI. This Sprite is provided to allow debugging of programs running on a bare board. You can use the Sprite to debug a program when there is no operating system on the board, or for debugging the operating system itself. If the board is running an operating system, and you wish to debug a program running on that OS, you should use the facilities provided by the OS itself (for instance, using **gdbserver**).

The Sprite acts as an interface between GDB and external debug devices and libraries. Refer to Section 6.3, “Invoking Sourcery G++ Debug Sprite” for information about the specific devices supported by this version of Sourcery G++ Lite.

Important

The Sourcery G++ Debug Sprite is not part of the GNU Debugger and is not free or open-source software. You may use the Sourcery G++ Debug Sprite only with the GNU Debugger. You may not distribute the Sourcery G++ Debug Sprite to any third party.

6.1. Probing for Debug Devices

Before running the Sourcery G++ Debug Sprite for the first time, or when attaching new debug devices to your host system, it is helpful to verify that the Sourcery G++ Debug Sprite recognizes your debug hardware. From the command line, invoke the Sprite with the `-i` option:

```
> arm-none-eabi-sprite -i
```

This prints out a list of supported device types. For devices that can be autodetected, it additionally probes for and prints out a list of attached devices. For instance:

```
CodeSourcery ARM Debug Sprite
  (Sourcery G++ Lite Sourcery G++ Lite 2009q3-68)
armusb: [speed=<n:0-7>] ARMUSB device
  armusb:///0B01000C - Stellaris Evaluation Board (0B01000C)
rdi: (rdi-library=<file>&rdi-config=<file>) RDI Device
  rdi:/// - RDI Device
```

This shows that ARMUSB and RDI devices are supported. The exact set of supported devices depends on your host system and the version of Sourcery G++ you have installed; refer to Section 6.3, “Invoking Sourcery G++ Debug Sprite” for complete information.

Note that it may take several seconds for the Debug Sprite to probe for all types of supported devices.

6.2. Debug Sprite Example

Start by compiling and linking a simple test program for your target board, following the instructions in Chapter 4, “Using Sourcery G++ from the Command Line”. Use the `-g` option to tell the compiler to generate debugging information.

To build the `factorial` program to run on the ARMulator simulator, which can communicate with the Sprite via the RDI protocol, use:

```
> arm-none-eabi-gcc -g -Tarmulator-ram-hosted.ld main.c \
  -o factorial
```

Next start the debugger on your host system:

```
> arm-none-eabi-gdb factorial
```

The command for connecting GDB to the board depends on the debug device you are using; this is described in more detail in Section 6.3, “Invoking Sourcery G++ Debug Sprite”. If you are connecting via RDI, you must specify the full path to the RDI library file and configuration file for that library. Use quotes to escape the Sprite argument syntax from the shell. For example, use a command like this to connect to the ARMulator:

```
(gdb) target remote | arm-none-eabi-sprite \
"rdi:///rdi-library=library&rdi-config=config" armulator
```

The Sprite prints some status messages as it connects to your debug device and target board. If the connection is successful, you should see output similar to:

```
arm-none-eabi-sprite:Target reset
0x00008936 in ?? ()
(gdb)
```

Next, use GDB to load your program onto the target board.

```
(gdb) load
```

At this point you can use GDB to control the execution of your program as required. For example:

```
(gdb) break main
(gdb) continue
```

6.3. Invoking Sourcery G++ Debug Sprite

The Debug Sprite is invoked as follows:

```
> arm-none-eabi-sprite [options] device-url board-file
```

The *device-url* specifies the debug device to use to communicate with the board. It follows the standard format:

```
scheme:scheme-specific-part[?device-options]
```

Most device URL schemes also follow the regular format:

```
scheme: [//hostname:[port]]/path[?device-options]
```

The meanings of *hostname*, *port*, *path* and *device-options* parts depend on the *scheme* and are described below. The following schemes are supported in Sourcery G++ Lite for ARM EABI:

rdi Use an RDI debugging device. Refer to Section 6.5, “Remote Debug Interface Devices”.

flashpro Use a FlashPro debugging device. Refer to Section 6.6, “Actel FlashPro Devices”.

altera Use an Altera FPGA. Refer to Section 6.7, “Altera Devices”.

The optional *?device-options* portion is allowed in all schemes. These allow additional device-specific options of the form *name=value*. Multiple options are concatenated using *&*.

The *board-file* specifies an XML file that describes how to initialize the target board, as well as other properties of the board used by the debugger. If *board-file* refers to a file (via a relative or absolute pathname), it is read. Otherwise, *board-file* can be a board name, and the toolchain’s

board directory is searched for a matching file. See Section 6.9, “Supported Board Files” for the list of supported boards, or invoke the Sprite with the `-b` option to list the available board files. You can also write a custom board file; see Section 6.10, “Board File Syntax” for more information about the file format.

Both the `device-url` and `board-file` command-line arguments are required to correctly connect the Sprite to a target board.

6.4. Sourcery G++ Debug Sprite Options

The following command-line options are supported by the Sourcery G++ Debug Sprite:

- `-b` Print a list of `board-file` files in the board config directory.
- `-h` Print a list of options and their meanings. A list of `device-url` syntaxes is also shown.
- `-i` Print a list of the accessible devices. If a `device-url` is also specified, only devices for that device type are scanned. Each supported device type is listed along with the options that can be appended to the `device-url`. For each discovered device, the `device-url` is printed along with a description of that device.
- `-l [host]:port` Specify the host address and port number to listen for a GDB connection. If this option is not given, the Debug Sprite communicates with GDB using stdin and stdout. If you start the Sprite from within GDB using the `target remote | arm-none-eabi-sprite ...` command, you do not need this option.
- `-m` Listen for multiple sequential connections. Normally the Debug Sprite terminates after the first connection from GDB terminates. This option instead makes it listen for a subsequent connection. To terminate the Sprite, open a connection and send the string `END\n`.
- `-q` Do not print any messages.
- `-v` Print additional messages.

If any of `-b`, `-i` or `-h` are given, the Debug Sprite terminates after providing the information rather than waiting for a debugger connection.

6.5. Remote Debug Interface Devices

Remote Debug Interface (RDI) devices are supported. The RDI device URL accepts no hostname, port or path components, so the `device-url` is specified as follows:

```
rdi:[:///][?device-options]
```

The following `device-options` are required:

- `rdi-library=library` Specify the library (DLL or shared object) implementing the RDI target you wish to use.
- `rdi-config=configfile` Specify a file containing configuration information for `library`. The format of this file is specific to the RDI library you are using,

but tends to constitute a list of *key=value* pairs. Consult the documentation of your RDI library for details.

6.6. Actel FlashPro Devices

On Windows hosts, Sourcery G++ Lite supports FlashPro devices used with Actel Cortex-M1 development kits.

For FlashPro devices, the *device-url* has the following form:

```
flashpro:[//usb12345/][?jtagclock=rate]
```

The optional *usb12345* part indicates the ID of the FlashPro device to connect to, which is useful if you have more than one such device attached to your computer. If the ID is omitted, the Debug Sprite connects automatically to the first detected FlashPro device. You can enumerate the connected FlashPro devices by invoking the Sprite with the *-i* switch, as follows:

```
> arm-none-eabi-sprite -i flashpro:
```

The *jtagclock* option allows the communication speed with the target board to be altered. The *rate* is specified in Hz and may range between 93750 and 4000000. The default is 93750, the slowest speed supported by the FlashPro device. Depending on your target board, you may be able to increase this rate, but beware that communication errors may occur above a certain threshold. If you encounter communication errors with a higher-than-default speed selected, try reducing the speed.

6.6.1. Installing FlashPro Windows drivers

Windows drivers for the FlashPro device are included with the FlashPro software provided by Actel. Refer to Actel's documentation for details on installing this software. You must use the Actel FlashPro software to configure the FPGA on your Cortex-M1 board, but it does not need to be running when using the Debug Sprite.

Once you have set up your board using the FlashPro software, you can check that it is recognized by the Sourcery G++ Debug Sprite by running the following command:

```
> arm-none-eabi-sprite -i
flashpro: [jtagclock=<n:93750-4000000>] FlashPro device
flashpro://usb12345/ - FlashPro Device
...
```

If output similar to the above does not appear, your FlashPro device is not working correctly. Contact CodeSourcery for further guidance in that case.

6.7. Altera Devices

The Debug Sprite can be used to debug applications running on a Cortex-M1 core embedded in an Altera FPGA supporting the System-Level Debug (SLD) architecture. Currently, the Sprite supports the Cyclone III FPGA Starter board on Microsoft Windows hosts.

The Debug Sprite accepts two forms of the *device-url* for Altera devices. For the common case where you have only one Altera Cortex-M1 device configured, you can use simply:

```
altera://
```

The full form of the *device-url* is:

```
altera://usbX/hubY/nodeZ
```

where *X*, *Y*, and *Z* are non-negative integers. The SLD architecture forms a hierarchy; there may be multiple USB Blaster devices (numbered by *X*), multiple Altera FPGAs (numbered by *Y*) per USB Blaster, and multiple nodes (numbered by *Z*) per FPGA.

The Debug Sprite can autodetect connected Altera Cortex-M1 devices. Invoking the Sprite with the `-i` option, as described in Section 6.1, “Probing for Debug Devices”, displays the *device-url* for each detected device:

```
> arm-none-eabi-sprite -i
...
altera: Altera SLD Hub device
  altera://usb0/hub0/node1 - Altera Cortex-M Device
```

6.7.1. Setting Up the Altera Device

Follow these steps for initial installation and set up of the Altera device.

1. Install Quartus II Web Edition (or any equivalent), available from Altera.
2. Install drivers for USB Blaster, also available from Altera.
3. Install Sourcery G++ Lite for ARM EABI. See Chapter 2, “Installation and Configuration”.
4. Connect the board and the host computer with a USB cable.
5. Turn on the board.
6. Use Quartus II to download a `.sof` file including a Cortex-M1 core to the FPGA.
7. Use `arm-none-eabi-sprite -i` to verify that the Sprite can detect the installed Cortex-M1 core.

6.7.2. Hardware Breakpoints

The Cortex-M1 core only permits hardware breakpoints to be set in the first 512MB of its address space. Because both external SRAM and flash memory are located at higher addresses, you cannot set hardware breakpoints in these memory regions.

6.8. Debugging a Remote Board

You can run the Sourcery G++ Debug Sprite on a different machine from the one on which GDB is running. For example, if your board is connected to a machine in your lab, you can run the debugger on your laptop and connect to the remote board. The Sourcery G++ Debug Sprite must run on the machine that is connected to the target board. You must have Sourcery G++ installed on both machines.

To use this mode, you must start the Sprite with the `-l` option and specify the port on which you want it to listen. For example:

```
> arm-none-eabi-sprite -l :10000 device-url board-file
```

starts the Sprite listening on port 10000.

When running GDB from the command line, use the following command to connect GDB to the remote Sprite:

```
(gdb) target remote host:10000
```

where *host* is the name of the remote machine. After this, debugging is just as if you are debugging a target board connected to your host machine.

For more detailed instructions on using the Sourcery G++ Debug Sprite in this way, please refer to the Sourcery G++ Knowledge Base¹.

6.9. Supported Board Files

The Sourcery G++ Debug Sprite for ARM EABI includes support for the following target boards. Specify the appropriate *board-file* as an argument when invoking the sprite from the command line.

Board	Config
Altera Cyclone III Cortex-M1	cycloneiii-cm1
ARMulator (RDI)	armulator

6.10. Board File Syntax

The *board-file* can be a user-written XML file to describe a non-standard board. The Sourcery G++ Debug Sprite searches for board files in the `arm-none-eabi/lib/boards` directory in the installation. Refer to the files in that directory for examples.

The file's DTD is:

```
<!-- Board description files

Copyright (c) 2007-2009 CodeSourcery, Inc.

THIS FILE CONTAINS PROPRIETARY, CONFIDENTIAL, AND TRADE
SECRET INFORMATION OF CODESOURCERY AND/OR ITS LICENSORS.

You may not use or distribute this file without the express
written permission of CodeSourcery or its authorized
distributor. This file is licensed only for use with
Sourcery G++. No other use is permitted.
-->

<!ELEMENT board
(properties?, feature?, initialize?, memory-map?)>

<!ELEMENT properties
(description?, property*)>

<!ELEMENT initialize
(write-register | write-memory | delay
| wait-until-memory-equal | wait-until-memory-not-equal)* >
```

¹ <https://support.codesourcery.com/GNUToolchain/kbentry132>

```

<!ELEMENT write-register EMPTY>
<!ATTLIST write-register
    address CDATA #REQUIRED
    value CDATA #REQUIRED
    bits CDATA #IMPLIED>
<!ELEMENT write-memory EMPTY>
<!ATTLIST write-memory
    address CDATA #REQUIRED
    value CDATA #REQUIRED
    bits CDATA #IMPLIED>
<!ELEMENT delay EMPTY>
<!ATTLIST delay
    time CDATA #REQUIRED>
<!ELEMENT wait-until-memory-equal EMPTY>
<!ATTLIST wait-until-memory-equal
    address CDATA #REQUIRED
    value CDATA #REQUIRED
    timeout CDATA #IMPLIED
    bits CDATA #IMPLIED>
<!ELEMENT wait-until-memory-not-equal EMPTY>
<!ATTLIST wait-until-memory-not-equal
    address CDATA #REQUIRED
    value CDATA #REQUIRED
    timeout CDATA #IMPLIED
    bits CDATA #IMPLIED>

<!ELEMENT memory-map (memory-device)*>
<!ELEMENT memory-device (property*, description?, sectors*)>
<!ATTLIST memory-device
    address CDATA #REQUIRED
    size CDATA #REQUIRED
    type CDATA #REQUIRED
    device CDATA #IMPLIED>

<!ELEMENT description (#PCDATA)>
<!ELEMENT property (#PCDATA)>
<!ATTLIST property name CDATA #REQUIRED>
<!ELEMENT sectors EMPTY>
<!ATTLIST sectors
    size CDATA #REQUIRED
    count CDATA #REQUIRED>

<!ENTITY % gdbtarget SYSTEM "gdb-target.dtd">
%gdbtarget;

```

All values can be provided in decimal, hex (with a 0x prefix) or octal (with a 0 prefix). Addresses and memory sizes can use a K, KB, M, MB, G or GB suffix to denote a unit of memory. Times must use a ms or us suffix.

The following elements are available:

<board> This top-level element encapsulates the entire description of the board. It can contain <properties>, <feature>, <initialize> and <memory-map> elements.

<properties>	<p>The <properties> element specifies specific properties of the target system. This element can occur at most once. It can contain a <description> element.</p> <p>It can also contain <property> elements with the following names:</p>
banked-regs	<p>The <code>banked-regs</code> property specifies that the CPU of the target board has banked registers for different processor modes (supervisor, IRQ, etc.).</p>
has-vfp	<p>The <code>has-vfp</code> property specifies that the CPU of the target board has VFP registers.</p>
system-v6-m	<p>The <code>system-v6-m</code> property specifies that the CPU of the target board has ARMv6-M architecture system registers.</p>
system-v7-m	<p>The <code>system-v7-m</code> property specifies that the CPU of the target board has ARMv7-M architecture system registers.</p>
core-family	<p>The <code>core-family</code> property specifies the ARM family of the target. The body of the <property> element may be one of <code>arm7</code>, <code>arm9</code>, <code>arm11</code>, and <code>cortex</code>.</p>
<initialize>	<p>The <initialize> element defines an initialization sequence for the board, which the Sprite performs before downloading a program. It can contain <write-register>, <write-memory> and <delay> elements.</p>
<feature>	<p>This element is used to inform GDB about additional registers and peripherals available on the board. It is passed directly to GDB; see the GDB manual for further details.</p>
<memory-map>	<p>This element describes the memory map of the target board. It is used by GDB to determine where software breakpoints may be used and when flash programming sequences must be used. This element can occur at most once. It can contain <memory-device> elements.</p>
<memory-device>	<p>This element specifies a region of memory. It has four attributes: <code>address</code>, <code>size</code>, <code>type</code> and <code>device</code>. The <code>address</code> and <code>size</code> attributes specify the location of the memory device. The <code>type</code> attribute specifies that device as <code>ram</code>, <code>rom</code> or <code>flash</code>. The <code>device</code> attribute is required for flash regions; it specifies the flash device type. The <memory-device> element can contain a <description> element.</p>
<write-register>	<p>This element writes a value to a control register. It has three attributes: <code>address</code>, <code>value</code> and <code>bits</code>. The <code>bits</code> attribute, specifying the bit width of the write operation, is optional; it defaults to 32.</p>
<write-memory>	<p>This element writes a value to a memory location. It has three attributes: <code>address</code>, <code>value</code> and <code>bits</code>. The <code>bits</code> attribute is optional and defaults to 32. Bit widths of 8, 16 and 32 bits are supported. The address written to must be naturally aligned for the size of the write being done.</p>

<code><delay></code>	This element introduces a delay. It has one attribute, <code>time</code> , which specifies the number of milliseconds, or microseconds to delay by.
<code><description></code>	This element encapsulates a human-readable description of its enclosing element.
<code><property></code>	The <code><property></code> element allows additional name/value pairs to be specified. The property name is specified in a <code>name</code> attribute. The property value is the body of the <code><property></code> element.

Chapter 7

Next Steps with Sourcery G++

This chapter describes where you can find additional documentation and information about using Sourcery G++ Lite and its components.

7.1. Sourcery G++ Subscriptions

CodeSourcery offers two levels of Sourcery G++ subscriptions. Professional Edition subscriptions include unlimited support, with no per-incident fees. CodeSourcery's support is provided by the same engineers who build Sourcery G++, and covers questions about installing and using Sourcery G++, the C and C++ programming languages, and all other topics relating to Sourcery G++. CodeSourcery provides updated versions of Sourcery G++ on demand to resolve critical problems reported by Professional Edition subscribers. Personal Edition subscriptions do not include support, but do include access to updates as long as the subscription remains active.

Subscription editions of Sourcery G++ also include many additional features not included in the free Lite editions:

- **Sourcery G++ IDE.** The Sourcery G++ IDE, based on Eclipse, provides a fully visual environment for developing applications, including an automated project builder, syntax-highlighting editor, and a graphical debugging interface. The debugger provides features especially useful to embedded systems programmers, including the ability to step through code at both the source and assembly level, view registers, and examine stack traces. CodeSourcery's enhancements to Eclipse include improved support for hardware debugging via JTAG or ICE units and complete integration with the rest of Sourcery G++.
- **Debug Sprites.** Sourcery G++ Debug Sprites provide hardware debugging support using JTAG and ICE devices. On some systems, Sourcery G++ Sprites can automatically program flash memory and display control registers. Debug Sprites included in Lite editions of Sourcery G++ include only a subset of the functionality of the Sprites in the subscription editions.
- **CS3.** CS3 provides a uniform, cross-platform approach to board initialization and interrupt handling on bare-metal ELF and EABI platforms. Subscription versions of Sourcery G++ include CS3 support for an expanded set of boards. In addition, the Sourcery G++ Board Builder allows you to extend the power of CS3 to cover custom board definitions. The Board Builder is fully integrated with the Sourcery G++ IDE and Debug Sprites.
- **CodeSourcery C Library.** Subscription versions of Sourcery G++ for bare-metal targets include the CodeSourcery C Library, a proprietary library implementation that is optimized to be smaller and faster than the Newlib C library included with Lite editions of Sourcery G++.
- **QEMU Instruction Set Simulator.** The QEMU instruction set simulator can be used to run — and debug — programs even without target hardware. Most bare-metal configurations of Sourcery G++ include QEMU and linker scripts targeting the simulator. Configurations of Sourcery G++ for GNU/Linux targets include a user-space QEMU emulator that runs on Linux hosts.
- **Sysroot Utilities.** Subscription editions of Sourcery G++ include a set of sysroot utilities for GNU/Linux targets. These utilities simplify use of the Sourcery G++ dynamic linker and shared libraries on the target and also support remote debugging with **gdbserver**.
- **GNU/Linux Prelinker.** For select GNU/Linux target systems, Sourcery G++ includes the GNU/Linux prelinker. The prelinker is a postprocessor for GNU/Linux applications which can dramatically reduce application launch time. CodeSourcery has modified the prelinker to operate on non-GNU/Linux host systems, including Microsoft Windows.
- **Library Reduction Utility.** Sourcery G++ also includes a Library Reduction Utility for GNU/Linux targets. This utility allows the GNU C Library to be relinked to include only those functions used by a given collection of binaries.

- **Additional Libraries.** For some platforms, additional run-time libraries optimized for particular CPUs are available. Pre-built binary versions of the libraries with debug information are also available to subscribers.
- **Additional Documentation.** Subscription customers receive expanded access to the Sourcery G++ Knowledge Base, covering many more tips, howtos, and application notes to help you make the best use of Sourcery G++.

If you would like more information about Sourcery G++ subscriptions, including a price quote or information about evaluating Sourcery G++, please send email to <sales@codesourcery.com>.

If you have a Sourcery G++ subscription, you may access your account by visiting the Sourcery G++ Portal¹. If you have a support account, but are unable to log in, send email to <support@codesourcery.com>.

7.2. Sourcery G++ Knowledge Base

The Sourcery G++ Knowledge Base is available to registered users at the Sourcery G++ Portal². Here you can find solutions to common problems including installing Sourcery G++, making it work with specific targets, and interoperability with third-party libraries. There are also additional example programs and tips for making the most effective use of the toolchain and for solving problems commonly encountered during debugging. The Knowledge Base is updated frequently with additional entries based on inquiries and feedback from customers.

7.3. Manuals for GNU Toolchain Components

Sourcery G++ Lite includes the full user manuals for each of the GNU toolchain components, such as the compiler, linker, assembler, and debugger. Most of the manuals include tutorial material for new users as well as serving as a complete reference for command-line options, supported extensions, and the like.

When you install Sourcery G++ Lite, links to both the PDF and HTML versions of the manuals are created in the shortcuts folder you select. If you elected not to create shortcuts when installing Sourcery G++ Lite, the documentation can be found in the `share/doc/sourceryg++-arm-none-eabi/` subdirectory of your installation directory.

In addition to the detailed reference manuals, Sourcery G++ Lite includes a Unix-style manual page for each toolchain component. You can view these by invoking the `man` command with the pathname of the file you want to view. For example, you can first go to the directory containing the man pages:

```
> cd $INSTALL/share/doc/sourceryg++-arm-none-eabi/man/man1
```

Then you can invoke `man` as:

```
> man ./arm-none-eabi-gcc.1
```

Alternatively, if you use `man` regularly, you'll probably find it more convenient to add the directory containing the Sourcery G++ man pages to your `MANPATH` environment variable. This should go in your `.profile` or equivalent shell startup file; see Section 2.6, "Setting up the Environment" for instructions. Then you can invoke `man` with just the command name rather than a pathname.

¹ <https://support.codesourcery.com/GNUToolchain/>

² <https://support.codesourcery.com/GNUToolchain/>

Finally, note that every command-line utility program included with Sourcery G++ Lite can be invoked with a `--help` option. This prints a brief description of the arguments and options to the program and exits without doing further processing.

Appendix A

Sourcery G++ Lite Release Notes

This appendix contains information about changes in this release of Sourcery G++ Lite for ARM EABI. You should read through these notes to learn about new features and bug fixes.

A.1. Changes in Sourcery G++ Lite for ARM EABI

This section documents Sourcery G++ Lite changes for each released revision.

A.1.1. Changes in Sourcery G++ Lite 2009q3-68

Out-of-range branch error. A compiler bug has been fixed that caused out-of-range branch errors from the assembler. The bug only affected code compiled in Thumb-2 mode.

A.1.2. Changes in Sourcery G++ Lite 2009q3-64

GDB crash fix. A GDB bug has been fixed that caused GDB to crash when unloading shared libraries or switching executables.

@FILE fix. A bug has been fixed in the processing of @FILE command-line options by GCC, GDB, and other tools. The bug caused any options in FILE following a blank line to be ignored.

Preprocessor error handling. The preprocessor now treats failing to find a file referenced via #include as a fatal error.

NEON improvements. The compiler now generates improved NEON vector code when copying memory or storing constants to memory using the NEON coprocessor. The compiler also generates better code for accessing data arrays that are not known to have 64-bit alignment. In addition, a bug that caused internal compiler errors when compiling for Thumb-2 with NEON enabled has been fixed, as has another bug that caused some vector shift NEON operations to be wrongly rejected.

ELF file corruption with strip. A bug that caused strip to corrupt unusual ELF files has been fixed.

GDB support for Cygwin pathnames. A bug in GDB's translation of Cygwin pathnames has been fixed.

Compiler errors with float32_t. A bug has been fixed that caused compiler errors when using the float32_t type from arm_neon.h.

Support for ARM Cortex-A5 cores. Sourcery G++ now includes basic support for ARM Cortex-A5 cores. Use the -mcpu=cortex-a5 command-line option.

Static variables and asm statements bug fix. A bug in GCC that caused functions containing static variables and asm statements to be miscompiled at -O2 or above has been fixed. The bug also occurred at other optimization levels when the -fremove-local-statics command-line option was used.

Linker script fixes. A bug in CS3 linker scripts for simulator profiles has been fixed. The bug resulted in data memory being too small, which sometimes caused the stack to be overwritten during initialization, or reduced space for malloc to allocate.

Warnings for naked functions. A compiler bug that resulted in incorrect warnings about missing return statements in non-void functions declared with the naked attribute has been fixed.

Optimizer bug fix. A bug in GCC that caused functions with complex loop nests to be miscompiled at -O2 or above has been fixed. The bug also occurred at other optimization levels when the -fpromote-loop-indices command-line option was used.

VFPv4 support. Sourcery G++ now includes support for VFPv4, VFPv4-D16 and NEON-VFPv4 coprocessors. Use the `-mfpu=vfpv4`, `-mfpu=vfpv4-d16` or `-mfpu=neon-vfpv4` options, respectively.

GCC internal compiler error. A bug has been fixed that caused the compiler to crash when optimizing code that casts between structure types and the type of the first field.

Flash programming support on Atmel AT91SAM7Sxxx. The Sourcery G++ Debug Sprite now supports flash programming on Atmel AT91SAM7Sxxx when using SEGGER J-Link devices.

ELF Program Headers. The linker now better diagnoses errors in the usage of `FILEHDR` and `PHDRS` keywords in `PHDRS` command of linker scripts. Refer to the linker manual for more information.

A.1.3. Changes in Sourcery G++ Lite 2009q3-37

Improved optimization for ARM. GCC now automatically enables loop unrolling and `-fpromote-loop-indices` when `-O2` or `-O3` is specified. Loop unrolling is limited at `-O2` to control code growth. These changes improve performance by more than 5%.

VFP assembly mnemonics. The assembler now accepts unified assembly mnemonics for VFP instructions (e.g. `VADD.f32 s0, s0`) in legacy syntax mode.

ARM Cortex-R4F assembler bug fix. The assembler now correctly recognizes the `-mcpu=cortex-r4f` command-line option to select the Cortex-R4F processor.

VFP half-precision extensions. Sourcery G++ now includes support for VFP coprocessors with half-precision floating-point extensions. This can be enabled with the `-mfpu=vfpv3-d16-fp16` or `-mfpu=vfpv3-fp16` command-line options.

Optimizer improvements. When optimizing for speed, the compiler now uses improved heuristics to limit certain types of optimizations that may adversely affect both code size and speed. This change also makes it possible to produce better code when optimizing for space rather than speed.

Improved optimization for Thumb-2. GCC now supports instruction scheduling for Thumb-2 code. This optimization is enabled when compiling with `-O2`, `-O3`, or `-Os`, and can improve performance substantially.

ARM VFP assembler bug fix. The assembler now correctly assembles the `vmls`, `vnmla` and `vnmls` mnemonics. Previously these were incorrectly assembled to different instructions.

GDB finish internal error. A bug has been fixed that caused a GDB internal error when using the `finish` command. The bug occurred when debugging optimized code.

Linking objects built without `-fPIC` into shared libraries. The linker now gives an error for attempts to link object files built without `-fPIC` or `-fpic` into shared libraries when those objects use the ARMv7 `MOVW` and `MOVT` instructions in ways that are unsafe in a shared library. Previously it built a shared library that behaved incorrectly when used.

GDB update. The included version of GDB has been updated to 6.8.50.20090630. This update adds numerous bug fixes and new features, including support for multi-byte and wide character sets and improved C++ template support.

New assembler directive `.inst`. The assembler now accepts the new `.inst` directive to generate an instruction from its integer encoding.

GDB and third-party compilers. Some bugs that caused GDB to crash when debugging programs compiled with third-party tools have been fixed. These bugs did not affect programs built with Sourcery G++.

Remote debugging hardware watchpoint bug fix. A GDB bug has been fixed that caused hardware watchpoint hits to be incorrectly reported in some cases.

Internal error in assembler. An assembler bug that caused an internal error when `.thumb` or `.arm` appears after an invalid instruction has been fixed.

GDB internal warning fix. A GDB bug has been fixed that caused warnings of the form `warning: (Internal error: pc address in read in psymtab, but not in symtab.)`.

Incorrect linker diagnostic removed. The linker has been corrected to not emit an error message when the load address of an output section with no contents overlaps an output section with contents. This can occur in linker scripts that use `MEMORY` regions and `AT>` to place initialized contents into ROM.

Improved bit counting operation. The `__builtin_ctz` built-in function, which returns the number of trailing zero bits in a value, has been improved to use a shorter instruction sequence for ARMv6T2 and later.

Out-of-range branch errors. A Thumb-2 code generation defect in the compiler that caused `branch out of range` errors from the assembler has been eliminated.

Binutils update. The binutils package has been updated to version 2.19.51.20090709 from the FSF trunk. This update includes numerous bug fixes.

Linker fix. The linker now correctly processes references to undefined local symbols. Such references are treated the same as references to undefined global symbols. Usually object files contain no such references, as they can never be satisfied.

Assembler validation improvements. The assembler now issues a warning when a section finishes with an unclosed IT instruction block at the end of the input file. It also now rejects unwinding directives that appear outside of a `.fnstart/.fnend` pair. Additionally, 32-bit Thumb instructions are now correctly rejected when assembling for cores that do not support these instructions.

Destructor function bug fix. A bug in CS3 has been fixed that caused functions with the `destructor` attribute not to be run on program termination.

Assembler validations fix. A bug in the assembler that caused some `addw` and `subw` instructions with `SP` or `PC` as operand to be wrongly rejected has been fixed.

-mauto-it assembler option replaced with -mimplicit-it. The `-mauto-it` command-line option to the assembler has been replaced with a more general `-mimplicit-it` option to control the behavior of the assembler when conditional instructions appear outside an IT instruction block. If you were previously using `-mauto-it`, you should now use `-mimplicit-it=always`. Other `-mimplicit-it` modes allow you to separately control implicit IT instruction insertion behavior in ARM and Thumb-2 code. For more information, refer to the assembler manual. In addition to renaming the option, a number of bugs in the implicit IT generation have been fixed.

GDB backwards compatibility fix. A bug has been fixed that caused GDB to crash when loading symbols from binaries built by very old versions of GCC.

Linker failure with Cortex-A8 erratum fix. A bug in the `--fix-cortex-a8` linker option has been fixed. The bug caused the linker either to produce a `bad value` error, or to silently generate an incorrect executable.

Debug information for variadic functions. A compiler bug that resulted in incorrect debug information for functions with variable arguments has been fixed.

Overlay sections. `arm-none-eabi-readelf` now correctly recognizes section headers for `ARM_DEBUGOVERLAY` and `ARM_OVERLAYSECTION` sections.

Code generation improvements. The compiler has been changed to make better use of VFP registers in mixed integer and floating-point code, resulting in faster code.

Register variable corruption. A compiler bug has been fixed that caused incorrect code to be generated when the frame pointer or other special-use registers are used as explicit local register variables, introduced via the `asm` keyword on their declarations.

Startup code debugging fixes. Two GDB bugs have been fixed that caused errors when debugging startup code. One bug caused an internal error message; the other caused the error `Cannot find bounds of current function`.

Assembler fix for mixed Thumb and ARM mode. A bug in the assembler has been fixed where mapping symbols were sometimes incorrectly placed at section boundaries. This could lead to incorrect disassembly in some cases.

C++ exception matching. A C++ conformance defect has been fixed. According to clause 15.3 of the standard, given a derived class `D` with base `B`, a thrown `D *` object is not caught by a handler with type `B *&` (that is, a reference to pointer `B`). The compiler formerly treated this case incorrectly as if the handler had type `B *`, which does catch `D *`.

-fremove-local-statics optimization. The `-fremove-local-statics` optimization is now enabled by default at `-O2` and higher optimization levels.

Elimination of spurious warnings about NULL. The C++ compiler no longer issues spurious warnings about comparisons between pointers to members and `NULL`.

Vectorizer improvements. The compiler now generates improved code for accesses to static nested array variables (e.g. `static int foo[8][8];`).

Linker bug fix. A bug that caused the linker to crash when `.ARM.exidx` sections were discarded by a linker script has been fixed.

Configuration file required for Debug Sprite. When invoking the Sourcery G++ Debug Sprite from the command line, it is now required to specify a board configuration file argument. This change eliminates a source of confusion and errors resulting from accidental omission of the configuration file argument, since recent improvements to debugger functionality depend on properties specified in the configuration file. Refer to Chapter 6, “Sourcery G++ Debug Sprite” for more details on invoking the Sourcery G++ Debug Sprite from the command line.

GCC version 4.4.1. Sourcery G++ Lite for ARM EABI is now based on GCC version 4.4.1. For more information about changes from GCC version 4.3 that was included in previous releases, see <http://gcc.gnu.org/gcc-4.4/changes.html>.

Watchpoint support. The Sourcery G++ Debug Sprite now implements watchpoints on all currently-supported debugging devices.

Linker map address sorting. The map generated by the linker `-Map` option now lists symbols sorted by address.

Assembler fix. The assembler now correctly diagnoses a missing operand to `b1` and `b1x` instructions. Previously, incorrect code was silently generated.

A.1.4. Changes in Sourcery G++ Lite 2009q1-161

Incorrect placement of linker-generated functions. A bug that caused some linker-generated functions (including stubs to support interworking from ARM mode to Thumb mode and stubs to avoid processor errata) to be placed in data sections has been fixed.

ARMulator support. CS3 now includes support for using the ARMulator via the Sourcery G++ Debug Sprite using the RDI protocol.

New option for automatically generating IT blocks. The assembler now allows use of conditional Thumb-2 instructions without requiring explicit IT instructions. Use the `-mauto-it` command-line option to enable this automatic generation of IT instructions.

Optimized memcpy. The Newlib implementation of `memcpy` has been optimized to increase performance on ARM targets that support prefetch instructions.

Support for Cortex-M0. Sourcery G++ Lite now includes support for Cortex-M0 processors. To compile for these processors, use `-mcpu=cortex-m0 -mthumb`.

Incorrect code when using `-falign-labels`. A bug that caused the compiler to generate incorrect code for `switch` statements when the `-falign-labels` option is used has been fixed.

Reduced compilation time. Compilation and build times when using Sourcery G++ Lite are now slightly faster. This performance improvement is the result of building the compilers and other host tools with a recent version of Sourcery G++, rather than an older GCC version.

Linker script load address processing. A bug in the linker has been fixed affecting linker scripts using `AT>region` to set the load address. This now follows the documented behavior of maintaining the virtual address to load address difference in output section statements. Refer to the "Output Section LMA" section of the linker manual for details of how to control the load address.

Debug section placement. A linker script bug in CS3 has been fixed that caused `.debug_ranges` debug sections to be misplaced.

Assembler bug fix. A bug in the assembler that caused duplicate and missing mapping symbols has been fixed. The bug caused incorrect `objdump` output and incorrect byte-swapping for BE8 configurations.

Multiple flash regions. A bug in the CS3 linker scripts affecting boards with multiple flash devices has been fixed. The bug caused initialization code to treat certain flash devices as normal memory.

Stack backtracing and C++ exception handling. Improvements have been made to the linker in support of C++ runtime exception handling and stack backtracing. A problem that caused crashes during the backtrace of C routines that were not compiled with the `-fexceptions` option has been fixed. In addition, the linker generates more compact stack unwinding tables which can lead to smaller executables.

Assembler floating-point format. The assembler now defaults to VFP format for floating-point numbers. It previously defaulted to the legacy FPA format if no `-mcpu` or `-march` option was

specified, or if a CPU with no floating-point unit was specified. This bug resulted in incorrect behavior of the `.double` and `.dcb.d` directives.

Incorrect linker-generated functions. A bug that caused some linker-generated functions (such as stubs to support interworking from ARM mode to Thumb mode) to contain only nop instructions instead of correct code sequences has been fixed.

Assembler diagnostics for invalid instructions. The assembler now issues diagnostics for invalid ADR and ADRL instructions. Formerly, these invalid instructions were silently mis-assembled. This assembler bug did not affect correct code.

Sprite's failure to reset the target. A bug has been fixed that sometimes caused the Sourcery G++ Debug Sprite to fail to reset the target when using the multiple sequential connection feature (enabled via the `-m` command-line option). This problem was specific to running the Debug Sprite on Microsoft Windows hosts.

Optimized memcpy and memset routines. The Newlib implementations of `memcpy` and `memset` have been optimized to increase performance on ARM targets.

Loop optimization improvements. A new option, `-fpromote-loop-indices`, has been added to the compiler. Specifying this option enables an optimization that improves the performance of loops with index variables of integer types narrower than the target machine word size, such as `char` or `short`. This optimization also applies to `int` on 64-bit targets.

Disassembler bug fix. A bug has been fixed that caused incorrect disassembly of some object files with multiple sections whose symbol tables included symbols in the middle of functions. These typically resulted from hand-written assembly.

DMB, DSB, and ISB instructions on ARMv6-M. The assembler now accepts the DMB, DSB, and ISB instructions on ARMv6-M CPUs, including Cortex-M0 and Cortex-M1. These instructions were incorrectly rejected on these CPUs in previous releases.

Extraneous linker error messages. A linker bug that caused extraneous error messages of the form `Dwarf Error: Offset (507) greater than or equal to .debug_str size (421).` has been corrected. This bug did not affect the correctness of output binaries.

Linker crash with very large applications. A linker bug that caused a crash when linking very large applications with the `--fix-cortex-a8` command-line option has been fixed.

Assembler marking of data. Data generated using the assembler directives `.ascii`, `.asciz`, `.dc.d`, `.dc.s`, `.dc.x`, `.dcb`, `.dcb.b`, `.dcb.d`, `.dcb.l`, `.dcb.s`, `.dcb.w`, `.dcb.x`, `.ds`, `.ds.b`, `.ds.d`, `.ds.l`, `.ds.p`, `.ds.s`, `.ds.w`, `.ds.x`, `.double`, `.fill`, `.float`, `.incbin`, `.single`, `.space`, `.skip`, `.string`, `.string8`, `.string16`, `.string32`, `.string64`, and `.zero` is now correctly marked by the assembler as data rather than code. This fixes incorrect byte-swapping of such data when linking for BE8 configurations.

arm-none-eabi-objcopy bug fix. A bug has been fixed that caused `arm-none-eabi-objcopy` to issue an error when generating output in the Intel HEX format and using `--change-section-lma` to change section addresses.

Linker script search path. The bug in the linker has been fixed that caused it not to follow its documented behavior for searching for linker scripts named with the `-T` option. Now scripts are looked up first in the current directory, then in library directories specified with `-L` command-line options, and finally in the default system linker script directory.

VFP ABI support. Sourcery G++ now supports the VFP variant of the Procedure Call Standard for the ARM® Architecture (AAPCS) in addition to the default soft-float ABI. The VFP ABI uses VFP registers to pass function arguments and return values, resulting in faster floating-point code. Code compiled with the VFP ABI is not compatible with the soft-float ABI libraries provided with Sourcery G++ Lite; however, VFP ABI libraries for little-endian ARM v7-A processors are now available as add-ons for Sourcery G++ Professional Edition. For further information about floating-point compiler, ABI and library support in Sourcery G++, refer to Section 3.6.1, “Enabling Hardware Floating Point”.

Cortex-A8 erratum workaround enabled for ARMv7-A. The workaround for the erratum in Cortex-A8 processors mentioned below is now enabled by default if you are targeting the ARMv7-A architecture profile. The workaround can be disabled by passing the `--no-fix-cortex-a8` option to the linker.

Improved vectorization. Automatic vectorization for NEON now uses the fused multiply-add (VMLA) and fused multiply-subtract (VMLS) instructions. These fused instructions are faster than the equivalent two-instruction sequence consisting of a multiply followed by an add or subtract.

Internal compiler error when optimizing. A bug has been fixed that caused internal compiler error: `in build2_stat` when compiling.

GDB quit error. A bug in GDB has been fixed that caused `quit` to report `Quitting: You can't do that without a process to debug. when debugging a core dump file.`

Out-of-bounds accesses to stack arrays. A bug has been fixed that caused internal compiler errors when some code involving out-of-bounds accesses to stack-allocated arrays was compiled with the `-mthumb` option. Such code is not valid C; although it is now accepted by the compiler and no diagnostic is issued, it has undefined behavior if executed.

Erratum workaround for Cortex-A8 processors. The linker now implements a workaround for an erratum in Cortex-A8 processors. If you are targeting an affected part and wish to use the workaround, pass the `--fix-cortex-a8` option to the linker. Please contact ARM for further details of the erratum.

Maximum code alignment increased. The maximum allowed code alignment has been increased from 32 to 64 bytes. This change affects the `.p2align` and `.align` directives in GAS and the `-falign-functions` GCC option.

Corruption of block-scope variables. A compiler optimization bug that sometimes caused corruption of stack-allocated variables has been fixed. The bug affected variables declared in a local block scope in functions containing multiple non-overlapping lexical block scopes, a technique commonly used by programmers to reduce stack frame size. In some rare cases, other optimizations performed by the compiler were ignoring the local extent of such block-scope variables.

A.1.5. Changes in Sourcery G++ Lite 2009q1-116

Linking big-endian programs for ARMv7-A. When linking for ARMv7-A targets with `-mbig-endian`, Sourcery G++ now implicitly assumes BE8 mode, rather than BE32.

GCC version 4.3.3. Sourcery G++ Lite for ARM EABI is now based on GCC version 4.3.3. This is a bug fix update to GCC. For more information about changes from GCC version 4.3.2 that was included in previous releases, see <http://gcc.gnu.org/gcc-4.3/changes.html>.

Improved NOP generation for Thumb-2 cores. The assembler now generates Thumb-2/ARMv6K architectural NOP instructions when alignment padding is required in code sections.

Internal compiler error with `-O3` or `-fpredictive-commoning`. A bug has been fixed that caused internal compiler errors when compiling some code with `-O3` or `-fpredictive-commoning`.

CS3 board and processor support. CS3 board and processor support has been cleaned up to remove entries that are not appropriate for or supported by Sourcery G++ Lite on ARM EABI targets. This includes processors for which Sourcery G++ Lite does not include appropriate run-time libraries. These changes are intended to simplify processor and board selection. For the full list of boards supported by CS3, refer to Chapter 5, “CS3™: The CodeSourcery Common Startup Code Sequence”.

C++ named operators bug fix. A bug has been fixed that caused the compiler to crash in some cases when the C++ operators `and_eq`, `bitand`, `bitor`, `compl`, `not_eq`, `or_eq` and `xor_eq` were used in contexts where the preprocessor converts their names to strings.

Debug information for anonymous structure types. A GCC bug in the generation of debug information for anonymous structure types in C++ code has been fixed. The bug caused printing the type information for such structures in the debugger (via the **ptype** command) to fail with an error message.

CS3 bug fix. A bug in CS3 has been fixed that caused a write to invalid address on program startup for all hosted targets.

Altera device detection. A bug in the Sourcery G++ Debug Sprite has been fixed. The bug showed a spurious error when an Altera device without an SLD hub was connected to the host.

Interrupting the target from the debugger. GDB has been improved to be more responsive to attempts to interrupt the target (as by a **Ctrl+C** when using GDB from the command line) during execution of programs using semihosting.

Linker errors on non-ELF input. A bug has been fixed that caused internal errors from the linker when linking non-ELF input files (with the `-b` or `--format` linker options).

Undefined weak references in shared libraries. A linker bug has been fixed affecting calls from Thumb code in shared libraries to functions that are undefined weak references when the shared library is linked. Such calls executed as nops whether or not the functions were defined at run time.

Improved code generation. The compiler has been improved to generate better code for an integer multiplication whose result feeds into an addition.

Newlib update. The Newlib package has been updated to version 1.17.0, with additions from the community CVS trunk as of 2009-02-24. This update provides new C99 wide-character functions; POSIX regex functions; string-function performance improvements; an improved `sprintf` implementation that no longer requires I/O functions like `_open`, `_write`, and `_close`; and other bug fixes and improvements. For more information, refer to the Newlib C Library and Math Library manuals, and to the Newlib web site at <http://sourceware.org/newlib/>.

Installer fails during upgrade. The Sourcery G++ installer for Microsoft Windows hosts could fail during an upgrade while waiting for the previous version to be uninstalled. This bug has been fixed.

Performance improvements. Tuning parameters for ARM code generation have been adjusted to improve performance of the generated code.

Uninstaller removed by upgrade. The uninstaller could be incorrectly deleted during an upgrade on Microsoft Windows hosts. This bug has been fixed.

Remote debugging connection auto-retry. The `target remote` command within GDB now uses a configurable auto-retry timeout when establishing TCP connections. This is useful in avoiding race conditions when the remote GDB stub or GDB server is launched simultaneously with GDB. The auto-retry behavior is enabled by default; refer to the GDB manual for details.

CMP Thumb-2 instruction. The assembler no longer issues an error about CMP instructions in which the second argument is the stack pointer (`r13`), as these are valid instructions. However, use of the stack pointer in this context is deprecated in the current ARM architecture specification and the assembler now warns about the deprecated use.

Thumb half-precision floating point bug fix. A compiler bug has been fixed that formerly caused incorrect code to be generated in Thumb mode for functions using half-precision floating-point constants. The bug did not affect Thumb-2 code.

Improved code generation. The compiler has been improved to generate better code for integer multiplication by certain constants.

Support for Keil MCB2130 and MCB2140 boards. CS3 now includes support for Keil MCB2130 and MCB2140 boards.

Thumb-2 switch code generation bug fix. A bug has been fixed that caused incorrect Thumb-2 code to be generated for some `switch` statements.

Internal compiler errors when optimizing. A defect that occasionally caused internal compiler errors when partial redundancy elimination (PRE) optimization was enabled has been corrected.

Install directory pathnames. Bugs in the install and uninstall scripts for Linux hosts that caused errors or incorrect behavior when the Sourcery G++ install directory pathname contains whitespace characters have been fixed.

Internal compiler error with large NEON types. A bug has been fixed that caused internal compiler errors when compiling code using NEON types at least 32 bytes wide.

Temporary files on Microsoft Windows. On Microsoft Windows hosts, Sourcery G++ Lite now uses the standard Windows algorithm to choose the directory in which to place temporary files. This change eliminates a crash that occurred if none of the `TEMP`, `TMP`, or `TMPDIR` variables were set to a suitable directory.

Vectorized shift fix. A bug has been fixed that caused incorrect code for loops containing a right shift by a constant. The bug affected code compiled with `-mfpu=neon` and loop vectorization enabled with `-O3` or `-ftree-vectorize`.

Incorrect code for nested functions. A bug in GCC that caused the compiler to generate incorrect code for nested functions has been fixed. The bug resulted in incorrect stack alignments in the affected functions.

Binutils update. The binutils package has been updated to version 2.19.51.20090205 from the FSF trunk. This update includes numerous bug fixes.

ARM build attributes conformance improvements. Several ARM EABI 2.07 conformance issues relating to the handling of build attributes in the assembler and linker have been fixed. All build attribute types are now recognized, and can now be declared by name, in addition to by number. Support for merging attributes in the linker has been improved, and the linking of incompatible objects is now detected and rejected in more cases.

VFP initialization for i.MX31. The CS3 startup code for i.MX31 now automatically enables the ARM VFP coprocessor.

Internal compiler error with `-fremove-local-statics`. An internal compiler error that occurred when using the `-fremove-local-statics` option has been fixed. The error occurred when compiling code with function-local `static` array or structure variables.

GDB update. The included version of GDB has been updated to 6.8.50.20081022. This update includes numerous bug fixes.

Linker crash on incompatible input files. Some third-party compilers, including ARM RealView® 4.0, produce a build attribute marking output files that are not compatible with the ABI for the ARM Architecture. This attribute sometimes caused the linker to crash. The linker now correctly issues an error message.

A.1.6. Changes in Sourcery G++ Lite 2008q3-66

Bug fix for assembly listing. A bug that caused the assembler to produce corrupted listings (via the `-a` option) on Windows hosts has been fixed.

Reduced RAM usage for semihosting. The previous change to add support for command-line arguments in semihosted applications increased RAM requirements for all programs using semihosting. This feature has been reimplemented to substantially reduce the amount of additional memory required.

Optimizer bug fix. A bug that caused an unrecognizable `insn` internal compiler error when compiling at optimization levels above `-O0` has been fixed.

VFP compiler fix. A compiler bug that resulted in `internal compiler error: output_operand: invalid expression as operand` when generating VFP code has been fixed.

GDB display of source. A bug has been fixed that prevented GDB from locating debug information in some cases. The debugger failed to display source code for or step into the affected functions.

Profiling support added. The `-pg` option is now supported by the compiler. You are required to provide a function named `__gnu_mcount_nc`. For more details, see Section 3.9, “ARM Profiling Implementation”.

Workaround for Cortex-M3 CPU errata. Errata present in some Cortex-M3 cores can cause data corruption when overlapping registers are used in `LDRD` instructions. The compiler avoids generating these problematic instructions when the `-mfix-cortex-m3-ldrd` or `-mcpu=cortex-m3` command-line options are used. The Sourcery G++ runtime libraries have also been updated to include this workaround.

GDB segment warning. Some compilers produce binaries including uninitialized data regions, such as the stack and heap. GDB incorrectly displayed the warning `Loadable segment "name" outside of ELF segments` for such binaries; the warning has now been fixed.

Misaligned NEON memory accesses. A bug has been fixed that caused the compiler to use aligned NEON load/store instructions to access misaligned data when autovectorizing certain loops. The bug affected code compiled with `-mfpu=neon` and loop vectorization enabled with `-O3` or `-ftree-vectorize`.

Sprite crash on error. A bug has been fixed which sometimes caused the Sourcery G++ Debug Sprite to crash when it attempted to send an error message to GDB.

Portable object file fixes. Bugs in the `limits.h`, `stdlib.h`, `time.h`, and `setjmp.h` headers have been fixed. These bugs caused compilation errors when using `-D_AEABI_PORTABILITY_LEVEL=1`, as described in Section 3.8, “Object File Portability”.

Persistent remote server connections. A GDB bug has been fixed that caused the **target extended-remote** command to fail to tell the remote server to make the connection persistent across program invocations.

A.1.7. Changes in Sourcery G++ Lite 2008q3-39

Definition of `va_list`. In order to conform to the ABI for the ARM Architecture, the definition of the type of `va_list` (defined in `stdarg.h`) has been changed. This change impacts only the mangled names of C++ entities. For example, the mangled name of a C++ function taking an argument of type `va_list`, or `va_list *`, or another type involving `va_list` has changed. Since this is an incompatible change, you must recompile and relink any modules defining or using affected `va_list`-typed entities.

Thumb-2 assembler fixes. The Thumb-2 encodings of `QADD`, `QDADD`, `QSUB`, and `QDSUB` have been corrected. Previous versions of the assembler generated incorrect object files for these instructions. The assembler now accepts the `ORN`, `QASX`, `QSAX`, `RRX`, `SHASX`, `SHSAX`, `SSAX`, `USAX`, `UHASX`, `UQSAX`, and `USAX` mnemonics. The assembler now detects and issues errors for invalid uses of register 13 (the stack pointer) and register 15 (the program counter) in many instructions.

Printing casted values in GDB. A GDB bug that caused incorrect output for expressions containing casts, such as in the `print *(Type *)ptr` command, has been fixed.

Bug fix for `objcopy/strip`. An `objcopy` bug that corrupted `COMDAT` groups when creating new binaries has been fixed. This bug also affected `strip -g`.

Improved support for debugging RealView® objects . GDB support for programs compiled by the ARM RealView® compiler has been improved.

Binutils support for DWARF Version 3. The `addr2line` command now supports binaries containing DWARF 3 debugging information. The `ld` command can display error messages with source locations for input files containing DWARF 3 debugging information.

NEON improvements. Several improvements and bug fixes have been made to the NEON Advanced SIMD Extension support in GCC. A problem that caused the autovectorizer to fail in some circumstances has been fixed. Also, many of the intrinsics available via the `arm_neon.h` header file now have improved error checking for out-of-bounds arguments, and the `vget_lane` intrinsics that return signed values now produce improved code.

NEON compiler fix. A compiler bug that resulted in incorrect NEON code being generated has been fixed. Typically the incorrect code occurred when NEON intrinsics were used inside small `if` statements.

Connecting to the target using a pipe. A bug in GDB's `target remote | program` command has been fixed. When launching the specified `program` failed, the bug caused GDB to crash, hang, or give a message `Error: No Error`.

Mixed-case NEON register aliases. An assembler bug that prevented NEON register aliases from being created with mixed-case names using the `.dn` and `.qn` directives has been fixed. Previously only aliases created with all-lowercase or all-uppercase names worked correctly.

Newlib manuals. The documentation packaged with Sourcery G++ Lite now includes the Newlib C Library and Math Library manuals.

Object file portability. Sourcery G++ for ARM EABI can now produce object files that are link-compatible with the GNU C Library included with Sourcery G++ for ARM GNU/Linux. These object files are additionally link-compatible with other ARM C Library ABI-compliant static linking environments and toolchains. For additional information, refer to Section 3.8, “Object File Portability”.

Janus 2CC support. GCC now includes a work-around for a hardware bug in Avalent Janus 2CC cores. To compile and link for these cores, use the `-mfix-janus-2cc` compiler option. If you are using the linker directly use the `--fix-janus-2cc` linker option.

ARM exception handling bug fix. A bug in the runtime library has been fixed that formerly caused throwing an unexpected exception in C++ to crash instead of calling the unexpected exception handler. The bug only affected C++ code compiled by non-GNU compilers such as ARM RealView®.

Simulation of programs with command-line arguments. When using the simulator, command-line arguments are now correctly passed to the simulator program via `argc` and `argv`.

Mangling of NEON type names. A bug in the algorithm used by the C++ compiler for mangling the names of NEON types, such as `int8x16_t`, has been fixed. These mangled names are used internally in object files to encode type information in addition to the programmer-visible names of the C++ variables and functions. The new mangled name encoding is more compact and conforms to the ARM C++ ABI.

Errors after loading the debugged program. An intermittent GDB bug has been fixed. The bug could cause a GDB internal error after the `load` command.

Half-precision floating point. Sourcery G++ now includes support for half-precision floating point via the `__fp16` type in C and C++. The compiler can generate code using either hardware support or library routines. For more information, see Section 3.6.3, “Half-Precision Floating Point”.

A.1.8. Changes in Sourcery G++ Lite 2008q3-12

Stellaris UDMAERR interrupt vector. The name of the UDMAERR (uDMA Error) interrupt vector provided by CS3 for Luminary Micro Stellaris configurations has been corrected. The correct name is `__cs3_isr_udmaerr`.

GDB update. The included version of GDB has been updated to 6.8.50.20080821. This update adds numerous bug fixes and new features, including support for decimal floating point, improved Thumb mode support, the new `find` command to search memory, the new `/m` (mixed source and assembly) option to the `disassemble` command, and the new `macro define` command to define C preprocessor macros interactively.

Uppercase operands to IT instructions. The assembler now accepts both uppercase and lowercase operands for the IT family of instructions.

NEON autovectorizer fix. A compiler bug that caused generation of bad VLD1 instructions has been fixed. The bug affected code compiled with `-mfpu=neon -ftree-vectorize`.

Bug fix in simulator reset code. A bug which could cause programs to crash at startup has been fixed in the CS3 reset code used for simulators. The affected linker scripts are `generic.ld`, `generic-vfp.ld`, `generic-m.ld`, and hosted versions of the above.

Output files removed on error. When GCC encounters an error, it now consistently removes any incomplete output files that it may have created.

Atmel AT91SAM7S-EK support. Support for Atmel AT91SAM7S-EK boards has been added.

Placing bss-like regions in load regions. The linker no longer issues an incorrect error message when a bss-like section is placed at specific load region. The linker formerly incorrectly considered the section as taking up space in the load region.

ARMv7 offset out of range errors. An assembler bug that resulted in `offset out of range` errors when compiling for ARMv7 processors has been fixed.

Thumb-2 MUL encoding. In Thumb-2 mode, the assembler now encodes `MUL` as a 16-bit instruction (rather than as a 32-bit instruction) when possible. This fix results in smaller code, with no loss of performance.

ARM C++ ABI utility functions. Vector utility functions required by the ARM C++ ABI no longer crash when passed null pointers. The affected functions are `__aeabi_vec_dtor_cookie`, `__aeabi_vec_delete`, `__aeabi_vec_delete3`, and `__aeabi_vec_delete3_nodtor`. These functions are not intended for use by application programmers; they are only called by compiler-generated code. They are not presently used by the GNU C++ compiler, but are used by some other compilers, including ARM's RealView® compiler.

GCC version 4.3.2. Sourcery G++ Lite for ARM EABI is now based on GCC version 4.3.2. For more information about changes from GCC version 4.2 that was included in previous releases, see <http://gcc.gnu.org/gcc-4.3/changes.html>.

Smaller Thumb-2 code. When optimizing for size (i.e., when `-Os` is in use), GCC now generates the 16-bit `MULS` Thumb-2 multiply instruction instead of the 32-bit `MUL` instruction.

Thumb-2 RBIT encoding. An assembler bug that resulted in incorrect encoding of the Thumb-2 `RBIT` instruction has been fixed.

Additional Stellaris interrupt vectors. The `USB0`, `PWM3`, `UDMA` and `UDMAERR` interrupt vectors have been added to the CS3 Luminary Micro Stellaris configurations.

Additional Luminary Micro CPUs. A large number of additional Luminary Micro CPUs are now supported. Linker scripts for these CPUs include the internal ROM region.

Sprite communication improvements. The Sourcery G++ Debug Sprite now uses a more efficient protocol for communicating with GDB. This can result in less latency when debugging, especially when running the Sprite on a remote machine over a network connection.

Marvell Feroceon compiler bug fix. A bug that caused an internal compiler error when optimizing for Marvell Feroceon CPUs has been fixed.

Misaligned accesses to packed structures fix. A bug that caused GCC to generate misaligned accesses to packed structures has been fixed.

SIZE_MIN and SIZE_MAX. `SIZE_MAX` is now correctly defined in `stdint.h` to be an unsigned quantity equal to the maximum possible value of a `size_t`, as required by the C standard. In addition, the definition of `SIZE_MIN` has been removed, as this constant is not prescribed by the C standard.

Bug fix for objdump on Windows. An `objdump` bug that caused the `-S` option not to work on Windows in some cases has been fixed.

A.1.9. Changes in Older Releases

For information about changes in older releases of Sourcery G++ Lite for ARM EABI, please refer to the Getting Started guide packaged with those releases.

Appendix B

Sourcery G++ Lite Licenses

Sourcery G++ Lite contains software provided under a variety of licenses. Some components are “free” or “open source” software, while other components are proprietary. This appendix explains what licenses apply to your use of Sourcery G++ Lite. You should read this appendix to understand your legal rights and obligations as a user of Sourcery G++ Lite.

B.1. Licenses for Sourcery G++ Lite Components

The table below lists the major components of Sourcery G++ Lite for ARM EABI and the license terms which apply to each of these components.

Some free or open-source components provide documentation or other files under terms different from those shown below. For definitive information about the license that applies to each component, consult the source package corresponding to this release of Sourcery G++ Lite. Sourcery G++ Lite may contain free or open-source components not included in the list below; for a definitive list, consult the source package corresponding to this release of Sourcery G++ Lite.

Component	License
GNU Compiler Collection	GNU General Public License 3.0 ¹
GNU Binary Utilities	GNU General Public License 3.0 ²
GNU Debugger	GNU General Public License 3.0 ³
Sourcery G++ Debug Sprite for ARM	CodeSourcery License
CodeSourcery Common Startup Code Sequence	CodeSourcery License
Newlib C Library	BSD License. For the text of the license and a complete list of copyright holders, see COPYING.NEWLIB included in the source package.
GNU Make	GNU General Public License 2.0 ⁴
GNU Core Utilities	GNU General Public License 2.0 ⁵

The CodeSourcery License is available in Section B.2, “Sourcery G++ Software License Agreement”.

Important

Although some of the licenses that apply to Sourcery G++ Lite are “free software” or “open source software” licenses, none of these licenses impose any obligation on you to reveal the source code of applications you build with Sourcery G++ Lite. You can develop proprietary applications and libraries with Sourcery G++ Lite.

Sourcery G++ Lite may include some third party example programs and libraries in the `share/sourceryg++-arm-none-eabi-examples` subdirectory. These examples are not covered by the Sourcery G++ Software License Agreement. To the extent permitted by law, these examples are provided by CodeSourcery as is with no warranty of any kind, including implied warranties of merchantability or fitness for a particular purpose. Your use of each example is governed by the license notice (if any) it contains.

¹ <http://www.gnu.org/licenses/gpl.html>

² <http://www.gnu.org/licenses/gpl.html>

³ <http://www.gnu.org/licenses/gpl.html>

⁴ <http://www.gnu.org/licenses/old-licenses/gpl-2.0.html>

⁵ <http://www.gnu.org/licenses/old-licenses/gpl-2.0.html>

B.2. Sourcery G++™ Software License Agreement

1. **Parties.** The parties to this Agreement are you, the licensee (“You” or “Licensee”) and CodeSourcery. If You are not acting on behalf of Yourself as an individual, then “You” means Your company or organization.
2. **The Software.** The Software licensed under this Agreement consists of computer programs and documentation referred to as Sourcery G++™ Lite Edition (the “Software”).
3. **Definitions.**
 - 3.1. **CodeSourcery Proprietary Components.** The components of the Software that are owned and/or licensed by CodeSourcery and are not subject to a “free software” or “open source” license, such as the GNU Public License. The CodeSourcery Proprietary Components of the Software include, without limitation, the Sourcery G++ Installer, any Sourcery G++ Eclipse plug-ins, and any Sourcery G++ Debug Sprite. For a complete list, refer to the *Getting Started Guide* included with the distribution.
 - 3.2. **Open Source Software Components.** The components of the Software that are subject to a “free software” or “open source” license, such as the GNU Public License.
 - 3.3. **Proprietary Rights.** All rights in and to copyrights, rights to register copyrights, trade secrets, inventions, patents, patent rights, trademarks, trademark rights, confidential and proprietary information protected under contract or otherwise under law, and other similar rights or interests in intellectual or industrial property.
 - 3.4. **Redistributable Components.** The CodeSourcery Proprietary Components that are intended to be incorporated or linked into Licensee object code developed with the Software. The Redistributable Components of the Software include, without limitation, the CSLIBC run-time library and the CodeSourcery Common Startup Code Sequence (CS3). For a complete list, refer to the *Getting Started Guide* included with the distribution.
4. **License Grant to Proprietary Components of the Software.** You are granted a non-exclusive, royalty-free license (a) to install and use the CodeSourcery Proprietary Components of the Software, (b) to transmit the CodeSourcery Proprietary Components over an internal computer network, (c) to copy the CodeSourcery Proprietary Components for Your internal use only, and (d) to distribute the Redistributable Component(s) in binary form only and only as part of Licensee object code developed with the Software that provides substantially different functionality than the Redistributable Component(s).
5. **Restrictions.** You may not: (i) copy or permit others to use the CodeSourcery Proprietary Components of the Software, except as expressly provided above; (ii) distribute the CodeSourcery Proprietary Components of the Software to any third party, except as expressly provided above; or (iii) reverse engineer, decompile, or disassemble the CodeSourcery Proprietary Components of the Software, except to the extent this restriction is expressly prohibited by applicable law.
6. **“Free Software” or “Open Source” License to Certain Components of the Software.** This Agreement does not limit Your rights under, or grant You rights that supersede, the license terms of any Open Source Software Component delivered to You by CodeSourcery. Sourcery G++ includes components provided under various different licenses. The *Getting Started Guide* provides an overview of which license applies to different components. Definitive licensing

information for each “free software” or “open source” component is available in the relevant source file.

7. **CodeSourcery Trademarks.** Notwithstanding any provision in a “free software” or “open source” license agreement applicable to a component of the Software that permits You to distribute such component to a third party in source or binary form, You may not use any CodeSourcery trademark, whether registered or unregistered, including without limitation, CodeSourcery™, Sourcery G++™, the CodeSourcery crystal ball logo, or the Sourcery G++ splash screen, or any confusingly similar mark, in connection with such distribution, and You may not recompile the Open Source Software Components with the `--with-pkgversion` or `--with-bugurl` configuration options that embed CodeSourcery trademarks in the resulting binary.
8. **Term and Termination.** This Agreement shall remain in effect unless terminated pursuant to this provision. CodeSourcery may terminate this Agreement upon seven (7) days written notice of a material breach of this Agreement if such breach is not cured; provided that the unauthorized use, copying, or distribution of the CodeSourcery Proprietary Components of the Software will be deemed a material breach that cannot be cured.
9. **Transfers.** You may not transfer any rights under this Agreement without the prior written consent of CodeSourcery, which consent shall not be unreasonably withheld. A condition to any transfer or assignment shall be that the recipient agrees to the terms of this Agreement. Any attempted transfer or assignment in violation of this provision shall be null and void.
10. **Ownership.** CodeSourcery owns and/or has licensed the CodeSourcery Proprietary Components of the Software and all intellectual property rights embodied therein, including copyrights and valuable trade secrets embodied in its design and coding methodology. The CodeSourcery Proprietary Components of the Software are protected by United States copyright laws and international treaty provisions. CodeSourcery also owns all rights, title and interest in and with respect to its trade names, domain names, trade dress, logos, trademarks, service marks, and other similar rights or interests in intellectual property. This Agreement provides You only a limited use license, and no ownership of any intellectual property.
11. **Warranty Disclaimer; Limitation of Liability.** CODESOURCERY AND ITS LICENSORS PROVIDE THE SOFTWARE “AS-IS” AND PROVIDED WITH ALL FAULTS. CODESOURCERY DOES NOT MAKE ANY WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. CODESOURCERY SPECIFICALLY DISCLAIMS THE IMPLIED WARRANTIES OF TITLE, NON-INFRINGEMENT, MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, SYSTEM INTEGRATION, AND DATA ACCURACY. THERE IS NO WARRANTY OR GUARANTEE THAT THE OPERATION OF THE SOFTWARE WILL BE UNINTERRUPTED, ERROR-FREE, OR VIRUS-FREE, OR THAT THE SOFTWARE WILL MEET ANY PARTICULAR CRITERIA OF PERFORMANCE, QUALITY, ACCURACY, PURPOSE, OR NEED. YOU ASSUME THE ENTIRE RISK OF SELECTION, INSTALLATION, AND USE OF THE SOFTWARE. THIS DISCLAIMER OF WARRANTY CONSTITUTES AN ESSENTIAL PART OF THIS AGREEMENT. NO USE OF THE SOFTWARE IS AUTHORIZED HEREUNDER EXCEPT UNDER THIS DISCLAIMER.
12. **Local Law.** If implied warranties may not be disclaimed under applicable law, then ANY IMPLIED WARRANTIES ARE LIMITED IN DURATION TO THE PERIOD REQUIRED BY APPLICABLE LAW.
13. **Limitation of Liability.** INDEPENDENT OF THE FORGOING PROVISIONS, IN NO EVENT AND UNDER NO LEGAL THEORY, INCLUDING WITHOUT LIMITATION, TORT, CONTRACT, OR STRICT PRODUCTS LIABILITY, SHALL CODESOURCERY BE LIABLE TO YOU OR ANY OTHER PERSON FOR ANY INDIRECT, SPECIAL, INCID-

ENTAL, OR CONSEQUENTIAL DAMAGES OF ANY KIND, INCLUDING WITHOUT LIMITATION, DAMAGES FOR LOSS OF GOODWILL, WORK STOPPAGE, COMPUTER MALFUNCTION, OR ANY OTHER KIND OF COMMERCIAL DAMAGE, EVEN IF CODESOURCERY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION SHALL NOT APPLY TO LIABILITY FOR DEATH OR PERSONAL INJURY TO THE EXTENT PROHIBITED BY APPLICABLE LAW. IN NO EVENT SHALL CODESOURCERY'S LIABILITY FOR ACTUAL DAMAGES FOR ANY CAUSE WHATSOEVER, AND REGARDLESS OF THE FORM OF ACTION, EXCEED THE AMOUNT PAID BY YOU IN FEES UNDER THIS AGREEMENT DURING THE PREVIOUS ONE YEAR PERIOD.

14. **Export Controls.** You agree to comply with all export laws and restrictions and regulations of the United States or foreign agencies or authorities, and not to export or re-export the Software or any direct product thereof in violation of any such restrictions, laws or regulations, or without all necessary approvals. As applicable, each party shall obtain and bear all expenses relating to any necessary licenses and/or exemptions with respect to its own export of the Software from the U.S. Neither the Software nor the underlying information or technology may be electronically transmitted or otherwise exported or re-exported (i) into Cuba, Iran, Iraq, Libya, North Korea, Sudan, Syria or any other country subject to U.S. trade sanctions covering the Software, to individuals or entities controlled by such countries, or to nationals or residents of such countries other than nationals who are lawfully admitted permanent residents of countries not subject to such sanctions; or (ii) to anyone on the U.S. Treasury Department's list of Specially Designated Nationals and Blocked Persons or the U.S. Commerce Department's Table of Denial Orders. By downloading or using the Software, Licensee agrees to the foregoing and represents and warrants that it complies with these conditions.
15. **U.S. Government End-Users.** The Software is a "commercial item," as that term is defined in 48 C.F.R. 2.101 (Oct. 1995), consisting of "commercial computer software" and "commercial computer software documentation," as such terms are used in 48 C.F.R. 12.212 (Sept. 1995). Consistent with 48 C.F.R. 12.212 and 48 C.F.R. 227.7202-1 through 227.7202-4 (June 1995), all U.S. Government End Users acquire the Software with only those rights set forth herein.
16. **Licensee Outside The U.S.** If You are located outside the U.S., then the following provisions shall apply: (i) Les parties aux presentes confirment leur volonte que cette convention de meme que tous les documents y compris tout avis qui s'y rattache, soient rediges en langue anglaise (translation: "The parties confirm that this Agreement and all related documentation is and will be in the English language."); and (ii) You are responsible for complying with any local laws in your jurisdiction which might impact your right to import, export or use the Software, and You represent that You have complied with any regulations or registration procedures required by applicable law to make this license enforceable.
17. **Severability.** If any provision of this Agreement is declared invalid or unenforceable, such provision shall be deemed modified to the extent necessary and possible to render it valid and enforceable. In any event, the unenforceability or invalidity of any provision shall not affect any other provision of this Agreement, and this Agreement shall continue in full force and effect, and be construed and enforced, as if such provision had not been included, or had been modified as above provided, as the case may be.
18. **Arbitration.** Except for actions to protect intellectual property rights and to enforce an arbitrator's decision hereunder, all disputes, controversies, or claims arising out of or relating to this Agreement or a breach thereof shall be submitted to and finally resolved by arbitration under the rules of the American Arbitration Association ("AAA") then in effect. There shall be one arbitrator, and such arbitrator shall be chosen by mutual agreement of the parties in accordance with AAA rules. The arbitration shall take place in Granite Bay, California, and may be conducted

by telephone or online. The arbitrator shall apply the laws of the State of California, USA to all issues in dispute. The controversy or claim shall be arbitrated on an individual basis, and shall not be consolidated in any arbitration with any claim or controversy of any other party. The findings of the arbitrator shall be final and binding on the parties, and may be entered in any court of competent jurisdiction for enforcement. Enforcements of any award or judgment shall be governed by the United Nations Convention on the Recognition and Enforcement of Foreign Arbitral Awards. Should either party file an action contrary to this provision, the other party may recover attorney's fees and costs up to \$1000.00.

19. **Jurisdiction And Venue.** The courts of Placer County in the State of California, USA and the nearest U.S. District Court shall be the exclusive jurisdiction and venue for all legal proceedings that are not arbitrated under this Agreement.
20. **Independent Contractors.** The relationship of the parties is that of independent contractor, and nothing herein shall be construed to create a partnership, joint venture, franchise, employment, or agency relationship between the parties. Licensee shall have no authority to enter into agreements of any kind on behalf of CodeSourcery and shall not have the power or authority to bind or obligate CodeSourcery in any manner to any third party.
21. **Force Majeure.** Neither CodeSourcery nor Licensee shall be liable for damages for any delay or failure of delivery arising out of causes beyond their reasonable control and without their fault or negligence, including, but not limited to, Acts of God, acts of civil or military authority, fires, riots, wars, embargoes, or communications failure.
22. **Miscellaneous.** This Agreement constitutes the entire understanding of the parties with respect to the subject matter of this Agreement and merges all prior communications, representations, and agreements. This Agreement may be modified only by a written agreement signed by the parties. If any provision of this Agreement is held to be unenforceable for any reason, such provision shall be reformed only to the extent necessary to make it enforceable. This Agreement shall be construed under the laws of the State of California, USA, excluding rules regarding conflicts of law. The application of the United Nations Convention of Contracts for the International Sale of Goods is expressly excluded. This license is written in English, and English is its controlling language.