

GNU gprof

The GNU Profiler
(Sourcery G++ Lite 4.5-97)
Version 2.20.51

Jay Fenlason and Richard Stallman

This manual describes the GNU profiler, `gprof`, and how you can use it to determine which parts of a program are taking most of the execution time. We assume that you know how to write, compile, and execute programs. GNU `gprof` was written by Jay Fenlason. Eric S. Raymond made some minor corrections and additions in 2003.

Copyright © 1988, 1992, 1997, 1998, 1999, 2000, 2003, 2008, 2009 Free Software Foundation, Inc.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation; with no Invariant Sections, with no Front-Cover Texts, and with no Back-Cover Texts. A copy of the license is included in the section entitled “GNU Free Documentation License”.

Table of Contents

1	Introduction to Profiling	1
2	Compiling a Program for Profiling	3
3	Executing the Program	5
4	gprof Command Summary	7
4.1	Output Options	7
4.2	Analysis Options	10
4.3	Miscellaneous Options	12
4.4	Deprecated Options	13
4.5	Symspecs	14
5	Interpreting gprof's Output	15
5.1	The Flat Profile	15
5.2	The Call Graph	17
5.2.1	The Primary Line	18
5.2.2	Lines for a Function's Callers	19
5.2.3	Lines for a Function's Subroutines	19
5.2.4	How Mutually Recursive Functions Are Described	20
5.3	Line-by-line Profiling	22
5.4	The Annotated Source Listing	24
6	Inaccuracy of gprof Output	27
6.1	Statistical Sampling Error	27
6.2	Estimating children Times	28
7	Answers to Common Questions	29
8	Incompatibilities with Unix gprof	31
9	Details of Profiling	33
9.1	Implementation of Profiling	33
9.2	Profiling Data File Format	34
9.2.1	Histogram Records	35
9.2.2	Call-Graph Records	35
9.2.3	Basic-Block Execution Count Records	36
9.3	gprof's Internal Operation	36
9.4	Debugging gprof	39

Appendix A GNU Free Documentation License
..... **41**

1 Introduction to Profiling

Profiling allows you to learn where your program spent its time and which functions called which other functions while it was executing. This information can show you which pieces of your program are slower than you expected, and might be candidates for rewriting to make your program execute faster. It can also tell you which functions are being called more or less often than you expected. This may help you spot bugs that had otherwise been unnoticed.

Since the profiler uses information collected during the actual execution of your program, it can be used on programs that are too large or too complex to analyze by reading the source. However, how your program is run will affect the information that shows up in the profile data. If you don't use some feature of your program while it is being profiled, no profile information will be generated for that feature.

Profiling has several steps:

- You must compile and link your program with profiling enabled. See [Chapter 2 \[Compiling a Program for Profiling\]](#), page 3.
- You must execute your program to generate a profile data file. See [Chapter 3 \[Executing the Program\]](#), page 5.
- You must run `gprof` to analyze the profile data. See [Chapter 4 \[gprof Command Summary\]](#), page 7.

The next three chapters explain these steps in greater detail.

Several forms of output are available from the analysis.

The *flat profile* shows how much time your program spent in each function, and how many times that function was called. If you simply want to know which functions burn most of the cycles, it is stated concisely here. See [Section 5.1 \[The Flat Profile\]](#), page 15.

The *call graph* shows, for each function, which functions called it, which other functions it called, and how many times. There is also an estimate of how much time was spent in the subroutines of each function. This can suggest places where you might try to eliminate function calls that use a lot of time. See [Section 5.2 \[The Call Graph\]](#), page 17.

The *annotated source* listing is a copy of the program's source code, labeled with the number of times each line of the program was executed. See [Section 5.4 \[The Annotated Source Listing\]](#), page 24.

To better understand how profiling works, you may wish to read a description of its implementation. See [Section 9.1 \[Implementation of Profiling\]](#), page 33.

2 Compiling a Program for Profiling

The first step in generating profile information for your program is to compile and link it with profiling enabled.

To compile a source file for profiling, specify the `-pg` option when you run the compiler. (This is in addition to the options you normally use.)

To link the program for profiling, if you use a compiler such as `cc` to do the linking, simply specify `-pg` in addition to your usual options. The same option, `-pg`, alters either compilation or linking to do what is necessary for profiling. Here are examples:

```
cc -g -c myprog.c utils.c -pg
cc -o myprog myprog.o utils.o -pg
```

The `-pg` option also works with a command that both compiles and links:

```
cc -o myprog myprog.c utils.c -g -pg
```

Note: The `-pg` option must be part of your compilation options as well as your link options. If it is not then no call-graph data will be gathered and when you run `gprof` you will get an error message like this:

```
gprof: gmon.out file is missing call-graph data
```

If you add the `-Q` switch to suppress the printing of the call graph data you will still be able to see the time samples:

Flat profile:

```
Each sample counts as 0.01 seconds.
%   cumulative   self           self         total
time  seconds    seconds   calls   Ts/call   Ts/call   name
44.12    0.07      0.07                44.12    0.07   zazLoop
35.29    0.14      0.06                35.29    0.06   main
20.59    0.17      0.04                20.59    0.04   bazMillion
```

If you run the linker `ld` directly instead of through a compiler such as `cc`, you may have to specify a profiling startup file `gcrt0.o` as the first input file instead of the usual startup file `crt0.o`. In addition, you would probably want to specify the profiling C library, `libc_p.a`, by writing `-lc_p` instead of the usual `-lc`. This is not absolutely necessary, but doing this gives you number-of-calls information for standard library functions such as `read` and `open`. For example:

```
ld -o myprog /lib/gcrt0.o myprog.o utils.o -lc_p
```

If you are running the program on a system which supports shared libraries you may run into problems with the profiling support code in a shared library being called before that library has been fully initialised. This is usually detected by the program encountering a segmentation fault as soon as it is run. The solution is to link against a static version of the library containing the profiling support code, which for `gcc` users can be done via the `-static` or `-static-libgcc` command line option. For example:

```
gcc -g -pg -static-libgcc myprog.c utils.c -o myprog
```

If you compile only some of the modules of the program with ‘-pg’, you can still profile the program, but you won’t get complete information about the modules that were compiled without ‘-pg’. The only information you get for the functions in those modules is the total time spent in them; there is no record of how many times they were called, or from where. This will not affect the flat profile (except that the `calls` field for the functions will be blank), but will greatly reduce the usefulness of the call graph.

If you wish to perform line-by-line profiling you should use the `gcov` tool instead of `gprof`. See that tool’s manual or info pages for more details of how to do this.

Note, older versions of `gcc` produce line-by-line profiling information that works with `gprof` rather than `gcov` so there is still support for displaying this kind of information in `gprof`. See [Section 5.3 \[Line-by-line Profiling\]](#), [page 22](#).

It also worth noting that `gcc` implements a ‘-finstrument-functions’ command line option which will insert calls to special user supplied instrumentation routines at the entry and exit of every function in their program. This can be used to implement an alternative profiling scheme.

3 Executing the Program

Once the program is compiled for profiling, you must run it in order to generate the information that `gprof` needs. Simply run the program as usual, using the normal arguments, file names, etc. The program should run normally, producing the same output as usual. It will, however, run somewhat slower than normal because of the time spent collecting and writing the profile data.

The way you run the program—the arguments and input that you give it—may have a dramatic effect on what the profile information shows. The profile data will describe the parts of the program that were activated for the particular input you use. For example, if the first command you give to your program is to quit, the profile data will show the time used in initialization and in cleanup, but not much else.

Your program will write the profile data into a file called `'gmon.out'` just before exiting. If there is already a file called `'gmon.out'`, its contents are overwritten. There is currently no way to tell the program to write the profile data under a different name, but you can rename the file afterwards if you are concerned that it may be overwritten.

In order to write the `'gmon.out'` file properly, your program must exit normally: by returning from `main` or by calling `exit`. Calling the low-level function `_exit` does not write the profile data, and neither does abnormal termination due to an unhandled signal.

The `'gmon.out'` file is written in the program's *current working directory* at the time it exits. This means that if your program calls `chdir`, the `'gmon.out'` file will be left in the last directory your program `chdir`'d to. If you don't have permission to write in this directory, the file is not written, and you will get an error message.

Older versions of the GNU profiling library may also write a file called `'bb.out'`. This file, if present, contains an human-readable listing of the basic-block execution counts. Unfortunately, the appearance of a human-readable `'bb.out'` means the basic-block counts didn't get written into `'gmon.out'`. The Perl script `bbconv.pl`, included with the `gprof` source distribution, will convert a `'bb.out'` file into a format readable by `gprof`. Invoke it like this:

```
bbconv.pl < bb.out > bh-data
```

This translates the information in `'bb.out'` into a form that `gprof` can understand. But you still need to tell `gprof` about the existence of this translated information. To do that, include `bb-data` on the `gprof` command line, *along with* `'gmon.out'`, like this:

```
gprof options executable-file gmon.out bb-data [yet-more-profile-data-files...] [> out-  
file]
```


4 gprof Command Summary

After you have a profile data file ‘gmon.out’, you can run `gprof` to interpret the information in it. The `gprof` program prints a flat profile and a call graph on standard output. Typically you would redirect the output of `gprof` into a file with ‘>’.

You run `gprof` like this:

```
gprof options [executable-file [profile-data-files...]] [> outfile]
```

Here square-brackets indicate optional arguments.

If you omit the executable file name, the file ‘a.out’ is used. If you give no profile data file name, the file ‘gmon.out’ is used. If any file is not in the proper format, or if the profile data file does not appear to belong to the executable file, an error message is printed.

You can give more than one profile data file by entering all their names after the executable file name; then the statistics in all the data files are summed together.

The order of these options does not matter.

4.1 Output Options

These options specify which of several output formats `gprof` should produce.

Many of these options take an optional *symspec* to specify functions to be included or excluded. These options can be specified multiple times, with different *symspecs*, to include or exclude sets of symbols. See [Section 4.5 \[Symspecs\]](#), page 14.

Specifying any of these options overrides the default (‘-p -q’), which prints a flat profile and call graph analysis for all functions.

`-A[symspec]`

`--annotated-source[=symspec]`

The ‘-A’ option causes `gprof` to print annotated source code. If *symspec* is specified, print output only for matching symbols. See [Section 5.4 \[The Annotated Source Listing\]](#), page 24.

`-b`

`--brief` If the ‘-b’ option is given, `gprof` doesn’t print the verbose blurbs that try to explain the meaning of all of the fields in the tables. This is useful if you intend to print out the output, or are tired of seeing the blurbs.

`-C[symspec]`

`--exec-counts[=symspec]`

The ‘-C’ option causes `gprof` to print a tally of functions and the number of times each was called. If *symspec* is specified, print tally only for matching symbols.

If the profile data file contains basic-block count records, specifying the `-l` option, along with `-C`, will cause basic-block execution counts to be tallied and displayed.

`-i`

`--file-info`

The `-i` option causes `gprof` to display summary information about the profile data file(s) and then exit. The number of histogram, call graph, and basic-block count records is displayed.

`-I dirs`

`--directory-path=dirs`

The `-I` option specifies a list of search directories in which to find source files. Environment variable `GPROF_PATH` can also be used to convey this information. Used mostly for annotated source output.

`-J[symspec]`

`--no-annotated-source[=symspec]`

The `-J` option causes `gprof` not to print annotated source code. If `symspec` is specified, `gprof` prints annotated source, but excludes matching symbols.

`-L`

`--print-path`

Normally, source filenames are printed with the path component suppressed. The `-L` option causes `gprof` to print the full pathname of source filenames, which is determined from symbolic debugging information in the image file and is relative to the directory in which the compiler was invoked.

`-p[symspec]`

`--flat-profile[=symspec]`

The `-p` option causes `gprof` to print a flat profile. If `symspec` is specified, print flat profile only for matching symbols. See [Section 5.1 \[The Flat Profile\]](#), page 15.

`-P[symspec]`

`--no-flat-profile[=symspec]`

The `-P` option causes `gprof` to suppress printing a flat profile. If `symspec` is specified, `gprof` prints a flat profile, but excludes matching symbols.

`-q[symspec]`

`--graph[=symspec]`

The `-q` option causes `gprof` to print the call graph analysis. If `symspec` is specified, print call graph only for matching symbols and their children. See [Section 5.2 \[The Call Graph\]](#), page 17.

`-Q[symspec]`

`--no-graph[=symspec]`

The ‘-Q’ option causes `gprof` to suppress printing the call graph. If *symspec* is specified, `gprof` prints a call graph, but excludes matching symbols.

`-t`

`--table-length=num`

The ‘-t’ option causes the *num* most active source lines in each source file to be listed when source annotation is enabled. The default is 10.

`-y`

`--separate-files`

This option affects annotated source output only. Normally, `gprof` prints annotated source files to standard-output. If this option is specified, annotated source for a file named ‘*path/filename*’ is generated in the file ‘*filename-ann*’. If the underlying file system would truncate ‘*filename-ann*’ so that it overwrites the original ‘*filename*’, `gprof` generates annotated source in the file ‘*filename.ann*’ instead (if the original file name has an extension, that extension is *replaced* with ‘.ann’).

`-Z[symspec]`

`--no-exec-counts[=symspec]`

The ‘-Z’ option causes `gprof` not to print a tally of functions and the number of times each was called. If *symspec* is specified, print tally, but exclude matching symbols.

`-r`

`--function-ordering`

The ‘`--function-ordering`’ option causes `gprof` to print a suggested function ordering for the program based on profiling data. This option suggests an ordering which may improve paging, tlb and cache behavior for the program on systems which support arbitrary ordering of functions in an executable.

The exact details of how to force the linker to place functions in a particular order is system dependent and out of the scope of this manual.

`-R map_file`

`--file-ordering map_file`

The ‘`--file-ordering`’ option causes `gprof` to print a suggested .o link line ordering for the program based on profiling data. This option suggests an ordering which may improve paging, tlb and cache behavior for the program on systems which do not support arbitrary ordering of functions in an executable. Use of the ‘-a’ argument is highly recommended with this option.

The *map_file* argument is a pathname to a file which provides function name to object file mappings. The format of the file is similar to the output of the program *nm*.

```
c-parse.o:00000000 T yyparse
c-parse.o:00000004 C yyerrflag
c-lang.o:00000000 T maybe_objc_method_name
c-lang.o:00000000 T print_lang_statistics
c-lang.o:00000000 T recognize_objc_keyword
c-decl.o:00000000 T print_lang_identifier
c-decl.o:00000000 T print_lang_type
...
```

To create a *map_file* with GNU *nm*, type a command like *nm --extern-only --defined-only -v --print-file-name program-name*.

-T

--traditional

The ‘-T’ option causes *gprof* to print its output in “traditional” BSD style.

-w *width*

--width=*width*

Sets width of output lines to *width*. Currently only used when printing the function index at the bottom of the call graph.

-x

--all-lines

This option affects annotated source output only. By default, only the lines at the beginning of a basic-block are annotated. If this option is specified, every line in a basic-block is annotated by repeating the annotation for the first line. This behavior is similar to *tcov*’s ‘-a’.

--demangle[=*style*]

--no-demangle

These options control whether C++ symbol names should be demangled when printing output. The default is to demangle symbols. The **--no-demangle** option may be used to turn off demangling. Different compilers have different mangling styles. The optional demangling style argument can be used to choose an appropriate demangling style for your compiler.

4.2 Analysis Options

-a

--no-static

The ‘-a’ option causes *gprof* to suppress the printing of statically declared (private) functions. (These are functions whose

names are not listed as global, and which are not visible outside the file/function/block where they were defined.) Time spent in these functions, calls to/from them, etc., will all be attributed to the function that was loaded directly before it in the executable file. This option affects both the flat profile and the call graph.

`-c`

`--static-call-graph`

The `'-c'` option causes the call graph of the program to be augmented by a heuristic which examines the text space of the object file and identifies function calls in the binary machine code. Since normal call graph records are only generated when functions are entered, this option identifies children that could have been called, but never were. Calls to functions that were not compiled with profiling enabled are also identified, but only if symbol table entries are present for them. Calls to dynamic library routines are typically *not* found by this option. Parents or children identified via this heuristic are indicated in the call graph with call counts of `'0'`.

`-D`

`--ignore-non-functions`

The `'-D'` option causes `gprof` to ignore symbols which are not known to be functions. This option will give more accurate profile data on systems where it is supported (Solaris and HP-UX for example).

`-k from/to`

The `'-k'` option allows you to delete from the call graph any arcs from symbols matching *symspec from* to those matching *symspec to*.

`-l`

`--line`

The `'-l'` option enables line-by-line profiling, which causes histogram hits to be charged to individual source code lines, instead of functions. This feature only works with programs compiled by older versions of the `gcc` compiler. Newer versions of `gcc` are designed to work with the `gcov` tool instead.

If the program was compiled with basic-block counting enabled, this option will also identify how many times each line of code was executed. While line-by-line profiling can help isolate where in a large function a program is spending its time, it also significantly increases the running time of `gprof`, and magnifies statistical inaccuracies. See [Section 6.1 \[Statistical Sampling Error\]](#), page 27.

`-m num`

`--min-count=num`

This option affects execution count output only. Symbols that are executed less than *num* times are suppressed.

`-nsymspec`

`--time=symspec`

The ‘-n’ option causes `gprof`, in its call graph analysis, to only propagate times for symbols matching *symspec*.

`-Nsymspec`

`--no-time=symspec`

The ‘-n’ option causes `gprof`, in its call graph analysis, not to propagate times for symbols matching *symspec*.

`-Sfilename`

`--external-symbol-table=filename`

The ‘-S’ option causes `gprof` to read an external symbol table file, such as ‘`/proc/kallsyms`’, rather than read the symbol table from the given object file (the default is `a.out`). This is useful for profiling kernel modules.

`-z`

`--display-unused-functions`

If you give the ‘-z’ option, `gprof` will mention all functions in the flat profile, even those that were never called, and that had no time spent in them. This is useful in conjunction with the ‘-c’ option for discovering which routines were never called.

4.3 Miscellaneous Options

`-d [num]`

`--debug [= num]`

The ‘-d *num*’ option specifies debugging options. If *num* is not specified, enable all debugging. See [Section 9.4 \[Debugging gprof\], page 39](#).

`-h`

`--help` The ‘-h’ option prints command line usage.

`-Oname`

`--file-format=name`

Selects the format of the profile data files. Recognized formats are ‘auto’ (the default), ‘bsd’, ‘4.4bsd’, ‘magic’, and ‘prof’ (not yet supported).

`-s`

`--sum`

The ‘-s’ option causes `gprof` to summarize the information in the profile data files it read in, and write out a profile data file called ‘`gmon.sum`’, which contains all the information from the

profile data files that `gprof` read in. The file `'gmon.sum'` may be one of the specified input files; the effect of this is to merge the data in the other input files into `'gmon.sum'`.

Eventually you can run `gprof` again without `'-s'` to analyze the cumulative data in the file `'gmon.sum'`.

`-v`

`--version`

The `'-v'` flag causes `gprof` to print the current version number, and then exit.

4.4 Deprecated Options

These options have been replaced with newer versions that use `symspecs`.

`-e function_name`

The `'-e function'` option tells `gprof` to not print information about the function `function_name` (and its children...) in the call graph. The function will still be listed as a child of any functions that call it, but its index number will be shown as `'[not printed]'`. More than one `'-e'` option may be given; only one `function_name` may be indicated with each `'-e'` option.

`-E function_name`

The `-E function` option works like the `-e` option, but time spent in the function (and children who were not called from anywhere else), will not be used to compute the percentages-of-time for the call graph. More than one `'-E'` option may be given; only one `function_name` may be indicated with each `'-E'` option.

`-f function_name`

The `'-f function'` option causes `gprof` to limit the call graph to the function `function_name` and its children (and their children...). More than one `'-f'` option may be given; only one `function_name` may be indicated with each `'-f'` option.

`-F function_name`

The `'-F function'` option works like the `-f` option, but only time spent in the function and its children (and their children...) will be used to determine total-time and percentages-of-time for the call graph. More than one `'-F'` option may be given; only one `function_name` may be indicated with each `'-F'` option. The `'-F'` option overrides the `'-E'` option.

Note that only one function can be specified with each `-e`, `-E`, `-f` or `-F` option. To specify more than one function, use multiple options. For example, this command:

`gprof -e boring -f foo -f bar myprogram > gprof.output`
 lists in the call graph all functions that were reached from either `foo` or `bar` and were not reachable from `boring`.

4.5 Symspecs

Many of the output options allow functions to be included or excluded using *symspecs* (symbol specifications), which observe the following syntax:

```
filename_containing_a_dot
| funcname_not_containing_a_dot
| linenumber
| ( [ any_filename ] ':' ( any_funcname | linenumber ) )
```

Here are some sample symspecs:

`'main.c'` Selects everything in file `'main.c'`—the dot in the string tells `gprof` to interpret the string as a filename, rather than as a function name. To select a file whose name does not contain a dot, a trailing colon should be specified. For example, `'odd:'` is interpreted as the file named `'odd'`.

`'main'` Selects all functions named `'main'`.

Note that there may be multiple instances of the same function name because some of the definitions may be local (i.e., static). Unless a function name is unique in a program, you must use the colon notation explained below to specify a function from a specific source file.

Sometimes, function names contain dots. In such cases, it is necessary to add a leading colon to the name. For example, `':.mul'` selects function `'mul'`.

In some object file formats, symbols have a leading underscore. `gprof` will normally not print these underscores. When you name a symbol in a symspec, you should type it exactly as `gprof` prints it in its output. For example, if the compiler produces a symbol `'_main'` from your `main` function, `gprof` still prints it as `'main'` in its output, so you should use `'main'` in symspecs.

`'main.c:main'`

Selects function `'main'` in file `'main.c'`.

`'main.c:134'`

Selects line 134 in file `'main.c'`.

5 Interpreting gprof's Output

`gprof` can produce several different output styles, the most important of which are described below. The simplest output styles (file information, execution count, and function and file ordering) are not described here, but are documented with the respective options that trigger them. See [Section 4.1 \[Output Options\]](#), page 7.

5.1 The Flat Profile

The *flat profile* shows the total amount of time your program spent executing each function. Unless the `-z` option is given, functions with no apparent time spent in them, and no apparent calls to them, are not mentioned. Note that if a function was not compiled for profiling, and didn't run long enough to show up on the program counter histogram, it will be indistinguishable from a function that was never called.

This is part of a flat profile for a small program:

Flat profile:

Each sample counts as 0.01 seconds.

% time	cumulative seconds	self seconds	calls	self ms/call	total ms/call	name
33.34	0.02	0.02	7208	0.00	0.00	open
16.67	0.03	0.01	244	0.04	0.12	offtime
16.67	0.04	0.01	8	1.25	1.25	memcpy
16.67	0.05	0.01	7	1.43	1.43	write
16.67	0.06	0.01				mcount
0.00	0.06	0.00	236	0.00	0.00	tzset
0.00	0.06	0.00	192	0.00	0.00	tolower
0.00	0.06	0.00	47	0.00	0.00	strlen
0.00	0.06	0.00	45	0.00	0.00	strchr
0.00	0.06	0.00	1	0.00	50.00	main
0.00	0.06	0.00	1	0.00	0.00	memcpy
0.00	0.06	0.00	1	0.00	10.11	print
0.00	0.06	0.00	1	0.00	0.00	profil
0.00	0.06	0.00	1	0.00	50.00	report

...

The functions are sorted first by decreasing run-time spent in them, then by decreasing number of calls, then alphabetically by name. The functions `'mcount'` and `'profil'` are part of the profiling apparatus and appear in every flat profile; their time gives a measure of the amount of overhead due to profiling.

Just before the column headers, a statement appears indicating how much time each sample counted as. This *sampling period* estimates the margin of error in each of the time figures. A time figure that is not much larger than this is not reliable. In this example, each sample counted as 0.01 seconds, suggesting a 100 Hz sampling rate. The program's total execution time was 0.06 seconds, as indicated by the `'cumulative seconds'` field. Since each

sample counted for 0.01 seconds, this means only six samples were taken during the run. Two of the samples occurred while the program was in the ‘open’ function, as indicated by the ‘self seconds’ field. Each of the other four samples occurred one each in ‘offtime’, ‘memccpy’, ‘write’, and ‘mcount’. Since only six samples were taken, none of these values can be regarded as particularly reliable. In another run, the ‘self seconds’ field for ‘mcount’ might well be ‘0.00’ or ‘0.02’. See [Section 6.1 \[Statistical Sampling Error\]](#), page 27, for a complete discussion.

The remaining functions in the listing (those whose ‘self seconds’ field is ‘0.00’) didn’t appear in the histogram samples at all. However, the call graph indicated that they were called, so therefore they are listed, sorted in decreasing order by the ‘calls’ field. Clearly some time was spent executing these functions, but the paucity of histogram samples prevents any determination of how much time each took.

Here is what the fields in each line mean:

% time This is the percentage of the total execution time your program spent in this function. These should all add up to 100%.

cumulative seconds

This is the cumulative total number of seconds the computer spent executing this functions, plus the time spent in all the functions above this one in this table.

self seconds

This is the number of seconds accounted for by this function alone. The flat profile listing is sorted first by this number.

calls

This is the total number of times the function was called. If the function was never called, or the number of times it was called cannot be determined (probably because the function was not compiled with profiling enabled), the *calls* field is blank.

self ms/call

This represents the average number of milliseconds spent in this function per call, if this function is profiled. Otherwise, this field is blank for this function.

total ms/call

This represents the average number of milliseconds spent in this function and its descendants per call, if this function is profiled. Otherwise, this field is blank for this function. This is the only field in the flat profile that uses call graph analysis.

name

This is the name of the function. The flat profile is sorted by this field alphabetically after the *self seconds* and *calls* fields are sorted.

5.2 The Call Graph

The *call graph* shows how much time was spent in each function and its children. From this information, you can find functions that, while they themselves may not have used much time, called other functions that did use unusual amounts of time.

Here is a sample call from a small program. This call came from the same gprof run as the flat profile example in the previous section.

```
granularity: each sample hit covers 2 byte(s) for 20.00% of 0.05 seconds
```

index	% time	self	children	called	name
[1]	100.0	0.00	0.05		<spontaneous>
		0.00	0.05	1/1	start [1]
		0.00	0.00	1/2	main [2]
		0.00	0.00	1/1	on_exit [28]
		0.00	0.00	1/1	exit [59]

[2]	100.0	0.00	0.05	1/1	start [1]
		0.00	0.05	1	main [2]
		0.00	0.05	1/1	report [3]

[3]	100.0	0.00	0.05	1/1	main [2]
		0.00	0.05	1	report [3]
		0.00	0.03	8/8	timelocal [6]
		0.00	0.01	1/1	print [9]
		0.00	0.01	9/9	fgets [12]
		0.00	0.00	12/34	strncmp <cycle 1> [40]
		0.00	0.00	8/8	lookup [20]
		0.00	0.00	1/1	fopen [21]
		0.00	0.00	8/8	chewtime [24]
		0.00	0.00	8/16	skipSPACE [44]

[4]	59.8	0.01	0.02	8+472	<cycle 2 as a whole> [4]
		0.01	0.02	244+260	offtime <cycle 2> [7]
		0.00	0.00	236+1	tzset <cycle 2> [26]

The lines full of dashes divide this table into *entries*, one for each function. Each entry has one or more lines.

In each entry, the primary line is the one that starts with an index number in square brackets. The end of this line says which function the entry is for. The preceding lines in the entry describe the callers of this function and the following lines describe its subroutines (also called *children* when we speak of the call graph).

The entries are sorted by time spent in the function and its subroutines.

The internal profiling function `mcount` (see [Section 5.1 \[The Flat Profile\]](#), [page 15](#)) is never mentioned in the call graph.

5.2.1 The Primary Line

The *primary line* in a call graph entry is the line that describes the function which the entry is about and gives the overall statistics for this function.

For reference, we repeat the primary line from the entry for function `report` in our main example, together with the heading line that shows the names of the fields:

```

    index % time    self  children called      name
    ...
    [3]   100.0    0.00   0.05      1      report [3]

```

Here is what the fields in the primary line mean:

<code>index</code>	<p>Entries are numbered with consecutive integers. Each function therefore has an index number, which appears at the beginning of its primary line.</p> <p>Each cross-reference to a function, as a caller or subroutine of another, gives its index number as well as its name. The index number guides you if you wish to look for the entry for that function.</p>
<code>% time</code>	<p>This is the percentage of the total time that was spent in this function, including time spent in subroutines called from this function.</p> <p>The time spent in this function is counted again for the callers of this function. Therefore, adding up these percentages is meaningless.</p>
<code>self</code>	<p>This is the total amount of time spent in this function. This should be identical to the number printed in the <code>seconds</code> field for this function in the flat profile.</p>
<code>children</code>	<p>This is the total amount of time spent in the subroutine calls made by this function. This should be equal to the sum of all the <code>self</code> and <code>children</code> entries of the children listed directly below this function.</p>
<code>called</code>	<p>This is the number of times the function was called.</p> <p>If the function called itself recursively, there are two numbers, separated by a '+'. The first number counts non-recursive calls, and the second counts recursive calls.</p> <p>In the example above, the function <code>report</code> was called once from <code>main</code>.</p>
<code>name</code>	<p>This is the name of the current function. The index number is repeated after it.</p> <p>If the function is part of a cycle of recursion, the cycle number is printed between the function's name and the index number (see Section 5.2.4 [How Mutually Recursive Functions Are Described], page 20). For example, if function <code>gnurr</code> is part of</p>

cycle number one, and has index number twelve, its primary line would be end like this:

```
gnurr <cycle 1> [12]
```

5.2.2 Lines for a Function's Callers

A function's entry has a line for each function it was called by. These lines' fields correspond to the fields of the primary line, but their meanings are different because of the difference in context.

For reference, we repeat two lines from the entry for the function `report`, the primary line and one caller-line preceding it, together with the heading line that shows the names of the fields:

```
index % time    self children called    name
...
          0.00   0.05     1/1         main [2]
[3]    100.0   0.00   0.05     1         report [3]
```

Here are the meanings of the fields in the caller-line for `report` called from `main`:

self An estimate of the amount of time spent in `report` itself when it was called from `main`.

children An estimate of the amount of time spent in subroutines of `report` when `report` was called from `main`.

The sum of the **self** and **children** fields is an estimate of the amount of time spent within calls to `report` from `main`.

called Two numbers: the number of times `report` was called from `main`, followed by the total number of non-recursive calls to `report` from all its callers.

name and index number

The name of the caller of `report` to which this line applies, followed by the caller's index number.

Not all functions have entries in the call graph; some options to `gprof` request the omission of certain functions. When a caller has no entry of its own, it still has caller-lines in the entries of the functions it calls.

If the caller is part of a recursion cycle, the cycle number is printed between the name and the index number.

If the identity of the callers of a function cannot be determined, a dummy caller-line is printed which has '`<spontaneous>`' as the "caller's name" and all other fields blank. This can happen for signal handlers.

5.2.3 Lines for a Function's Subroutines

A function's entry has a line for each of its subroutines—in other words, a line for each other function that it called. These lines' fields correspond to

the fields of the primary line, but their meanings are different because of the difference in context.

For reference, we repeat two lines from the entry for the function `main`, the primary line and a line for a subroutine, together with the heading line that shows the names of the fields:

```

index % time    self children called    name
...
[2]   100.0    0.00   0.05      1      main [2]
                0.00   0.05      1/1      report [3]
```

Here are the meanings of the fields in the subroutine-line for `main` calling `report`:

self An estimate of the amount of time spent directly within `report` when `report` was called from `main`.

children An estimate of the amount of time spent in subroutines of `report` when `report` was called from `main`.
The sum of the `self` and `children` fields is an estimate of the total time spent in calls to `report` from `main`.

called Two numbers, the number of calls to `report` from `main` followed by the total number of non-recursive calls to `report`. This ratio is used to determine how much of `report`'s `self` and `children` time gets credited to `main`. See [Section 6.2 \[Estimating children Times\]](#), page 28.

name The name of the subroutine of `main` to which this line applies, followed by the subroutine's index number.
If the caller is part of a recursion cycle, the cycle number is printed between the name and the index number.

5.2.4 How Mutually Recursive Functions Are Described

The graph may be complicated by the presence of *cycles of recursion* in the call graph. A cycle exists if a function calls another function that (directly or indirectly) calls (or appears to call) the original function. For example: if `a` calls `b`, and `b` calls `a`, then `a` and `b` form a cycle.

Whenever there are call paths both ways between a pair of functions, they belong to the same cycle. If `a` and `b` call each other and `b` and `c` call each other, all three make one cycle. Note that even if `b` only calls `a` if it was not called from `a`, `gprof` cannot determine this, so `a` and `b` are still considered a cycle.

The cycles are numbered with consecutive integers. When a function belongs to a cycle, each time the function name appears in the call graph it is followed by '`<cycle number>`'.

The reason cycles matter is that they make the time values in the call graph paradoxical. The "time spent in children" of `a` should include the time

spent in its subroutine **b** and in **b**'s subroutines—but one of **b**'s subroutines is **a**! How much of **a**'s time should be included in the children of **a**, when **a** is indirectly recursive?

The way **gprof** resolves this paradox is by creating a single entry for the cycle as a whole. The primary line of this entry describes the total time spent directly in the functions of the cycle. The “subroutines” of the cycle are the individual functions of the cycle, and all other functions that were called directly by them. The “callers” of the cycle are the functions, outside the cycle, that called functions in the cycle.

Here is an example portion of a call graph which shows a cycle containing functions **a** and **b**. The cycle was entered by a call to **a** from **main**; both **a** and **b** called **c**.

index	% time	self	children	called	name
		1.77	0	1/1	main [2]
[3]	91.71	1.77	0	1+5	<cycle 1 as a whole> [3]
		1.02	0	3	b <cycle 1> [4]
		0.75	0	2	a <cycle 1> [5]
				3	a <cycle 1> [5]
[4]	52.85	1.02	0	0	b <cycle 1> [4]
				2	a <cycle 1> [5]
		0	0	3/6	c [6]
		1.77	0	1/1	main [2]
				2	b <cycle 1> [4]
[5]	38.86	0.75	0	1	a <cycle 1> [5]
				3	b <cycle 1> [4]
		0	0	3/6	c [6]

(The entire call graph for this program contains in addition an entry for **main**, which calls **a**, and an entry for **c**, with callers **a** and **b**.)

index	% time	self	children	called	name
					<spontaneous>
[1]	100.00	0	1.93	0	start [1]
		0.16	1.77	1/1	main [2]
		0.16	1.77	1/1	start [1]
[2]	100.00	0.16	1.77	1	main [2]
		1.77	0	1/1	a <cycle 1> [5]
		1.77	0	1/1	main [2]
[3]	91.71	1.77	0	1+5	<cycle 1 as a whole> [3]
		1.02	0	3	b <cycle 1> [4]
		0.75	0	2	a <cycle 1> [5]
		0	0	6/6	c [6]
				3	a <cycle 1> [5]
[4]	52.85	1.02	0	0	b <cycle 1> [4]

				2	a <cycle 1> [5]
		0	0	3/6	c [6]

		1.77	0	1/1	main [2]
				2	b <cycle 1> [4]
[5]	38.86	0.75	0	1	a <cycle 1> [5]
				3	b <cycle 1> [4]
		0	0	3/6	c [6]

		0	0	3/6	b <cycle 1> [4]
		0	0	3/6	a <cycle 1> [5]
[6]	0.00	0	0	6	c [6]

The **self** field of the cycle's primary line is the total time spent in all the functions of the cycle. It equals the sum of the **self** fields for the individual functions in the cycle, found in the entry in the subroutine lines for these functions.

The **children** fields of the cycle's primary line and subroutine lines count only subroutines outside the cycle. Even though **a** calls **b**, the time spent in those calls to **b** is not counted in **a**'s **children** time. Thus, we do not encounter the problem of what to do when the time in those calls to **b** includes indirect recursive calls back to **a**.

The **children** field of a caller-line in the cycle's entry estimates the amount of time spent *in the whole cycle*, and its other subroutines, on the times when that caller called a function in the cycle.

The **called** field in the primary line for the cycle has two numbers: first, the number of times functions in the cycle were called by functions outside the cycle; second, the number of times they were called by functions in the cycle (including times when a function in the cycle calls itself). This is a generalization of the usual split into non-recursive and recursive calls.

The **called** field of a subroutine-line for a cycle member in the cycle's entry says how many time that function was called from functions in the cycle. The total of all these is the second number in the primary line's **called** field.

In the individual entry for a function in a cycle, the other functions in the same cycle can appear as subroutines and as callers. These lines show how many times each function in the cycle called or was called from each other function in the cycle. The **self** and **children** fields in these lines are blank because of the difficulty of defining meanings for them when recursion is going on.

5.3 Line-by-line Profiling

gprof's '-1' option causes the program to perform *line-by-line* profiling. In this mode, histogram samples are assigned not to functions, but to individual lines of source code. This only works with programs compiled with older

versions of the `gcc` compiler. Newer versions of `gcc` use a different program - `gcov` - to display line-by-line profiling information.

With the older versions of `gcc` the program usually has to be compiled with a `-g` option, in addition to `-pg`, in order to generate debugging symbols for tracking source code lines. Note, in much older versions of `gcc` the program had to be compiled with the `-a` command line option as well.

The flat profile is the most useful output table in line-by-line mode. The call graph isn't as useful as normal, since the current version of `gprof` does not propagate call graph arcs from source code lines to the enclosing function. The call graph does, however, show each line of code that called each function, along with a count.

Here is a section of `gprof`'s output, without line-by-line profiling. Note that `ct_init` accounted for four histogram hits, and 13327 calls to `init_block`.

Flat profile:

Each sample counts as 0.01 seconds.

% time	cumulative seconds	self seconds	calls	self us/call	total us/call	name
30.77	0.13	0.04	6335	6.31	6.31	ct_init

Call graph (explanation follows)

granularity: each sample hit covers 4 byte(s) for 7.69% of 0.13 seconds

index	% time	self	children	called	name
		0.00	0.00	1/13496	name_too_long
		0.00	0.00	40/13496	deflate
		0.00	0.00	128/13496	deflate_fast
		0.00	0.00	13327/13496	ct_init
[7]	0.0	0.00	0.00	13496	init_block

Now let's look at some of `gprof`'s output from the same program run, this time with line-by-line profiling enabled. Note that `ct_init`'s four histogram hits are broken down into four lines of source code—one hit occurred on each of lines 349, 351, 382 and 385. In the call graph, note how `ct_init`'s 13327 calls to `init_block` are broken down into one call from line 396, 3071 calls from line 384, 3730 calls from line 385, and 6525 calls from 387.

Flat profile:

Each sample counts as 0.01 seconds.

% time	cumulative seconds	self seconds	calls	name
7.69	0.10	0.01		ct_init (trees.c:349)
7.69	0.11	0.01		ct_init (trees.c:351)

```

7.69      0.12      0.01      ct_init (trees.c:382)
7.69      0.13      0.01      ct_init (trees.c:385)

```

Call graph (explanation follows)

granularity: each sample hit covers 4 byte(s) for 7.69% of 0.13 seconds

% time	self	children	called	name
	0.00	0.00	1/13496	name_too_long (gzip.c:1440)
	0.00	0.00	1/13496	deflate (deflate.c:763)
	0.00	0.00	1/13496	ct_init (trees.c:396)
	0.00	0.00	2/13496	deflate (deflate.c:727)
	0.00	0.00	4/13496	deflate (deflate.c:686)
	0.00	0.00	5/13496	deflate (deflate.c:675)
	0.00	0.00	12/13496	deflate (deflate.c:679)
	0.00	0.00	16/13496	deflate (deflate.c:730)
	0.00	0.00	128/13496	deflate_fast (deflate.c:654)
	0.00	0.00	3071/13496	ct_init (trees.c:384)
	0.00	0.00	3730/13496	ct_init (trees.c:385)
	0.00	0.00	6525/13496	ct_init (trees.c:387)
[6] 0.0	0.00	0.00	13496	init_block (trees.c:408)

5.4 The Annotated Source Listing

`gprof`'s `'-A'` option triggers an annotated source listing, which lists the program's source code, each function labeled with the number of times it was called. You may also need to specify the `'-I'` option, if `gprof` can't find the source code files.

With older versions of `gcc` compiling with `'gcc ... -g -pg -a'` augments your program with basic-block counting code, in addition to function counting code. This enables `gprof` to determine how many times each line of code was executed. With newer versions of `gcc` support for displaying basic-block counts is provided by the `gcov` program.

For example, consider the following function, taken from `gzip`, with line numbers added:

```

1 ulg updcrc(s, n)
2     uch *s;
3     unsigned n;
4 {
5     register ulg c;
6
7     static ulg crc = (ulg)0xffffffffL;
8
9     if (s == NULL) {
10        c = 0xffffffffL;
11    } else {

```

```

12         c = crc;
13         if (n) do {
14             c = crc_32_tab[...];
15         } while (--n);
16     }
17     crc = c;
18     return c ^ 0xffffffffL;
19 }

```

updcrc has at least five basic-blocks. One is the function itself. The `if` statement on line 9 generates two more basic-blocks, one for each branch of the `if`. A fourth basic-block results from the `if` on line 13, and the contents of the `do` loop form the fifth basic-block. The compiler may also generate additional basic-blocks to handle various special cases.

A program augmented for basic-block counting can be analyzed with `'gprof -l -A'`. The `'-x'` option is also helpful, to ensure that each line of code is labeled at least once. Here is `updcrc`'s annotated source listing for a sample `gzip` run:

```

        ulg updcrc(s, n)
        uch *s;
        unsigned n;
2 ->{
        register ulg c;

        static ulg crc = (ulg)0xffffffffL;

2 ->    if (s == NULL) {
1 ->        c = 0xffffffffL;
1 ->    } else {
1 ->        c = crc;
1 ->        if (n) do {
26312 ->            c = crc_32_tab[...];
26312,1,26311 ->        } while (--n);
        }
2 ->    crc = c;
2 ->    return c ^ 0xffffffffL;
2 ->}

```

In this example, the function was called twice, passing once through each branch of the `if` statement. The body of the `do` loop was executed a total of 26312 times. Note how the `while` statement is annotated. It began execution 26312 times, once for each iteration through the loop. One of those times (the last time) it exited, while it branched back to the beginning of the loop 26311 times.

6 Inaccuracy of gprof Output

6.1 Statistical Sampling Error

The run-time figures that `gprof` gives you are based on a sampling process, so they are subject to statistical inaccuracy. If a function runs only a small amount of time, so that on the average the sampling process ought to catch that function in the act only once, there is a pretty good chance it will actually find that function zero times, or twice.

By contrast, the number-of-calls and basic-block figures are derived by counting, not sampling. They are completely accurate and will not vary from run to run if your program is deterministic and single threaded. In multi-threaded applications, or single threaded applications that link with multi-threaded libraries, the counts are only deterministic if the counting function is thread-safe. (Note: beware that the `mcount` counting function in `glibc` is *not* thread-safe). See [Section 9.1 \[Implementation of Profiling\]](#), [page 33](#).

The *sampling period* that is printed at the beginning of the flat profile says how often samples are taken. The rule of thumb is that a run-time figure is accurate if it is considerably bigger than the sampling period.

The actual amount of error can be predicted. For n samples, the *expected* error is the square-root of n . For example, if the sampling period is 0.01 seconds and `foo`'s run-time is 1 second, n is 100 samples (1 second/0.01 seconds), $\text{sqrt}(n)$ is 10 samples, so the expected error in `foo`'s run-time is 0.1 seconds (10*0.01 seconds), or ten percent of the observed value. Again, if the sampling period is 0.01 seconds and `bar`'s run-time is 100 seconds, n is 10000 samples, $\text{sqrt}(n)$ is 100 samples, so the expected error in `bar`'s run-time is 1 second, or one percent of the observed value. It is likely to vary this much *on the average* from one profiling run to the next. (*Sometimes* it will vary more.)

This does not mean that a small run-time figure is devoid of information. If the program's *total* run-time is large, a small run-time for one function does tell you that that function used an insignificant fraction of the whole program's time. Usually this means it is not worth optimizing.

One way to get more accuracy is to give your program more (but similar) input data so it will take longer. Another way is to combine the data from several runs, using the `-s` option of `gprof`. Here is how:

1. Run your program once.
2. Issue the command `'mv gmon.out gmon.sum'`.
3. Run your program again, the same as before.
4. Merge the new data in `'gmon.out'` into `'gmon.sum'` with this command:

```
gprof -s executable-file gmon.out gmon.sum
```
5. Repeat the last two steps as often as you wish.

6. Analyze the cumulative data using this command:

```
gprof executable-file gmon.sum > output-file
```

6.2 Estimating children Times

Some of the figures in the call graph are estimates—for example, the `children` time values and all the time figures in caller and subroutine lines.

There is no direct information about these measurements in the profile data itself. Instead, `gprof` estimates them by making an assumption about your program that might or might not be true.

The assumption made is that the average time spent in each call to any function `foo` is not correlated with who called `foo`. If `foo` used 5 seconds in all, and 2/5 of the calls to `foo` came from `a`, then `foo` contributes 2 seconds to `a`'s `children` time, by assumption.

This assumption is usually true enough, but for some programs it is far from true. Suppose that `foo` returns very quickly when its argument is zero; suppose that `a` always passes zero as an argument, while other callers of `foo` pass other arguments. In this program, all the time spent in `foo` is in the calls from callers other than `a`. But `gprof` has no way of knowing this; it will blindly and incorrectly charge 2 seconds of time in `foo` to the children of `a`.

We hope some day to put more complete data into '`gmon.out`', so that this assumption is no longer needed, if we can figure out how. For the novice, the estimated figures are usually more useful than misleading.

