

# Debugging with GDB

---

The GNU Source-Level Debugger

Fifth Edition, for GDB version 4.17.gnat.3.11p  
April 1998

**Richard M. Stallman and Roland H. Pesch**

---

(Send bugs and comments on GDB to [bug-gdb@prep.ai.mit.edu](mailto:bug-gdb@prep.ai.mit.edu).)

*Debugging with GDB*

TeXinfo 2.196

[doc@cygnus.com](mailto:doc@cygnus.com)

Copyright © 1988, 1989, 1990, 1991, 1992, 1993, 1994, 1995, 1996, 1997, 1998 Free Software Foundation, Inc.

Published by the Free Software Foundation

59 Temple Place - Suite 330,

Boston, MA 02111-1307 USA

Printed copies are available for \$20 each.

ISBN 1-882114-11-6

Permission is granted to make and distribute verbatim copies of this manual provided the copyright notice and this permission notice are preserved on all copies.

Permission is granted to copy and distribute modified versions of this manual under the conditions for verbatim copying, provided also that the entire resulting derived work is distributed under the terms of a permission notice identical to this one.

Permission is granted to copy and distribute translations of this manual into another language, under the above conditions for modified versions.

## Summary of GDB

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

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

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

You can use GDB to debug programs written in C or C++. For more information, see [Section 8.3.1 \[C and C++\], page 59](#).

Support for Modula-2 and Chill is partial. For information on Modula-2, see [\(undefined\) \[Modula-2\], page \(undefined\)](#). There is no further documentation on Chill yet.

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

## Free software

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

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

## Contributors to GDB

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

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

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

So that they may not regard their long labor as thankless, we particularly thank those who shepherded GDB through major releases: Stan Shebs (release 4.14), Fred Fish (releases

4.13, 4.12, 4.11, 4.10, and 4.9), Stu Grossman and John Gilmore (releases 4.8, 4.7, 4.6, 4.5, and 4.4), John Gilmore (releases 4.3, 4.2, 4.1, 4.0, and 3.9); Jim Kingdon (releases 3.5, 3.4, and 3.3); and Randy Smith (releases 3.2, 3.1, and 3.0). As major maintainer of GDB for some period, each contributed significantly to the structure, stability, and capabilities of the entire debugger.

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

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

GDB 4 uses the BFD subroutine library to examine multiple object-file formats; BFD was a joint project of David V. Henkel-Wallace, Rich Pixley, Steve Chamberlain, and John Gilmore.

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

Brent Benson of Harris Computer Systems contributed DWARF 2 support.

Adam de Boor and Bradley Davis contributed the ISI Optimum V support. Per Bothner, Noboyuki Hikichi, and Alessandro Forin contributed MIPS support. Jean-Daniel Fekete contributed Sun 386i support. Chris Hanson improved the HP9000 support. Noboyuki Hikichi and Tomoyuki Hasei contributed Sony/News OS 3 support. David Johnson contributed Encore Umax support. Jyrki Kuoppala contributed Altos 3068 support. Jeff Law contributed HP PA and SOM support. Keith Packard contributed NS32K support. Doug Rabson contributed Acorn Risc Machine support. Bob Rusk contributed Harris Nighthawk CX-UX support. Chris Smith contributed Convex support (and Fortran debugging). Jonathan Stone contributed Pyramid support. Michael Tiemann contributed SPARC support. Tim Tucker contributed support for the Gould NP1 and Gould Povernode. Pace Willison contributed Intel 386 support. Jay Vosburgh contributed Symmetry support.

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

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

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

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

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

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

Hitachi America, Ltd. sponsored the support for Hitachi microprocessors.

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

Michael Snyder added support for tracepoints.

Stu Grossman wrote gdbserver.

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

Cygnus Solutions has sponsored GDB maintenance and much of its development since 1991.



# 1 Getting In and Out of GDB

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

- type ‘`gdb`’ to start GDB.
- type `quit` or `C-d` to exit.

## 1.1 Invoking GDB

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

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

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

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

```
gdb program
```

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

```
gdb -silent
```

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

Type

```
gdb -help
```

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

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

### 1.1.1 Choosing files

When GDB starts, it reads any argument other than options as specifying an executable file. This is the same as if the argument was specified by the ‘`-se`’ option.

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

```
-symbols file
```

```
-s file      Read symbol table from file file.
```

```
-exec file
```

```
-e file      Use file file as the executable file to execute when appropriate.
```

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

- `-command file`
- `-x file` Execute GDB commands from file *file*. See [Section 14.3 \[Command files\]](#), page 102.
- `-directory directory`
- `-d directory`  
Add *directory* to the path to search for source files.
- `-r`
- `-readnow` Read each symbol file's entire symbol table immediately, rather than the default, which is to read it incrementally as it is needed. This makes startup slower, but makes future operations faster.

### 1.1.2 Choosing modes

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

- `-nx`
- `-n` Do not execute commands from any initialization files (normally called '`.gdbinit`'). Normally, the commands in these files are executed after all the command options and arguments have been processed. See [Section 14.3 \[Command files\]](#), page 102.
- `-quiet`
- `-q` “Quiet”. Do not print the introductory and copyright messages. These messages are also suppressed in batch mode.
- `-batch` Run in batch mode. Exit with status 0 after processing all the command files specified with '`-x`' (and all commands from initialization files, if not inhibited with '`-n`'). Exit with nonzero status if an error occurs in executing the GDB commands in the command files.  
  
Batch mode may be useful for running GDB as a filter, for example to download and run a program on another computer; in order to make this more useful, the message  

```
Program exited normally.
```

(which is ordinarily issued whenever a program running under GDB control terminates) is not issued when running in batch mode.
- `-cd directory`  
Run GDB using *directory* as its working directory, instead of the current directory.
- `-fullname`
- `-f` GNU Emacs sets this option when it runs GDB as a subprocess. It tells GDB to output the full file name and line number in a standard, recognizable fashion each time a stack frame is displayed (which includes each time your program stops). This recognizable format looks like two '`\032`' characters, followed by the file name, line number and character position separated by colons, and a newline. The Emacs-to-GDB interface program uses the two '`\032`' characters as a signal to display the source code for the frame.



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

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

## 1.2 Quitting GDB

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

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

## 1.3 Shell commands

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

**shell** *command string*

Invoke a the standard shell to execute *command string*. If it exists, the environment variable **SHELL** determines which shell to run. Otherwise GDB uses **/bin/sh**.

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

**make** *make-args*

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



## 2 GDB Commands

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

### 2.1 Command syntax

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

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

A blank line as input to GDB (typing just `RET`) means to repeat the previous command. Certain commands (for example, `run`) will not repeat this way; these are commands whose unintentional repetition might cause trouble and which you are unlikely to want to repeat.

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

GDB can also use `RET` in another way: to partition lengthy output, in a way similar to the common utility `more` (see [Section 13.4 \[Screen size\], page 99](#)). Since it is easy to press one `RET` too many in this situation, GDB disables command repetition after any command that generates this sort of display.

Any text from a `#` to the end of the line is a comment; it does nothing. This is useful mainly in command files (see [Section 14.3 \[Command files\], page 102](#)).

### 2.2 Command completion

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

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

```
( gdb ) info bre TAB
```

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

```
( gdb) info breakpoints
```

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

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

```
( gdb) b make_ (TAB)
GDB sounds bell; press (TAB) again, to see:
make_a_section_from_file      make_environ
make_abs_section              make_function_type
make_blockvector              make_pointer_type
make_cleanup                   make_reference_type
make_command                   make_symbol_completion_list
( gdb) b make_
```

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

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

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

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

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

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

```
( gdb) b bub TAB
```

GDB alters your input line to the following, and rings a bell:

```
( gdb) b 'bubble(
```

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

## 2.3 Getting help

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

```
help
```

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

```
( gdb) help
```

```
List of classes of commands:
```

```
running -- Running the program
```

```
stack -- Examining the stack
```

```
data -- Examining data
```

```
breakpoints -- Making program stop at certain points
```

```
files -- Specifying and examining files
```

```
status -- Status inquiries
```

```
support -- Support facilities
```

```
user-defined -- User-defined commands
```

```
aliases -- Aliases of other commands
```

```
obscure -- Obscure features
```

```
Type "help" followed by a class name for a list of
commands in that class.
```

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

```
Command name abbreviations are allowed if unambiguous.
```

```
( gdb)
```

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

```
( gdb) help status
```

```
Status inquiries.
```

```
List of commands:
```

```
show -- Generic command for showing things set
with "set"
```

```
info -- Generic command for printing status
```

```
Type "help" followed by command name for full
```

```
documentation.
Command name abbreviations are allowed if unambiguous.
( gdb)
```

**help *command***

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

**complete *args***

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

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

This is intended for use by GNU Emacs.

In addition to **help**, you can use the GDB commands **info** and **show** to inquire about the state of your program, or the state of GDB itself. Each command supports many topics of inquiry; this manual introduces each of them in the appropriate context. The listings under **info** and under **show** in the Index point to all the sub-commands. See [\[Index\]](#), page 129.

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

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

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

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

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

**show version**

Show what version of GDB is running. You should include this information in GDB bug-reports. If multiple versions of GDB are in use at your site, you may occasionally want to determine which version of GDB you are running; as GDB

evolves, new commands are introduced, and old ones may wither away. The version number is also announced when you start GDB.

`show copying`

Display information about permission for copying GDB.

`show warranty`

Display the GNU “NO WARRANTY” statement.





## 3 Running Programs Under GDB

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

### 3.1 Compiling for debugging

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

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

Many C compilers are unable to handle the `-g` and `-O` options together. Using those compilers, you cannot generate optimized executables containing debugging information.

GCC, the GNU C compiler, supports `-g` with or without `-O`, making it possible to debug optimized code. We recommend that you *always* use `-g` whenever you compile a program. You may think your program is correct, but there is no sense in pushing your luck.

When you debug a program compiled with `-g -O`, remember that the optimizer is rearranging your code; the debugger shows you what is really there. Do not be too surprised when the execution path does not exactly match your source file! An extreme example: if you define a variable, but never use it, GDB never sees that variable—because the compiler optimizes it out of existence.

Some things do not work as well with `-g -O` as with just `-g`, particularly on machines with instruction scheduling. If in doubt, recompile with `-g` alone, and if this fixes the problem, please report it to us as a bug (including a test case!).

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

### 3.2 Starting your program

**run**

**r** Use the `run` command to start your program under GDB. You must first specify the program name (except on VxWorks) with an argument to GDB (see [Chapter 1 \[Getting In and Out of GDB\], page 5](#)), or by using the `file` or `exec-file` command (see [Section 11.1 \[Commands to specify files\], page 79](#)).

When you issue the `run` command, your program begins to execute immediately. See [Chapter 4 \[Stopping and continuing\], page 17](#), for discussion of how to arrange for your program to stop. Once your program has stopped, you may call functions in your program, using the `print` or `call` commands. See [Chapter 7 \[Examining Data\], page 41](#).

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

### 3.3 Debugging programs with multiple processes

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

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

## 4 Stopping and Continuing

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

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

`info program`

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

### 4.1 Breakpoints, watchpoints, and exceptions

A *breakpoint* makes your program stop whenever a certain point in the program is reached. For each breakpoint, you can add conditions to control in finer detail whether your program stops. You can set breakpoints with the `break` command and its variants (see [Section 4.1.1 \[Setting breakpoints\], page 17](#)), to specify the place where your program should stop by line number, function name or exact address in the program. In languages with exception handling (such as GNU C++), you can also set breakpoints where an exception is raised (see [Section 4.1.3 \[Breakpoints and exceptions\], page 21](#)).

In SunOS 4.x, SVR4, and Alpha OSF/1 configurations, you can now set breakpoints in shared libraries before the executable is run.

A *watchpoint* is a special breakpoint that stops your program when the value of an expression changes. You must use a different command to set watchpoints (see [Section 4.1.2 \[Setting watchpoints\], page 21](#)), but aside from that, you can manage a watchpoint like any other breakpoint: you enable, disable, and delete both breakpoints and watchpoints using the same commands.

You can arrange to have values from your program displayed automatically whenever GDB stops at a breakpoint. See [Section 7.6 \[Automatic display\], page 46](#).

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

#### 4.1.1 Setting breakpoints

Breakpoints are set with the `break` command (abbreviated `b`). The debugger convenience variable `‘$bpnum’` records the number of the breakpoints you’ve set most recently; see [Section 7.9 \[Convenience variables\], page 53](#), for a discussion of what you can do with convenience variables.

You have several ways to say where the breakpoint should go.

**break** *function*

Set a breakpoint at entry to function *function*. When using source languages that permit overloading of symbols, such as C++, *function* may refer to more than one possible place to break. See [Section 4.1.8 \[Breakpoint menus\]](#), page 26, for a discussion of that situation.

**break** *+offset***break** *-offset*

Set a breakpoint some number of lines forward or back from the position at which execution stopped in the currently selected frame.

**break** *linenum*

Set a breakpoint at line *linenum* in the current source file. That file is the last file whose source text was printed. This breakpoint stops your program just before it executes any of the code on that line.

**break** *filename:linenum*

Set a breakpoint at line *linenum* in source file *filename*.

**break** *filename:function*

Set a breakpoint at entry to function *function* found in file *filename*. Specifying a file name as well as a function name is superfluous except when multiple files contain similarly named functions.

**break** *\*address*

Set a breakpoint at address *address*. You can use this to set breakpoints in parts of your program which do not have debugging information or source files.

**break**

When called without any arguments, **break** sets a breakpoint at the next instruction to be executed in the selected stack frame (see [Chapter 5 \[Examining the Stack\]](#), page 33). In any selected frame but the innermost, this makes your program stop as soon as control returns to that frame. This is similar to the effect of a **finish** command in the frame inside the selected frame—except that **finish** does not leave an active breakpoint. If you use **break** without an argument in the innermost frame, GDB stops the next time it reaches the current location; this may be useful inside loops.

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

**break** ... *if cond*

Set a breakpoint with condition *cond*; evaluate the expression *cond* each time the breakpoint is reached, and stop only if the value is nonzero—that is, if *cond* evaluates as true. ‘...’ stands for one of the possible arguments described above (or no argument) specifying where to break. See [Section 4.1.6 \[Break conditions\]](#), page 24, for more information on breakpoint conditions.

**tbreak** *args*

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

automatically deleted after the first time your program stops there. See [Section 4.1.5 \[Disabling breakpoints\], page 23](#).

#### `hbreak args`

Set a hardware-assisted breakpoint. *args* are the same as for the `break` command and the breakpoint is set in the same way, but the breakpoint requires hardware support and some target hardware may not have this support. The main purpose of this is EPROM/ROM code debugging, so you can set a breakpoint at an instruction without changing the instruction. This can be used with the new trap-generation provided by SPARClike DSU. DSU will generate traps when a program accesses some data or instruction address that is assigned to the debug registers. However the hardware breakpoint registers can only take two data breakpoints, and GDB will reject this command if more than two are used. Delete or disable unused hardware breakpoints before setting new ones. See [Section 4.1.6 \[Break conditions\], page 24](#).

#### `thbreak args`

Set a hardware-assisted breakpoint enabled only for one stop. *args* are the same as for the `hbreak` command and the breakpoint is set in the same way. However, like the `tbreak` command, the breakpoint is automatically deleted after the first time your program stops there. Also, like the `hbreak` command, the breakpoint requires hardware support and some target hardware may not have this support. See [Section 4.1.5 \[Disabling breakpoints\], page 23](#). Also See [Section 4.1.6 \[Break conditions\], page 24](#).

#### `rbreak regex`

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

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

#### `info breakpoints [n]`

#### `info break [n]`

#### `info watchpoints [n]`

Print a table of all breakpoints and watchpoints set and not deleted, with the following columns for each breakpoint:

#### *Breakpoint Numbers*

*Type*            Breakpoint or watchpoint.

#### *Disposition*

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

#### *Enabled or Disabled*

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

<i>Address</i>	Where the breakpoint is in your program, as a memory address
<i>What</i>	Where the breakpoint is in the source for your program, as a file and line number.

If a breakpoint is conditional, `info break` shows the condition on the line following the affected breakpoint; breakpoint commands, if any, are listed after that.

`info break` with a breakpoint number *n* as argument lists only that breakpoint. The convenience variable `$_` and the default examining-address for the `x` command are set to the address of the last breakpoint listed (see [Section 7.5 \[Examining memory\]](#), page 45).

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

GDB allows you to set any number of breakpoints at the same place in your program. There is nothing silly or meaningless about this. When the breakpoints are conditional, this is even useful (see [Section 4.1.6 \[Break conditions\]](#), page 24).

GDB itself sometimes sets breakpoints in your program for special purposes, such as proper handling of `longjmp` (in C programs). These internal breakpoints are assigned negative numbers, starting with `-1`; ‘`info breakpoints`’ does not display them.

You can see these breakpoints with the GDB maintenance command ‘`maint info breakpoints`’.

#### `maint info breakpoints`

Using the same format as ‘`info breakpoints`’, display both the breakpoints you’ve set explicitly, and those GDB is using for internal purposes. Internal breakpoints are shown with negative breakpoint numbers. The type column identifies what kind of breakpoint is shown:

<code>breakpoint</code>	Normal, explicitly set breakpoint.
<code>watchpoint</code>	Normal, explicitly set watchpoint.
<code>longjmp</code>	Internal breakpoint, used to handle correctly stepping through <code>longjmp</code> calls.
<code>longjmp resume</code>	Internal breakpoint at the target of a <code>longjmp</code> .
<code>until</code>	Temporary internal breakpoint used by the GDB <code>until</code> command.
<code>finish</code>	Temporary internal breakpoint used by the GDB <code>finish</code> command.

### 4.1.2 Setting watchpoints

You can use a watchpoint to stop execution whenever the value of an expression changes, without having to predict a particular place where this may happen.

Watchpoints currently execute two orders of magnitude more slowly than other breakpoints, but this can be well worth it to catch errors where you have no clue what part of your program is the culprit.

#### `watch expr`

Set a watchpoint for an expression. GDB will break when `expr` is written into by the program and its value changes. This can be used with the new trap-generation provided by SPARClite DSU. DSU will generate traps when a program accesses some data or instruction address that is assigned to the debug registers. For the data addresses, DSU facilitates the `watch` command. However the hardware breakpoint registers can only take two data watchpoints, and both watchpoints must be the same kind. For example, you can set two watchpoints with `watch` commands, two with `rwatch` commands, **or** two with `awatch` commands, but you cannot set one watchpoint with one command and the other with a different command. {No value for “GBDN”} will reject the command if you try to mix watchpoints. Delete or disable unused watchpoint commands before setting new ones.

#### `rwatch expr`

Set a watchpoint that will break when `watch args` is read by the program. If you use both watchpoints, both must be set with the `rwatch` command.

#### `awatch expr`

Set a watchpoint that will break when `args` is read and written into by the program. If you use both watchpoints, both must be set with the `awatch` command.

#### `info watchpoints`

This command prints a list of watchpoints and breakpoints; it is the same as `info break`.

### 4.1.3 Breakpoints and exceptions

Some languages, such as GNU C++, implement exception handling. You can use GDB to examine what caused your program to raise an exception, and to list the exceptions your program is prepared to handle at a given point in time.

#### `break exception`

Ada mode uses `break exception` to set a breakpoint on exception handlers. See [Section 8.3.2.5 \[Ada Exceptions\], page 66](#).

#### `break assert`

Ada mode uses `break assert` to set a breakpoint on the raising of the `Assert_Failure` exception.

### catch exceptions

You can set breakpoints at active exception handlers by using the `catch` command. `exceptions` is a list of names of exceptions to catch.

You can use `info catch` to list active exception handlers. See [Section 5.4 \[Information about a frame\]](#), page 36.

There are currently some limitations to exception handling in GDB:

- If you call a function interactively, GDB normally returns control to you when the function has finished executing. If the call raises an exception, however, the call may bypass the mechanism that returns control to you and cause your program to simply continue running until it hits a breakpoint, catches a signal that GDB is listening for, or exits.
- You cannot raise an exception interactively.
- You cannot install an exception handler interactively.

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

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

```
/* addr is where the exception identifier is stored.
   ID is the exception identifier. */
void __raise_exception (void **addr, void *id);
```

To make the debugger catch all exceptions before any stack unwinding takes place, set a breakpoint on `__raise_exception` (see [Section 4.1 \[Breakpoints; watchpoints; and exceptions\]](#), page 17).

With a conditional breakpoint (see [Section 4.1.6 \[Break conditions\]](#), page 24) that depends on the value of `id`, you can stop your program when a specific exception is raised. You can use multiple conditional breakpoints to stop your program when any of a number of exceptions are raised.

### 4.1.4 Deleting breakpoints

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

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

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



**clear** Delete any breakpoints at the next instruction to be executed in the selected stack frame (see [Section 5.3 \[Selecting a frame\], page 35](#)). When the innermost frame is selected, this is a good way to delete a breakpoint where your program just stopped.

**clear** *function*

**clear** *filename: function*

Delete any breakpoints set at entry to the function *function*.

**clear** *linenum*

**clear** *filename: linenum*

Delete any breakpoints set at or within the code of the specified line.

**delete** [**breakpoints**] [*bnums...*]

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

### 4.1.5 Disabling breakpoints

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

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

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

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

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

**disable** [**breakpoints**] [*bnums...*]

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

**enable** [**breakpoints**] [*bnums...*]

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

`enable [breakpoints] once bnums...`

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

`enable [breakpoints] delete bnums...`

Enable the specified breakpoints to work once, then die. GDB deletes any of these breakpoints as soon as your program stops there.

Except for a breakpoint set with `tbreak` (see [Section 4.1.1 \[Setting breakpoints\], page 17](#)), breakpoints that you set are initially enabled; subsequently, they become disabled or enabled only when you use one of the commands above. (The command `until` can set and delete a breakpoint of its own, but it does not change the state of your other breakpoints; see [Section 4.2 \[Continuing and stepping\], page 27](#).)

### 4.1.6 Break conditions

The simplest sort of breakpoint breaks every time your program reaches a specified place. You can also specify a *condition* for a breakpoint. A condition is just a Boolean expression in your programming language (see [Section 7.1 \[Expressions\], page 41](#)). A breakpoint with a condition evaluates the expression each time your program reaches it, and your program stops only if the condition is *true*.

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

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

Break conditions can have side effects, and may even call functions in your program. This can be useful, for example, to activate functions that log program progress, or to use your own print functions to format special data structures. The effects are completely predictable unless there is another enabled breakpoint at the same address. (In that case, GDB might see the other breakpoint first and stop your program without checking the condition of this one.) Note that breakpoint commands are usually more convenient and flexible for the purpose of performing side effects when a breakpoint is reached (see [Section 4.1.7 \[Breakpoint command lists\], page 25](#)).

Break conditions can be specified when a breakpoint is set, by using ‘`if`’ in the arguments to the `break` command. See [Section 4.1.1 \[Setting breakpoints\], page 17](#). They can also be changed at any time with the `condition` command. The `watch` command does not recognize the `if` keyword; `condition` is the only way to impose a further condition on a watchpoint.

`condition bnum expression`

Specify *expression* as the break condition for breakpoint or watchpoint number *bnum*. After you set a condition, breakpoint *bnum* stops your program only if the value of *expression* is true (nonzero, in C). When you use `condition`,

GDB checks *expression* immediately for syntactic correctness, and to determine whether symbols in it have referents in the context of your breakpoint. GDB does not actually evaluate *expression* at the time the `condition` command is given, however. See [Section 7.1 \[Expressions\]](#), page 41.

`condition` *bnum*

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

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

`ignore` *bnum count*

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

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

When you use `continue` to resume execution of your program from a breakpoint, you can specify an ignore count directly as an argument to `continue`, rather than using `ignore`. See [Section 4.2 \[Continuing and stepping\]](#), page 27.

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

You could achieve the effect of the ignore count with a condition such as `'$foo-- <= 0'` using a debugger convenience variable that is decremented each time. See [Section 7.9 \[Convenience variables\]](#), page 53.

#### 4.1.7 Breakpoint command lists

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

`commands` [*bnum*]

... *command-list* ...

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

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

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

Pressing `(RET)` as a means of repeating the last GDB command is disabled within a *command-list*.

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

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

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

The commands `echo`, `output`, and `printf` allow you to print precisely controlled output, and are often useful in silent breakpoints. See [Section 14.4 \[Commands for controlled output\], page 103](#).

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

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

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

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

### 4.1.8 Breakpoint menus

Some programming languages (notably C++) permit a single function name to be defined several times, for application in different contexts. This is called *overloading*. When a function name is overloaded, ‘`break function`’ is not enough to tell GDB where you want a breakpoint. If you realize this is a problem, you can use something like ‘`break function(types)`’ to specify which particular version of the function you want. Otherwise, GDB offers you a menu of numbered choices for different possible breakpoints, and waits for your

selection with the prompt ‘>’. The first two options are always ‘[0] cancel’ and ‘[1] all’. Typing *1* sets a breakpoint at each definition of *function*, and typing *0* aborts the **break** command without setting any new breakpoints.

For example, the following session excerpt shows an attempt to set a breakpoint at the overloaded symbol `String::after`. We choose three particular definitions of that function name:

```
( gdb) b String::after
[0] cancel
[1] all
[2] file:String.cc; line number:867
[3] file:String.cc; line number:860
[4] file:String.cc; line number:875
[5] file:String.cc; line number:853
[6] file:String.cc; line number:846
[7] file:String.cc; line number:735
> 2 4 6
Breakpoint 1 at 0xb26c: file String.cc, line 867.
Breakpoint 2 at 0xb344: file String.cc, line 875.
Breakpoint 3 at 0xafcc: file String.cc, line 846.
Multiple breakpoints were set.
Use the "delete" command to delete unwanted
breakpoints.
( gdb)
```

## 4.2 Continuing and stepping

*Continuing* means resuming program execution until your program completes normally. In contrast, *stepping* means executing just one more “step” of your program, where “step” may mean either one line of source code, or one machine instruction (depending on what particular command you use). Either when continuing or when stepping, your program may stop even sooner, due to a breakpoint.

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

Resume program execution, at the address where your program last stopped; any breakpoints set at that address are bypassed. The optional argument *ignore-count* allows you to specify a further number of times to ignore a breakpoint at this location; its effect is like that of **ignore** (see [Section 4.1.6 \[Break conditions\]](#), page 24).

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

The synonyms **c** and **fg** are provided purely for convenience, and have exactly the same behavior as **continue**.

To resume execution at a different place, you can use **return** (see [Section 10.3 \[Returning from a function\]](#), page 76) to go back to the calling function; or **jump** (see [Section 10.2 \[Continuing at a different address\]](#), page 76) to go to an arbitrary location in your program.

A typical technique for using stepping is to set a breakpoint (see [Section 4.1 \[Breakpoints; watchpoints; and exceptions\]](#), page 17) at the beginning of the function or the section of your program where a problem is believed to lie, run your program until it stops at that breakpoint, and then step through the suspect area, examining the variables that are interesting, until you see the problem happen.

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

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

The **step** command now only stops at the first instruction of a source line. This prevents the multiple stops that used to occur in switch statements, for loops, etc. **step** continues to stop if a function that has debugging information is called within the line.

Also, the **step** command now only enters a subroutine if there is line number information for the subroutine. Otherwise it acts like the **next** command. This avoids problems when using **cc -g1** on MIPS machines. Previously, **step** entered subroutines if there was any debugging information about the routine.

**step count**

Continue running as in **step**, but do so *count* times. If a breakpoint is reached, stepping stops right away.

**next** [*count*]

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

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

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

**finish**

Continue running until just after function in the selected stack frame returns. Print the returned value (if any).

Contrast this with the **return** command (see [Section 10.3 \[Returning from a function\]](#), page 76).

**u**

**until**

Continue running until a source line past the current line, in the current stack frame, is reached. This command is used to avoid single stepping through a loop

more than once. It is like the `next` command, except that when `until` encounters a jump, it automatically continues execution until the program counter is greater than the address of the jump.

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

`until` always stops your program if it attempts to exit the current stack frame.

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

```
( gdb) f
#0  main (argc=4, argv=0xf7ffae8) at m4.c:206
206          expand_input();
( gdb) until
195          for ( ; argc > 0; NEXTARG) {
```

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

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

#### `until location`

`u location` Continue running your program until either the specified location is reached, or the current stack frame returns. *location* is any of the forms of argument acceptable to `break` (see [Section 4.1.1 \[Setting breakpoints\], page 17](#)). This form of the command uses breakpoints, and hence is quicker than `until` without an argument.

#### `stepi`

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

It is often useful to do `'display/i $pc'` when stepping by machine instructions. This makes GDB automatically display the next instruction to be executed, each time your program stops. See [Section 7.6 \[Automatic display\], page 46](#).

An argument is a repeat count, as in `step`.

#### `nexti`

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

An argument is a repeat count, as in `next`.

## 4.3 Signals

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

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

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

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

`info signals`

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

`info handle` is the new alias for `info signals`.

`handle signal keywords...`

Change the way GDB handles signal *signal*. *signal* can be the number of a signal or its name (with or without the ‘`SIG`’ at the beginning). The *keywords* say what change to make.

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

<code>nostop</code>	GDB should not stop your program when this signal happens. It may still print a message telling you that the signal has come in.
<code>stop</code>	GDB should stop your program when this signal happens. This implies the <code>print</code> keyword as well.
<code>print</code>	GDB should print a message when this signal happens.
<code>noprint</code>	GDB should not mention the occurrence of the signal at all. This implies the <code>nostop</code> keyword as well.
<code>pass</code>	GDB should allow your program to see this signal; your program can handle the signal, or else it may terminate if the signal is fatal and not handled.
<code>nopass</code>	GDB should not allow your program to see this signal.



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

You can also use the `signal` command to prevent your program from seeing a signal, or cause it to see a signal it normally would not see, or to give it any signal at any time. For example, if your program stopped due to some sort of memory reference error, you might store correct values into the erroneous variables and continue, hoping to see more execution; but your program would probably terminate immediately as a result of the fatal signal once it saw the signal. To prevent this, you can continue with `'signal 0'`. See [\[Giving your program a signal\]](#), page [\[undefined\]](#).



## 5 Examining the Stack

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

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

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

One of the stack frames is *selected* by GDB and many GDB commands refer implicitly to the selected frame. In particular, whenever you ask GDB for the value of a variable in your program, the value is found in the selected frame. There are special GDB commands to select whichever frame you are interested in. See [Section 5.3 \[Selecting a frame\], page 35](#).

When your program stops, GDB automatically selects the currently executing frame and describes it briefly, similar to the `frame` command (see [Section 5.4 \[Information about a frame\], page 36](#)).

### 5.1 Stack frames

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

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

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

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

Some compilers provide a way to compile functions so that they operate without stack frames. (For example, the `gcc` option `'-fomit-frame-pointer'` generates functions without a frame.) This is occasionally done with heavily used library functions to save the frame setup time. GDB has limited facilities for dealing with these function invocations. If the

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

**frame** *args*

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

**select-frame**

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

## 5.2 Backtraces

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

**backtrace**

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

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

**backtrace** *n*

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

**backtrace** *-n*

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

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

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

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

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

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

### 5.3 Selecting a frame

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

**frame** *n*

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

**frame** *addr*

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

On the SPARC architecture, **frame** needs two addresses to select an arbitrary frame: a frame pointer and a stack pointer.

On the MIPS and Alpha architecture, it needs two addresses: a stack pointer and a program counter.

On the 29k architecture, it needs three addresses: a register stack pointer, a program counter, and a memory stack pointer.

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

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

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

For example:

```
( gdb) up
#1 0x22f0 in main (argc=1, argv=0xf7ffbf4, env=0xf7ffbfc)
    at env.c:10
10          read_input_file (argv[i]);
```

After such a printout, the **list** command with no arguments prints ten lines centered on the point of execution in the frame. See [Section 6.1 \[Printing source lines\], page 37](#).

**up-silently** *n*

**down-silently** *n*

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

## 5.4 Information about a frame

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

`frame`

`f` When used without any argument, this command does not change which frame is selected, but prints a brief description of the currently selected stack frame. It can be abbreviated `f`. With an argument, this command is used to select a stack frame. See [Section 5.3 \[Selecting a frame\], page 35](#).

`info frame`

`info f` This command prints a verbose description of the selected stack frame, including:

- the address of the frame
- the address of the next frame down (called by this frame)
- the address of the next frame up (caller of this frame)
- the language in which the source code corresponding to this frame is written
- the address of the frame's arguments
- the program counter saved in it (the address of execution in the caller frame)
- which registers were saved in the frame

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

`info frame addr`

`info f addr`

Print a verbose description of the frame at address *addr*, without selecting that frame. The selected frame remains unchanged by this command. This requires the same kind of address (more than one for some architectures) that you specify in the `frame` command. See [Section 5.3 \[Selecting a frame\], page 35](#).

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

`info locals`

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

`info catch`

Print a list of all the exception handlers that are active in the current stack frame at the current point of execution. To see other exception handlers, visit the associated frame (using the `up`, `down`, or `frame` commands); then type `info catch`. See [Section 4.1.3 \[Breakpoints and exceptions\], page 21](#).

## 6 Examining Source Files

GDB can print parts of your program's source, since the debugging information recorded in the program tells GDB what source files were used to build it. When your program stops, GDB spontaneously prints the line where it stopped. Likewise, when you select a stack frame (see [Section 5.3 \[Selecting a frame\], page 35](#)), GDB prints the line where execution in that frame has stopped. You can print other portions of source files by explicit command.

If you use GDB through its GNU Emacs interface, you may prefer to use Emacs facilities to view source; see [Chapter 15 \[Using GDB under GNU Emacs\], page 105](#).

### 6.1 Printing source lines

To print lines from a source file, use the `list` command (abbreviated `l`). By default, ten lines are printed. There are several ways to specify what part of the file you want to print.

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

- `list linenum`  
Print lines centered around line number *linenum* in the current source file.
- `list function`  
Print lines centered around the beginning of function *function*.
- `list`  
Print more lines. If the last lines printed were printed with a `list` command, this prints lines following the last lines printed; however, if the last line printed was a solitary line printed as part of displaying a stack frame (see [Chapter 5 \[Examining the Stack\], page 33](#)), this prints lines centered around that line.
- `list -`  
Print lines just before the lines last printed.

By default, GDB prints ten source lines with any of these forms of the `list` command. You can change this using `set listsize`:

- `set listsize count`  
Make the `list` command display *count* source lines (unless the `list` argument explicitly specifies some other number).
- `show listsize`  
Display the number of lines that `list` prints.

Repeating a `list` command with `(RET)` discards the argument, so it is equivalent to typing just `list`. This is more useful than listing the same lines again. An exception is made for an argument of '-'; that argument is preserved in repetition so that each repetition moves up in the source file.

In general, the `list` command expects you to supply zero, one or two *linespecs*. Linespecs specify source lines; there are several ways of writing them but the effect is always to specify some source line. Here is a complete description of the possible arguments for `list`:

- `list linespec`  
Print lines centered around the line specified by *linespec*.

- `list first,last` Print lines from *first* to *last*. Both arguments are linespecs.
- `list ,last` Print lines ending with *last*.
- `list first,` Print lines starting with *first*.
- `list +` Print lines just after the lines last printed.
- `list -` Print lines just before the lines last printed.
- `list` As described in the preceding table.

Here are the ways of specifying a single source line—all the kinds of linespec.

- `number` Specifies line *number* of the current source file. When a `list` command has two linespecs, this refers to the same source file as the first linespec.
- `+offset` Specifies the line *offset* lines after the last line printed. When used as the second linespec in a `list` command that has two, this specifies the line *offset* lines down from the first linespec.
- `-offset` Specifies the line *offset* lines before the last line printed.
- `filename:number` Specifies line *number* in the source file *filename*.
- `function` Specifies the line that begins the body of the function *function*. For example: in C, this is the line with the open brace.
- `filename:function` Specifies the line of the open-brace that begins the body of the function *function* in the file *filename*. You only need the file name with a function name to avoid ambiguity when there are identically named functions in different source files.
- `*address` Specifies the line containing the program address *address*. *address* may be any expression.

## 6.2 Searching source files

There are two commands for searching through the current source file for a regular expression.

`forward-search regexp`

`search regexp`

The command ‘`forward-search regexp`’ checks each line, starting with the one following the last line listed, for a match for *regexp*. It lists the line that is found. You can use the synonym ‘`search regexp`’ or abbreviate the command name as `fo`.

`reverse-search regexp`

The command ‘`reverse-search regexp`’ checks each line, starting with the one before the last line listed and going backward, for a match for *regexp*. It lists the line that is found. You can abbreviate this command as `rev`.



### 6.3 Specifying source directories

Executable programs sometimes do not record the directories of the source files from which they were compiled, just the names. Even when they do, the directories could be moved between the compilation and your debugging session. GDB has a list of directories to search for source files; this is called the *source path*. Each time GDB wants a source file, it tries all the directories in the list, in the order they are present in the list, until it finds a file with the desired name. Note that the executable search path is *not* used for this purpose. Neither is the current working directory, unless it happens to be in the source path.

If GDB cannot find a source file in the source path, and the object program records a directory, GDB tries that directory too. If the source path is empty, and there is no record of the compilation directory, GDB looks in the current directory as a last resort.

Whenever you reset or rearrange the source path, GDB clears out any information it has cached about where source files are found and where each line is in the file.

When you start GDB, its source path is empty. To add other directories, use the `directory` command.

`directory dirname ...`

`dir dirname ...`

Add directory *dirname* to the front of the source path. Several directory names may be given to this command, separated by ‘:’ or whitespace. You may specify a directory that is already in the source path; this moves it forward, so GDB searches it sooner.

You can use the string ‘`$cdir`’ to refer to the compilation directory (if one is recorded), and ‘`$cwd`’ to refer to the current working directory. ‘`$cwd`’ is not the same as ‘.’—the former tracks the current working directory as it changes during your GDB session, while the latter is immediately expanded to the current directory at the time you add an entry to the source path.

`directory`

Reset the source path to empty again. This requires confirmation.

`show directories`

Print the source path: show which directories it contains.

If your source path is cluttered with directories that are no longer of interest, GDB may sometimes cause confusion by finding the wrong versions of source. You can correct the situation as follows:

1. Use `directory` with no argument to reset the source path to empty.
2. Use `directory` with suitable arguments to reinstall the directories you want in the source path. You can add all the directories in one command.

### 6.4 Source and machine code

You can use the command `info line` to map source lines to program addresses (and vice versa), and the command `disassemble` to display a range of addresses as machine

instructions. When run under GNU Emacs mode, the `info line` command now causes the arrow to point to the line specified. Also, `info line` prints addresses in symbolic form as well as hex.

#### `info line linespec`

Print the starting and ending addresses of the compiled code for source line *linespec*. You can specify source lines in any of the ways understood by the `list` command (see [Section 6.1 \[Printing source lines\]](#), page 37).

For example, we can use `info line` to discover the location of the object code for the first line of function `m4_changequote`:

```
( gdb) info line m4_changecom
```

```
Line 895 of "builtin.c" starts at pc 0x634c and ends at 0x6350.
```

We can also inquire (using `*addr` as the form for *linespec*) what source line covers a particular address:

```
( gdb) info line *0x63ff
```

```
Line 926 of "builtin.c" starts at pc 0x63e4 and ends at 0x6404.
```

After `info line`, the default address for the `x` command is changed to the starting address of the line, so that `'x/i'` is sufficient to begin examining the machine code (see [Section 7.5 \[Examining memory\]](#), page 45). Also, this address is saved as the value of the convenience variable `$_` (see [Section 7.9 \[Convenience variables\]](#), page 53).

#### `disassemble`

This specialized command dumps a range of memory as machine instructions. The default memory range is the function surrounding the program counter of the selected frame. A single argument to this command is a program counter value; GDB dumps the function surrounding this value. Two arguments specify a range of addresses (first inclusive, second exclusive) to dump.

We can use `disassemble` to inspect the object code range shown in the last `info line` example (the example shows SPARC machine instructions):

```
( gdb) disas 0x63e4 0x6404
```

```
Dump of assembler code from 0x63e4 to 0x6404:
```

```
0x63e4 <builtin_init+5340>:    ble 0x63f8 <builtin_init+5360>
0x63e8 <builtin_init+5344>:    sethi %hi(0x4c00), %o0
0x63ec <builtin_init+5348>:    ld [%i1+4], %o0
0x63f0 <builtin_init+5352>:    b 0x63fc <builtin_init+5364>
0x63f4 <builtin_init+5356>:    ld [%o0+4], %o0
0x63f8 <builtin_init+5360>:    or %o0, 0x1a4, %o0
0x63fc <builtin_init+5364>:    call 0x9288 <path_search>
0x6400 <builtin_init+5368>:    nop
```

```
End of assembler dump.
```

#### `set assembly-language instruction-set`

This command selects the instruction set to use when disassembling the program via the `disassemble` or `x/i` commands. It is useful for architectures that have more than one native instruction set.

Currently it is only defined for the Intel x86 family. You can set *instruction-set* to either `i386` or `i8086`. The default is `i386`.

## 7 Examining Data

The usual way to examine data in your program is with the `print` command (abbreviated `p`), or its synonym `inspect`. It evaluates and prints the value of an expression of the language your program is written in (see [Chapter 8 \[Using GDB with Different Languages\]](#), [page 57](#)).

```
print exp
print /f exp
```

`exp` is an expression (in the source language). By default the value of `exp` is printed in a format appropriate to its data type; you can choose a different format by specifying `'/f'`, where `f` is a letter specifying the format; see [Section 7.4 \[Output formats\]](#), [page 44](#).

```
print
print /f
```

If you omit `exp`, GDB displays the last value again (from the *value history*; see [Section 7.8 \[Value history\]](#), [page 52](#)). This allows you to conveniently inspect the same value in an alternative format.

A more low-level way of examining data is with the `x` command. It examines data in memory at a specified address and prints it in a specified format. See [Section 7.5 \[Examining memory\]](#), [page 45](#).

If you are interested in information about types, or about how the fields of a struct or class are declared, use the `ptype exp` command rather than `print`. See [Chapter 9 \[Examining the Symbol Table\]](#), [page 71](#).

### 7.1 Expressions

`print` and many other GDB commands accept an expression and compute its value. Any kind of constant, variable or operator defined by the programming language you are using is valid in an expression in GDB. This includes conditional expressions, function calls, casts and string constants. It unfortunately does not include symbols defined by preprocessor `#define` commands.

GDB now supports array constants in expressions input by the user. The syntax is `{element, element...}`. For example, you can now use the command `print {1, 2, 3}` to build up an array in memory that is `malloc`'d in the target program.

Because C is so widespread, most of the expressions shown in examples in this manual are in C. See [Chapter 8 \[Using GDB with Different Languages\]](#), [page 57](#), for information on how to use expressions in other languages.

In this section, we discuss operators that you can use in GDB expressions regardless of your programming language.

Casts are supported in all languages, not just in C, because it is so useful to cast a number into a pointer in order to examine a structure at that address in memory.

GDB supports these operators, in addition to those common to programming languages:

@        '@' is a binary operator for treating parts of memory as arrays. See [Section 7.3 \[Artificial arrays\]](#), [page 43](#), for more information.

`::` ‘`::`’ allows you to specify a variable in terms of the file or function where it is defined. See [Section 7.2 \[Program variables\]](#), page 42.

`{type} addr`

Refers to an object of type *type* stored at address *addr* in memory. *addr* may be any expression whose value is an integer or pointer (but parentheses are required around binary operators, just as in a cast). This construct is allowed regardless of what kind of data is normally supposed to reside at *addr*.

## 7.2 Program variables

The most common kind of expression to use is the name of a variable in your program.

Variables in expressions are understood in the selected stack frame (see [Section 5.3 \[Selecting a frame\]](#), page 35); they must be either:

- global (or static)

or

- visible according to the scope rules of the programming language from the point of execution in that frame

This means that in the function

```
foo (a)
    int a;
{
    bar (a);
    {
        int b = test ();
        bar (b);
    }
}
```

you can examine and use the variable `a` whenever your program is executing within the function `foo`, but you can only use or examine the variable `b` while your program is executing inside the block where `b` is declared.

There is an exception: you can refer to a variable or function whose scope is a single source file even if the current execution point is not in this file. But it is possible to have more than one such variable or function with the same name (in different source files). If that happens, referring to that name has unpredictable effects. If you wish, you can specify a static variable in a particular function or file, using the colon-colon notation:

```
file :: variable
function :: variable
```

Here *file* or *function* is the name of the context for the static *variable*. In the case of file names, you can use quotes to make sure GDB parses the file name as a single word—for example, to print a global value of `x` defined in ‘`f2.c`’:

```
( gdb ) p 'f2.c'::x
```

This use of ‘`::`’ is very rarely in conflict with the very similar use of the same notation in C++. GDB also supports use of the C++ scope resolution operator in GDB expressions.

*Warning:* Occasionally, a local variable may appear to have the wrong value at certain points in a function—just after entry to a new scope, and just before exit.

You may see this problem when you are stepping by machine instructions. This is because, on most machines, it takes more than one instruction to set up a stack frame (including local variable definitions); if you are stepping by machine instructions, variables may appear to have the wrong values until the stack frame is completely built. On exit, it usually also takes more than one machine instruction to destroy a stack frame; after you begin stepping through that group of instructions, local variable definitions may be gone.

### 7.3 Artificial arrays

It is often useful to print out several successive objects of the same type in memory; a section of an array, or an array of dynamically determined size for which only a pointer exists in the program.

You can do this by referring to a contiguous span of memory as an *artificial array*, using the binary operator '@'. The left operand of '@' should be the first element of the desired array and be an individual object. The right operand should be the desired length of the array. The result is an array value whose elements are all of the type of the left argument. The first element is actually the left argument; the second element comes from bytes of memory immediately following those that hold the first element, and so on. Here is an example. If a program says

```
int *array = (int *) malloc (len * sizeof (int));
```

you can print the contents of `array` with

```
p *array@len
```

The left operand of '@' must reside in memory. Array values made with '@' in this way behave just like other arrays in terms of subscripting, and are coerced to pointers when used in expressions. Artificial arrays most often appear in expressions via the value history (see [Section 7.8 \[Value history\]](#), page 52), after printing one out.

Another way to create an artificial array is to use a cast. This re-interprets a value as if it were an array. The value need not be in memory:

```
( gdb) p/x (short[2])0x12345678
$1 = {0x1234, 0x5678}
```

As a convenience, if you leave the array length out (as in '(*type*) [] *value*') gdb calculates the size to fill the value (as '`sizeof(value)/sizeof(type)`')

```
( gdb) p/x (short[])0x12345678
$2 = {0x1234, 0x5678}
```

Sometimes the artificial array mechanism is not quite enough; in moderately complex data structures, the elements of interest may not actually be adjacent—for example, if you are interested in the values of pointers in an array. One useful work-around in this situation is to use a convenience variable (see [Section 7.9 \[Convenience variables\]](#), page 53) as a counter in an expression that prints the first interesting value, and then repeat that expression via `(RET)`. For instance, suppose you have an array `dtab` of pointers to structures, and you are interested in the values of a field `fv` in each structure. Here is an example of what you might type:

```

set $i = 0
p dtab[$i++]>->fv
(RET)
(RET)
...

```

## 7.4 Output formats

By default, GDB prints a value according to its data type. Sometimes this is not what you want. For example, you might want to print a number in hex, or a pointer in decimal. Or you might want to view data in memory at a certain address as a character string or as an instruction. To do these things, specify an *output format* when you print a value.

The simplest use of output formats is to say how to print a value already computed. This is done by starting the arguments of the `print` command with a slash and a format letter. The format letters supported are:

- `x`      Regard the bits of the value as an integer, and print the integer in hexadecimal.
- `d`      Print as integer in signed decimal.
- `u`      Print as integer in unsigned decimal.
- `o`      Print as integer in octal.
- `t`      Print as integer in binary. The letter ‘`t`’ stands for “two”.<sup>1</sup>
- `a`      Print as an address, both absolute in hexadecimal and as an offset from the nearest preceding symbol. You can use this format used to discover where (in what function) an unknown address is located:
 

```
( gdb) p/a 0x54320
$3 = 0x54320 <_initialize_vx+396>
```
- `c`      Regard as an integer and print it as a character constant.
- `f`      Regard the bits of the value as a floating point number and print using typical floating point syntax.

For example, to print the program counter in hex (see [Section 7.10 \[Registers\], page 54](#)), type

```
p/x $pc
```

Note that no space is required before the slash; this is because command names in GDB cannot contain a slash.

To reprint the last value in the value history with a different format, you can use the `print` command with just a format and no expression. For example, ‘`p/x`’ reprints the last value in hex.

---

<sup>1</sup> ‘`b`’ cannot be used because these format letters are also used with the `x` command, where ‘`b`’ stands for “byte”; see [Section 7.5 \[Examining memory\], page 45](#).

## 7.5 Examining memory

You can use the command `x` (for “examine”) to examine memory in any of several formats, independently of your program’s data types.

`x/nfu addr`

`x addr`

`x`            Use the `x` command to examine memory.

`n`, `f`, and `u` are all optional parameters that specify how much memory to display and how to format it; `addr` is an expression giving the address where you want to start displaying memory. If you use defaults for `nfu`, you need not type the slash `/`. Several commands set convenient defaults for `addr`.

`n`, the repeat count

The repeat count is a decimal integer; the default is 1. It specifies how much memory (counting by units `u`) to display.

`f`, the display format

The display format is one of the formats used by `print`, `'s'` (null-terminated string), or `'i'` (machine instruction). The default is `'x'` (hexadecimal) initially. The default changes each time you use either `x` or `print`.

`u`, the unit size

The unit size is any of

`b`            Bytes.

`h`            Halfwords (two bytes).

`w`            Words (four bytes). This is the initial default.

`g`            Giant words (eight bytes).

Each time you specify a unit size with `x`, that size becomes the default unit the next time you use `x`. (For the `'s'` and `'i'` formats, the unit size is ignored and is normally not written.)

`addr`, starting display address

`addr` is the address where you want GDB to begin displaying memory. The expression need not have a pointer value (though it may); it is always interpreted as an integer address of a byte of memory. See [Section 7.1 \[Expressions\], page 41](#), for more information on expressions. The default for `addr` is usually just after the last address examined—but several other commands also set the default address: `info breakpoints` (to the address of the last breakpoint listed), `info line` (to the starting address of a line), and `print` (if you use it to display a value from memory).

For example, `'x/3uh 0x54320'` is a request to display three halfwords (`h`) of memory, formatted as unsigned decimal integers (`'u'`), starting at address `0x54320`. `'x/4xw $sp'` prints the four words (`'w'`) of memory above the stack pointer (here, `$sp`; see [Section 7.10 \[Registers\], page 54](#)) in hexadecimal (`'x'`).

Since the letters indicating unit sizes are all distinct from the letters specifying output formats, you do not have to remember whether unit size or format comes first; either order works. The output specifications ‘4xw’ and ‘4wx’ mean exactly the same thing. (However, the count *n* must come first; ‘wx4’ does not work.)

Even though the unit size *u* is ignored for the formats ‘s’ and ‘i’, you might still want to use a count *n*; for example, ‘3i’ specifies that you want to see three machine instructions, including any operands. The command `disassemble` gives an alternative way of inspecting machine instructions; see [Section 6.4 \[Source and machine code\], page 39](#).

All the defaults for the arguments to `x` are designed to make it easy to continue scanning memory with minimal specifications each time you use `x`. For example, after you have inspected three machine instructions with ‘x/3i *addr*’, you can inspect the next seven with just ‘x/7’. If you use `RET` to repeat the `x` command, the repeat count *n* is used again; the other arguments default as for successive uses of `x`.

The addresses and contents printed by the `x` command are not saved in the value history because there is often too much of them and they would get in the way. Instead, GDB makes these values available for subsequent use in expressions as values of the convenience variables `$_` and `$__`. After an `x` command, the last address examined is available for use in expressions in the convenience variable `$_`. The contents of that address, as examined, are available in the convenience variable `$__`.

If the `x` command has a repeat count, the address and contents saved are from the last memory unit printed; this is not the same as the last address printed if several units were printed on the last line of output.

## 7.6 Automatic display

If you find that you want to print the value of an expression frequently (to see how it changes), you might want to add it to the *automatic display list* so that GDB prints its value each time your program stops. Each expression added to the list is given a number to identify it; to remove an expression from the list, you specify that number. The automatic display looks like this:

```
2: foo = 38
3: bar[5] = (struct hack *) 0x3804
```

This display shows item numbers, expressions and their current values. As with displays you request manually using `x` or `print`, you can specify the output format you prefer; in fact, `display` decides whether to use `print` or `x` depending on how elaborate your format specification is—it uses `x` if you specify a unit size, or one of the two formats (‘i’ and ‘s’) that are only supported by `x`; otherwise it uses `print`.

`display exp`

Add the expression *exp* to the list of expressions to display each time your program stops. See [Section 7.1 \[Expressions\], page 41](#).

`display` does not repeat if you press `RET` again after using it.



**display/fmt exp**

For *fmt* specifying only a display format and not a size or count, add the expression *exp* to the auto-display list but arrange to display it each time in the specified format *fmt*. See [Section 7.4 \[Output formats\]](#), page 44.

**display/fmt addr**

For *fmt* ‘i’ or ‘s’, or including a unit-size or a number of units, add the expression *addr* as a memory address to be examined each time your program stops. Examining means in effect doing ‘x/fmt *addr*’. See [Section 7.5 \[Examining memory\]](#), page 45.

For example, ‘display/i \$pc’ can be helpful, to see the machine instruction about to be executed each time execution stops (‘\$pc’ is a common name for the program counter; see [Section 7.10 \[Registers\]](#), page 54).

**undisplay dnums...****delete display dnums...**

Remove item numbers *dnums* from the list of expressions to display.

**undisplay** does not repeat if you press **(RET)** after using it. (Otherwise you would just get the error ‘No display number ...’.)

**disable display dnums...**

Disable the display of item numbers *dnums*. A disabled display item is not printed automatically, but is not forgotten. It may be enabled again later.

**enable display dnums...**

Enable display of item numbers *dnums*. It becomes effective once again in auto display of its expression, until you specify otherwise.

**display** Display the current values of the expressions on the list, just as is done when your program stops.

**info display**

Print the list of expressions previously set up to display automatically, each one with its item number, but without showing the values. This includes disabled expressions, which are marked as such. It also includes expressions which would not be displayed right now because they refer to automatic variables not currently available.

If a display expression refers to local variables, then it does not make sense outside the lexical context for which it was set up. Such an expression is disabled when execution enters a context where one of its variables is not defined. For example, if you give the command **display last\_char** while inside a function with an argument **last\_char**, GDB displays this argument while your program continues to stop inside that function. When it stops elsewhere—where there is no variable **last\_char**—the display is disabled automatically. The next time your program stops where **last\_char** is meaningful, you can enable the display expression once again.

## 7.7 Print settings

GDB provides the following ways to control how arrays, structures, and symbols are printed.

These settings are useful for debugging programs in any language:

`set print address`

`set print address on`

GDB prints memory addresses showing the location of stack traces, structure values, pointer values, breakpoints, and so forth, even when it also displays the contents of those addresses. The default is `on`. For example, this is what a stack frame display looks like with `set print address on`:

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

`set print address off`

Do not print addresses when displaying their contents. For example, this is the same stack frame displayed with `set print address off`:

```
( gdb) set print addr off
( gdb) f
#0  set_quotes (lq="<<", rq=">>") at input.c:530
530          if (lquote != def_lquote)
```

You can use ‘`set print address off`’ to eliminate all machine dependent displays from the GDB interface. For example, with `print address off`, you should get the same text for backtraces on all machines—whether or not they involve pointer arguments.

`show print address`

Show whether or not addresses are to be printed.

When GDB prints a symbolic address, it normally prints the closest earlier symbol plus an offset. If that symbol does not uniquely identify the address (for example, it is a name whose scope is a single source file), you may need to clarify. One way to do this is with `info line`, for example ‘`info line *0x4537`’. Alternately, you can set GDB to print the source file and line number when it prints a symbolic address:

`set print symbol-filename on`

Tell GDB to print the source file name and line number of a symbol in the symbolic form of an address.

`set print symbol-filename off`

Do not print source file name and line number of a symbol. This is the default.

`show print symbol-filename`

Show whether or not GDB will print the source file name and line number of a symbol in the symbolic form of an address.

Another situation where it is helpful to show symbol filenames and line numbers is when disassembling code; GDB shows you the line number and source file that corresponds to each instruction.

Also, you may wish to see the symbolic form only if the address being printed is reasonably close to the closest earlier symbol:

```
set print max-symbolic-offset max-offset
```

Tell GDB to only display the symbolic form of an address if the offset between the closest earlier symbol and the address is less than *max-offset*. The default is 0, which tells GDB to always print the symbolic form of an address if any symbol precedes it.

```
show print max-symbolic-offset
```

Ask how large the maximum offset is that GDB prints in a symbolic address.

If you have a pointer and you are not sure where it points, try ‘`set print symbol-filename on`’. Then you can determine the name and source file location of the variable where it points, using ‘`p/a pointer`’. This interprets the address in symbolic form. For example, here GDB shows that a variable `ptt` points at another variable `t`, defined in ‘`hi2.c`’:

```
( gdb) set print symbol-filename on
( gdb) p/a ptt
$4 = 0xe008 <t in hi2.c>
```

*Warning:* For pointers that point to a local variable, ‘`p/a`’ does not show the symbol name and filename of the referent, even with the appropriate `set print` options turned on.

Other settings control how different kinds of objects are printed:

```
set print array
set print array on
```

Pretty print arrays. This format is more convenient to read, but uses more space. The default is off.

```
set print array off
```

Return to compressed format for arrays.

```
show print array
```

Show whether compressed or pretty format is selected for displaying arrays.

```
set print elements number-of-elements
```

Set a limit on how many elements of an array GDB will print. If GDB is printing a large array, it stops printing after it has printed the number of elements set by the `set print elements` command. This limit also applies to the display of strings. Setting *number-of-elements* to zero means that the printing is unlimited.

```
show print elements
```

Display the number of elements of a large array that GDB will print. If the number is 0, then the printing is unlimited.

**set print null-stop**

Cause GDB to stop printing the characters of an array when the first NULL is encountered. This is useful when large arrays actually contain only short strings.

**set print pretty on**

Cause GDB to print structures in an indented format with one member per line, like this:

```
$1 = {
  next = 0x0,
  flags = {
    sweet = 1,
    sour = 1
  },
  meat = 0x54 "Pork"
}
```

**set print pretty off**

Cause GDB to print structures in a compact format, like this:

```
$1 = {next = 0x0, flags = {sweet = 1, sour = 1}, \
  meat = 0x54 "Pork"}
```

This is the default format.

**show print pretty**

Show which format GDB is using to print structures.

**set print sevenbit-strings on**

Print using only seven-bit characters; if this option is set, GDB displays any eight-bit characters (in strings or character values) using the notation `\nnn`. This setting is best if you are working in English (ASCII) and you use the high-order bit of characters as a marker or “meta” bit.

**set print sevenbit-strings off**

Print full eight-bit characters. This allows the use of more international character sets, and is the default.

**show print sevenbit-strings**

Show whether or not GDB is printing only seven-bit characters.

**set print union on**

Tell GDB to print unions which are contained in structures. This is the default setting.

**set print union off**

Tell GDB not to print unions which are contained in structures.

**show print union**

Ask GDB whether or not it will print unions which are contained in structures.

For example, given the declarations

```
typedef enum {Tree, Bug} Species;
typedef enum {Big_tree, Acorn, Seedling} Tree_forms;
```

```
typedef enum {Caterpillar, Cocoon, Butterfly}
            Bug_forms;
```

```
struct thing {
    Species it;
    union {
        Tree_forms tree;
        Bug_forms bug;
    } form;
};
```

```
struct thing foo = {Tree, {Acorn}};
```

with `set print union on` in effect ‘p foo’ would print

```
$1 = {it = Tree, form = {tree = Acorn, bug = Cocoon}}
```

and with `set print union off` in effect it would print

```
$1 = {it = Tree, form = {...}}
```

These settings are of interest when debugging C++ programs:

`set print demangle`

`set print demangle on`

Print C++ names in their source form rather than in the encoded (“mangled”) form passed to the assembler and linker for type-safe linkage. The default is ‘on’.

`show print demangle`

Show whether C++ names are printed in mangled or demangled form.

`set print asm-demangle`

`set print asm-demangle on`

Print C++ names in their source form rather than their mangled form, even in assembler code printouts such as instruction disassemblies. The default is off.

`show print asm-demangle`

Show whether C++ names in assembly listings are printed in mangled or demangled form.

`set demangle-style style`

Choose among several encoding schemes used by different compilers to represent C++ names. The choices for *style* are currently:

**auto** Allow GDB to choose a decoding style by inspecting your program.

**gnu** Decode based on the GNU C++ compiler (`g++`) encoding algorithm. This is the default.

**lucid** Decode based on the Lucid C++ compiler (`lcc`) encoding algorithm.

**arm** Decode using the algorithm in the *C++ Annotated Reference Manual*. **Warning:** this setting alone is not sufficient to allow debugging `cfront`-generated executables. GDB would require further enhancement to permit that.

```

foo          Show the list of formats.

show demangle-style
    Display the encoding style currently in use for decoding C++ symbols.

set print object
set print object on
    When displaying a pointer to an object, identify the actual (derived) type of
    the object rather than the declared type, using the virtual function table.

set print object off
    Display only the declared type of objects, without reference to the virtual func-
    tion table. This is the default setting.

show print object
    Show whether actual, or declared, object types are displayed.

set print static-members
set print static-members on
    Print static members when displaying a C++ object. The default is on.

set print static-members off
    Do not print static members when displaying a C++ object.

show print static-members
    Show whether C++ static members are printed, or not.

set print vtbl
set print vtbl on
    Pretty print C++ virtual function tables. The default is off.

set print vtbl off
    Do not pretty print C++ virtual function tables.

show print vtbl
    Show whether C++ virtual function tables are pretty printed, or not.

```

## 7.8 Value history

Values printed by the `print` command are saved in the GDB *value history*. This allows you to refer to them in other expressions. Values are kept until the symbol table is re-read or discarded (for example with the `file` or `symbol-file` commands). When the symbol table changes, the value history is discarded, since the values may contain pointers back to the types defined in the symbol table.

The values printed are given *history numbers* by which you can refer to them. These are successive integers starting with one. `print` shows you the history number assigned to a value by printing `'$num = '` before the value; here *num* is the history number.

To refer to any previous value, use `'$'` followed by the value's history number. The way `print` labels its output is designed to remind you of this. Just `$` refers to the most recent value in the history, and `$$` refers to the value before that. `$$n` refers to the *n*th value from

the end; `$$2` is the value just prior to `$$`, `$$1` is equivalent to `$$`, and `$$0` is equivalent to `$`.

For example, suppose you have just printed a pointer to a structure and want to see the contents of the structure. It suffices to type

```
p *$
```

If you have a chain of structures where the component `next` points to the next one, you can print the contents of the next one with this:

```
p *$.next
```

You can print successive links in the chain by repeating this command—which you can do by just typing `(RET)`.

Note that the history records values, not expressions. If the value of `x` is 4 and you type these commands:

```
print x
set x=5
```

then the value recorded in the value history by the `print` command remains 4 even though the value of `x` has changed.

**show values**

Print the last ten values in the value history, with their item numbers. This is like `'p $$9'` repeated ten times, except that `show values` does not change the history.

**show values n**

Print ten history values centered on history item number *n*.

**show values +**

Print ten history values just after the values last printed. If no more values are available, `show values +` produces no display.

Pressing `(RET)` to repeat `show values n` has exactly the same effect as `'show values +'`.

## 7.9 Convenience variables

GDB provides *convenience variables* that you can use within GDB to hold on to a value and refer to it later. These variables exist entirely within GDB; they are not part of your program, and setting a convenience variable has no direct effect on further execution of your program. That is why you can use them freely.

Convenience variables are prefixed with `'$'`. Any name preceded by `'$'` can be used for a convenience variable, unless it is one of the predefined machine-specific register names (see [Section 7.10 \[Registers\]](#), page 54). (Value history references, in contrast, are *numbers* preceded by `'$'`. See [Section 7.8 \[Value history\]](#), page 52.)

You can save a value in a convenience variable with an assignment expression, just as you would set a variable in your program. For example:

```
set $foo = *object_ptr
```

would save in `$foo` the value contained in the object pointed to by `object_ptr`.

Using a convenience variable for the first time creates it, but its value is `void` until you assign a new value. You can alter the value with another assignment at any time.

Convenience variables have no fixed types. You can assign a convenience variable any type of value, including structures and arrays, even if that variable already has a value of a different type. The convenience variable, when used as an expression, has the type of its current value.

#### `show convenience`

Print a list of convenience variables used so far, and their values. Abbreviated `show con`.

One of the ways to use a convenience variable is as a counter to be incremented or a pointer to be advanced. For example, to print a field from successive elements of an array of structures:

```
set $i = 0
print bar[$i++]>contents
```

Repeat that command by typing `(RET)`.

Some convenience variables are created automatically by GDB and given values likely to be useful.

`$_` The variable `$_` is automatically set by the `x` command to the last address examined (see [Section 7.5 \[Examining memory\], page 45](#)). Other commands which provide a default address for `x` to examine also set `$_` to that address; these commands include `info line` and `info breakpoint`. The type of `$_` is `void *` except when set by the `x` command, in which case it is a pointer to the type of `$__`.

`$__` The variable `$__` is automatically set by the `x` command to the value found in the last address examined. Its type is chosen to match the format in which the data was printed.

#### `$_exitcode`

The variable `$_exitcode` is automatically set to the exit code when the program being debugged terminates.

## 7.10 Registers

You can refer to machine register contents, in expressions, as variables with names starting with `'$'`. The names of registers are different for each machine; use `info registers` to see the names used on your machine.

#### `info registers`

Print the names and values of all registers except floating-point registers (in the selected stack frame).

#### `info all-registers`

Print the names and values of all registers, including floating-point registers.



`info registers regname ...`

Print the *relativized* value of each specified register *regname*. As discussed in detail below, register values are normally relative to the selected stack frame. *regname* may be any register name valid on the machine you are using, with or without the initial ‘\$’.

GDB has four “standard” register names that are available (in expressions) on most machines—whenever they do not conflict with an architecture’s canonical mnemonics for registers. The register names `$pc` and `$sp` are used for the program counter register and the stack pointer. `$fp` is used for a register that contains a pointer to the current stack frame, and `$ps` is used for a register that contains the processor status. For example, you could print the program counter in hex with

```
p/x $pc
```

or print the instruction to be executed next with

```
x/i $pc
```

or add four to the stack pointer<sup>2</sup> with

```
set $sp += 4
```

Whenever possible, these four standard register names are available on your machine even though the machine has different canonical mnemonics, so long as there is no conflict. The `info registers` command shows the canonical names. For example, on the SPARC, `info registers` displays the processor status register as `$psr` but you can also refer to it as `$ps`.

GDB always considers the contents of an ordinary register as an integer when the register is examined in this way. Some machines have special registers which can hold nothing but floating point; these registers are considered to have floating point values. There is no way to refer to the contents of an ordinary register as floating point value (although you can *print* it as a floating point value with ‘`print/f $regname`’).

Some registers have distinct “raw” and “virtual” data formats. This means that the data format in which the register contents are saved by the operating system is not the same one that your program normally sees. For example, the registers of the 68881 floating point coprocessor are always saved in “extended” (raw) format, but all C programs expect to work with “double” (virtual) format. In such cases, GDB normally works with the virtual format only (the format that makes sense for your program), but the `info registers` command prints the data in both formats.

Normally, register values are relative to the selected stack frame (see [Section 5.3 \[Selecting a frame\]](#), page 35). This means that you get the value that the register would contain if all stack frames farther in were exited and their saved registers restored. In order to see the true contents of hardware registers, you must select the innermost frame (with ‘`frame 0`’).

---

<sup>2</sup> This is a way of removing one word from the stack, on machines where stacks grow downward in memory (most machines, nowadays). This assumes that the innermost stack frame is selected; setting `$sp` is not allowed when other stack frames are selected. To pop entire frames off the stack, regardless of machine architecture, use `return`; see [Section 10.3 \[Returning from a function\]](#), page 76.

However, GDB must deduce where registers are saved, from the machine code generated by your compiler. If some registers are not saved, or if GDB is unable to locate the saved registers, the selected stack frame makes no difference.

## 7.11 Floating point hardware

Depending on the configuration, GDB may be able to give you more information about the status of the floating point hardware.

`info float`

Display hardware-dependent information about the floating point unit. The exact contents and layout vary depending on the floating point chip. Currently, ‘`info float`’ is supported on the ARM and x86 machines.

## 8 Using GDB with Different Languages

Although programming languages generally have common aspects, they are rarely expressed in the same manner. For instance, in ANSI C, dereferencing a pointer `p` is accomplished by `*p`, but in Modula-2, it is accomplished by `p^`. Values can also be represented (and displayed) differently. Hex numbers in C appear as `'0x1ae'`, while in Modula-2 they appear as `'1AEH'`.

Language-specific information is built into GDB for some languages, allowing you to express operations like the above in your program's native language, and allowing GDB to output values in a manner consistent with the syntax of your program's native language. The language you use to build expressions is called the *working language*.

### 8.1 Switching between source languages

There are two ways to control the working language—either have GDB set it automatically, or select it manually yourself. You can use the `set language` command for either purpose. On startup, GDB defaults to setting the language automatically. The working language is used to determine how expressions you type are interpreted, how values are printed, etc.

In addition to the working language, every source file that GDB knows about has its own working language. For some object file formats, the compiler might indicate which language a particular source file is in. However, most of the time GDB infers the language from the name of the file. The language of a source file controls whether C++ names are demangled—this way `backtrace` can show each frame appropriately for its own language. There is no way to set the language of a source file from within GDB.

This is most commonly a problem when you use a program, such as `cfront` or `f2c`, that generates C but is written in another language. In that case, make the program use `#line` directives in its C output; that way GDB will know the correct language of the source code of the original program, and will display that source code, not the generated C code.

#### 8.1.1 List of filename extensions and languages

If a source file name ends in one of the following extensions, then GDB infers that its language is the one indicated.

<code>' .ada'</code>	
<code>' .ads'</code>	
<code>' .adb'</code>	
<code>' .a'</code>	Ada source file.
<code>' .c'</code>	C source file
<code>' .C'</code>	
<code>' .cc'</code>	
<code>' .cxx'</code>	
<code>' .cpp'</code>	
<code>' .cp'</code>	
<code>' .c++'</code>	C++ source file

`‘.ch’`  
`‘.c186’`  
`‘.c286’` CHILL source file.  
`‘.s’`  
`‘.S’` Assembler source file. This actually behaves almost like C, but GDB does not skip over function prologues when stepping.

### 8.1.2 Setting the working language

If you allow GDB to set the language automatically, expressions are interpreted the same way in your debugging session and your program.

If you wish, you may set the language manually. To do this, issue the command `‘set language lang’`, where *lang* is the name of a language, such as `c` or `modula-2`. For a list of the supported languages, type `‘set language’`.

Setting the language manually prevents GDB from updating the working language automatically. This can lead to confusion if you try to debug a program when the working language is not the same as the source language, when an expression is acceptable to both languages—but means different things. For instance, if the current source file were written in C, and GDB was parsing Modula-2, a command such as:

```
print a = b + c
```

might not have the effect you intended. In C, this means to add `b` and `c` and place the result in `a`. The result printed would be the value of `a`. In Modula-2, this means to compare `a` to the result of `b+c`, yielding a `BOOLEAN` value.

### 8.1.3 Having GDB infer the source language

To have GDB set the working language automatically, use `‘set language local’` or `‘set language auto’`. GDB then infers the working language. That is, when your program stops in a frame (usually by encountering a breakpoint), GDB sets the working language to the language recorded for the function in that frame. If the language for a frame is unknown (that is, if the function or block corresponding to the frame was defined in a source file that does not have a recognized extension), the current working language is not changed, and GDB issues a warning.

This may not seem necessary for most programs, which are written entirely in one source language. However, program modules and libraries written in one source language can be used by a main program written in a different source language. Using `‘set language auto’` in this case frees you from having to set the working language manually.

## 8.2 Displaying the language

The following commands help you find out which language is the working language, and also what language source files were written in.

`show language`

Display the current working language. This is the language you can use with commands such as `print` to build and compute expressions that may involve variables in your program.

**info frame**

Display the source language for this frame. This language becomes the working language if you use an identifier from this frame. See [Section 5.4 \[Information about a frame\]](#), page 36, to identify the other information listed here.

**info source**

Display the source language of this source file. See [Chapter 9 \[Examining the Symbol Table\]](#), page 71, to identify the other information listed here.

## 8.3 Supported languages

GDB 4 supports Ada, C, and C++. Some GDB features may be used in expressions regardless of the language you use: the GDB @ and :: operators, and the ‘{type}addr’ construct (see [Section 7.1 \[Expressions\]](#), page 41) can be used with the constructs of any supported language.

The following sections detail to what degree each source language is supported by GDB. These sections are not meant to be language tutorials or references, but serve only as a reference guide to what the GDB expression parser accepts, and what input and output formats should look like for different languages. There are many good books written on each of these languages; please look to these for a language reference or tutorial.

### 8.3.1 C and C++

Since C and C++ are so closely related, many features of GDB apply to both languages. Whenever this is the case, we discuss those languages together.

The C++ debugging facilities are jointly implemented by the GNU C++ compiler and GDB. Therefore, to debug your C++ code effectively, you must compile your C++ programs with the GNU C++ compiler, g++.

For best results when debugging C++ programs, use the stabs debugging format. You can select that format explicitly with the g++ command-line options ‘-gstabs’ or ‘-gstabs+’. See [section “Options for Debugging Your Program or GNU CC” in \*Using GNU CC\*](#), for more information.

#### 8.3.1.1 C and C++ operators

Operators must be defined on values of specific types. For instance, + is defined on numbers, but not on structures. Operators are often defined on groups of types.

For the purposes of C and C++, the following definitions hold:

- *Integral types* include `int` with any of its storage-class specifiers; `char`; and `enum`.
- *Floating-point types* include `float` and `double`.
- *Pointer types* include all types defined as `(type *)`.
- *Scalar types* include all of the above.

The following operators are supported. They are listed here in order of increasing precedence:

,	The comma or sequencing operator. Expressions in a comma-separated list are evaluated from left to right, with the result of the entire expression being the last expression evaluated.
=	Assignment. The value of an assignment expression is the value assigned. Defined on scalar types.
op=	Used in an expression of the form $a\ op=\ b$ , and translated to $a = a\ op\ b$ . $op=$ and $=$ have the same precedence. $op$ is any one of the operators $ $ , $\wedge$ , $\&$ , $\ll$ , $\gg$ , $+$ , $-$ , $*$ , $/$ , $\%$ .
?:	The ternary operator. $a\ ?\ b\ :\ c$ can be thought of as: if $a$ then $b$ else $c$ . $a$ should be of an integral type.
	Logical OR. Defined on integral types.
&&	Logical AND. Defined on integral types.
	Bitwise OR. Defined on integral types.
^	Bitwise exclusive-OR. Defined on integral types.
&	Bitwise AND. Defined on integral types.
==, !=	Equality and inequality. Defined on scalar types. The value of these expressions is 0 for false and non-zero for true.
<, >, <=, >=	Less than, greater than, less than or equal, greater than or equal. Defined on scalar types. The value of these expressions is 0 for false and non-zero for true.
<<, >>	left shift, and right shift. Defined on integral types.
@	The GDB “artificial array” operator (see <a href="#">Section 7.1 [Expressions]</a> , page 41).
+, -	Addition and subtraction. Defined on integral types, floating-point types and pointer types.
*, /, %	Multiplication, division, and modulus. Multiplication and division are defined on integral and floating-point types. Modulus is defined on integral types.
++, --	Increment and decrement. When appearing before a variable, the operation is performed before the variable is used in an expression; when appearing after it, the variable’s value is used before the operation takes place.
*	Pointer dereferencing. Defined on pointer types. Same precedence as $++$ .
&	Address operator. Defined on variables. Same precedence as $++$ . For debugging C++, GDB implements a use of ‘&’ beyond what is allowed in the C++ language itself: you can use ‘&(&ref)’ (or, if you prefer, simply ‘&&ref’) to examine the address where a C++ reference variable (declared with ‘&ref’) is stored.
-	Negative. Defined on integral and floating-point types. Same precedence as $++$ .
!	Logical negation. Defined on integral types. Same precedence as $++$ .

~	Bitwise complement operator. Defined on integral types. Same precedence as ++.
., ->	Structure member, and pointer-to-structure member. For convenience, GDB regards the two as equivalent, choosing whether to dereference a pointer based on the stored type information. Defined on <code>struct</code> and <code>union</code> data.
[]	Array indexing. <code>a[i]</code> is defined as <code>*(a+i)</code> . Same precedence as <code>-&gt;</code> .
()	Function parameter list. Same precedence as <code>-&gt;</code> .
::	C++ scope resolution operator. Defined on <code>struct</code> , <code>union</code> , and <code>class</code> types.
::	Doubled colons also represent the GDB scope operator (see <a href="#">Section 7.1 [Expressions]</a> , page 41). Same precedence as <code>::</code> , above.

### 8.3.1.2 C and C++ constants

GDB allows you to express the constants of C and C++ in the following ways:

- Integer constants are a sequence of digits. Octal constants are specified by a leading ‘0’ (i.e. zero), and hexadecimal constants by a leading ‘0x’ or ‘0X’. Constants may also end with a letter ‘l’, specifying that the constant should be treated as a `long` value.
- Floating point constants are a sequence of digits, followed by a decimal point, followed by a sequence of digits, and optionally followed by an exponent. An exponent is of the form: ‘e[+|-]nnn’, where *nnn* is another sequence of digits. The ‘+’ is optional for positive exponents.
- Enumerated constants consist of enumerated identifiers, or their integral equivalents.
- Character constants are a single character surrounded by single quotes (’), or a number—the ordinal value of the corresponding character (usually its ASCII value). Within quotes, the single character may be represented by a letter or by *escape sequences*, which are of the form ‘\nnn’, where *nnn* is the octal representation of the character’s ordinal value; or of the form ‘\x’, where ‘x’ is a predefined special character—for example, ‘\n’ for newline.
- String constants are a sequence of character constants surrounded by double quotes (“”).
- Pointer constants are an integral value. You can also write pointers to constants using the C operator ‘&’.
- Array constants are comma-separated lists surrounded by braces ‘{’ and ‘}’; for example, ‘{1,2,3}’ is a three-element array of integers, ‘{{1,2}, {3,4}, {5,6}}’ is a three-by-two array, and ‘{&"hi", &"there", &"fred"}’ is a three-element array of pointers.

### 8.3.1.3 C++ expressions

GDB expression handling has a number of extensions to interpret a significant subset of C++ expressions.

*Warning:* GDB can only debug C++ code if you compile with the GNU C++ compiler. Moreover, C++ debugging depends on the use of additional debugging

information in the symbol table, and thus requires special support. GDB has this support *only* with the stabs debug format. In particular, if your compiler generates a.out, MIPS ECOFF, RS/6000 XCOFF, or ELF with stabs extensions to the symbol table, these facilities are all available. (With GNU CC, you can use the ‘-gstabs’ option to request stabs debugging extensions explicitly.) Where the object code format is standard COFF or DWARF in ELF, on the other hand, most of the C++ support in GDB does *not* work.

1. Member function calls are allowed; you can use expressions like
 

```
count = aml->GetOriginal(x, y)
```
2. While a member function is active (in the selected stack frame), your expressions have the same namespace available as the member function; that is, GDB allows implicit references to the class instance pointer `this` following the same rules as C++.
3. You can call overloaded functions; GDB resolves the function call to the right definition, with one restriction—you must use arguments of the type required by the function that you want to call. GDB does not perform conversions requiring constructors or user-defined type operators.
4. GDB understands variables declared as C++ references; you can use them in expressions just as you do in C++ source—they are automatically dereferenced.

In the parameter list shown when GDB displays a frame, the values of reference variables are not displayed (unlike other variables); this avoids clutter, since references are often used for large structures. The *address* of a reference variable is always shown, unless you have specified ‘set print address off’.

5. GDB supports the C++ name resolution operator `::`—your expressions can use it just as expressions in your program do. Since one scope may be defined in another, you can use `::` repeatedly if necessary, for example in an expression like ‘*scope1::scope2::name*’. GDB also allows resolving name scope by reference to source files, in both C and C++ debugging (see [Section 7.2 \[Program variables\]](#), [page 42](#)).

### 8.3.1.4 C and C++ defaults

If you allow GDB to set type and range checking automatically, they both default to `off` whenever the working language changes to C or C++. This happens regardless of whether you or GDB selects the working language.

If you allow GDB to set the language automatically, it recognizes source files whose names end with ‘.c’, ‘.C’, or ‘.cc’, and when GDB enters code compiled from one of these files, it sets the working language to C or C++. See [Section 8.1.3 \[Having GDB infer the source language\]](#), [page 58](#), for further details.

### 8.3.1.5 GDB and C

The `set print union` and `show print union` commands apply to the `union` type. When set to ‘on’, any `union` that is inside a `struct` or `class` is also printed. Otherwise, it appears as ‘{...}’.

The `@` operator aids in the debugging of dynamic arrays, formed with pointers and a memory allocation function. See [Section 7.1 \[Expressions\]](#), [page 41](#).



### 8.3.1.6 GDB features for C++

Some GDB commands are particularly useful with C++, and some are designed specifically for use with C++. Here is a summary:

#### breakpoint menus

When you want a breakpoint in a function whose name is overloaded, GDB breakpoint menus help you specify which function definition you want. See [Section 4.1.8 \[Breakpoint menus\]](#), page 26.

#### rbreak *regex*

Setting breakpoints using regular expressions is helpful for setting breakpoints on overloaded functions that are not members of any special classes. See [Section 4.1.1 \[Setting breakpoints\]](#), page 17.

#### catch *exceptions*

#### info catch

Debug C++ exception handling using these commands. See [Section 4.1.3 \[Breakpoints and exceptions\]](#), page 21.

#### pptype *typename*

Print inheritance relationships as well as other information for type *typename*. See [Chapter 9 \[Examining the Symbol Table\]](#), page 71.

#### set print demangle

#### show print demangle

#### set print asm-demangle

#### show print asm-demangle

Control whether C++ symbols display in their source form, both when displaying code as C++ source and when displaying disassemblies. See [Section 7.7 \[Print settings\]](#), page 48.

#### set print object

#### show print object

Choose whether to print derived (actual) or declared types of objects. See [Section 7.7 \[Print settings\]](#), page 48.

#### set print vtbl

#### show print vtbl

Control the format for printing virtual function tables. See [Section 7.7 \[Print settings\]](#), page 48.

#### Overloaded symbol names

You can specify a particular definition of an overloaded symbol, using the same notation that is used to declare such symbols in C++: type *symbol(types)* rather than just *symbol*. You can also use the GDB command-line word completion facilities to list the available choices, or to finish the type list for you. See [Section 2.2 \[Command completion\]](#), page 9, for details on how to do this.

## 8.3.2 Ada

The extensions made to GDB for Ada only support output from the GNU Ada (GNAT) compiler (which is currently being developed). Other Ada compilers are not currently supported, and attempting to debug executables produced by them is most likely to be difficult.

### 8.3.2.1 Introduction

The Ada mode of GDB supports a fairly large subset of Ada expression syntax, with some extensions. The philosophy behind the design of this subset is

- That GDB should provide basic literals and access to operations for arithmetic, dereferencing, field selection, indexing, and subprogram calls, leaving more sophisticated computations to subprograms written into the program (which therefore may be called from GDB).
- That type safety and strict adherence to Ada language restrictions are not particularly important to the GDB user.
- That brevity is important to the GDB user.

Thus, for brevity, the debugger acts as if there were implicit `with` and `use` clauses in effect for all user-written packages, making it unnecessary to fully qualify most names with their packages, regardless of context. Where this causes ambiguity, GDB asks the user's intent.

The debugger will start in Ada mode if it detects an Ada main program. As for other languages, it will enter Ada mode when stopped in a program that was translated from an Ada source file.

While in Ada mode, you may use `--` for comments. This is useful mostly for documenting command files. The standard GDB comment (`#`) still works at the beginning of a line in Ada mode, but not in the middle (to allow based literals).

The debugger supports limited overloading. Given a subprogram call in which the function symbol has multiple definitions, it will use the number of actual parameters and some information about their types to attempt to narrow the set of definitions. It also makes very limited use of context, preferring procedures to functions in the context of the `call` command, and functions to procedures elsewhere.

### 8.3.2.2 Omissions from Ada

Here are the notable omissions from the subset:

- Only a subset of the attributes are supported:
  - `'First`, `'Last`, and `'Length` on array objects (not on types and subtypes).
  - `'Min` and `'Max`.
  - `'Pos` and `'Val`.
  - `'Tag`.
  - `'Range` on array objects (not subtypes), but only as the right operand of the membership (`in`) operator.

- 'Access, 'Unchecked\_Access, and 'Unrestricted\_Access (a GNAT extension).
- 'Address.
- The names in `Characters.Latin_1` are not available and concatenation is not implemented. Thus, escape characters in strings are not currently available.
- The component-by-component array operations (**and**, **or**, **xor**, **not**, and relational and equality tests) are not implemented.
- There are no record or array aggregates.
- Dispatching subprogram calls are not implemented.
- The overloading algorithm is much more limited (i.e., less selective) than that of real Ada. It does not use the context in which a subexpression appears to resolve its meaning, and it is much looser in its rules for allowing type matches. As a result, some function calls will be ambiguous, and the user will be asked to choose the proper resolution.
- The **new** operator is not implemented.
- Entry calls are not implemented.
- Aside from printing, arithmetic operations on the native VAX floating-point formats are not supported.

### 8.3.2.3 Additions to Ada

As it does for other languages, GDB makes certain generic extensions to Ada: the operators '@', '::', and '{type} addr' (see [Section 7.1 \[Expressions\], page 41](#)); convenience variables (see [Section 7.9 \[Convenience Vars\], page 53](#)); and machine registers (see [Section 7.10 \[Registers\], page 54](#)).

In addition, it provides a few other shortcuts and outright additions specific to Ada:

- The assignment statement is allowed as an expression, returning its right-hand operand as its value. Thus, you may enter
 

```
set x := y + 3
print A(tmp := y + 1)
```
- The semicolon is allowed as an “operator,” returning as its value the value of its right-hand operand. This allows, for example, complex conditional breaks:
 

```
break f
condition 1 (report(i); k += 1; A(k) > 100)
```
- Rather than use catenation and symbolic character names to introduce special characters into strings, one may instead use a special bracket notation, which is also used to print strings. A sequence of characters of the form '['XX']' within a string or character literal denotes the (single) character whose numeric encoding is XX in hexadecimal. The sequence of characters '['"]' also denotes a single quotation mark in strings. For example,
 

```
"One line.[0a]Next line.[0a]"
```

 Contains an ASCII newline character (`Ada.Characters.Latin_1.LF`) after each period.
- The subtype used as a prefix for the attributes 'Pos, 'Min, and 'Max is optional (and is ignored in any case). For example, it is legal to write

```
print 'max(x, y)
```

- When printing arrays, GDB uses positional notation when the array has a lower bound of 1, and uses a modified named notation otherwise. For example, a one-dimensional array of three integers with a lower bound of 3 might print as

```
(3 => 10, 17, 1)
```

That is, in contrast to valid Ada, only the first component has a => clause.

- You may abbreviate attributes in expressions with any unique, multi-character subsequence of their names (an exact match gets preference). For example, you may use `a'len`, `a'gth`, or `a'lh` in place of `a'length`.

### 8.3.2.4 Stopping at the Beginning

The main procedure in Ada has no fixed name, and attempts to break on `main` will position you to an utterly unhelpful location, in general. Therefore, Ada mode provides a convenient way to begin execution of the program and to stop at the beginning.

**begin** Does the equivalent of setting a temporary breakpoint at the beginning of the main procedure and then performing `run`. Since in general there is package elaboration code that runs before the main procedure begins, it is possible that the program will stop before reaching the main procedure. However, the temporary breakpoint will remain to halt execution.

### 8.3.2.5 Breaking on Ada Exceptions

In Ada mode, you can set breakpoints that trip when your program raises selected exceptions.

**break exception**

Set a breakpoint that trips whenever (any task in the) program raises any exception.

**break exception name**

Set a breakpoint that trips whenever (any task in the) program raises the exception *name*.

**break exception unhandled**

Set a breakpoint that trips whenever (any task in the) program raises an exception for which there is no handler.

**info exceptions**

**info exceptions *regexp***

The `info exceptions` command permits the user to examine all defined exceptions within Ada programs. With a regular expression, *regexp*, as argument, prints out only those exceptions whose name matches *regexp*.

### 8.3.2.6 Extensions for Ada Tasks

Support for Ada tasks is analogous to that for threads (see [\[Threads\]](#), page [\[Threads\]](#)). When in Ada mode (that is, when the “current language” is Ada), GDB allows the following task-related commands:

`info tasks`

This command shows a list of current Ada tasks, as in the following example:

```
(gdb) info tasks
      ID      TID P-ID Pri State          Name
      1  8088000  0  15 Child Activation Wait main_task
      2  80a4000  1  15 Accept Statement      b
      3  809a800  1  15 Child Activation Wait a
*     4  80ae800  3  15 Running                c
```

In this listing, the asterisk before the first task indicates it to be the currently running task.

**ID** Represents gdb's internal task number, used to refer to the task in conditional breakpoints and task switching.

**TID** The Ada `Task_ID`.

**P-ID** The parent's task ID (gdb's internal task number).

**Pri** The base priority of the task.

**State** Current state of the task.

- Unactivated: the task has been created but has not been activated. It cannot be executing.
- Running: the task currently running.
- Runnable: the task is not blocked for any reason known to Ada. (It may be waiting for a mutex, though.) It is conceptually "executing" in normal mode.
- Terminated: the task is terminated, in the sense of ARM 9.3 (5). Any dependents that were waiting on terminate alternatives have been awakened and have terminated themselves.
- Child Activation Wait: the task is waiting for created tasks to complete activation.
- Accept Statement: the task is waiting on an accept or selective wait statement.
- Waiting on entry call: the task is waiting on an entry call.
- Async Select Wait: the task is waiting to start the abortable part of an asynchronous select statement.
- Delay Sleep: the task is waiting on a select statement with only a delay alternative open.
- Child Termination Wait: the task is sleeping having completed a master within itself, and is waiting for the tasks dependent on that master to become terminated or waiting on a terminate Phase.
- Wait Child in Term Alt: the task is sleeping waiting for tasks on terminate alternatives to finish terminating.
- Accepting RV with *taskno*: the task is accepting a rendez-vous with the task *taskno*

Name            Name of the task in the program.

`info task taskno`

This command shows detailed informations on the specified task, as in the following example:

```
(gdb) info tasks
  ID      TID P-ID Pri State           Name
   1    8077880   0  15 Child Activation Wait  main_task
*  2    807c468   1  15 Running                task_1
(gdb) info task 2
Ada Task: 0x807c468
Name: task_1
Thread: 0x807f378
Parent: 1 (main_task)
Base Priority: 15
State: Runnable
```

`task taskno`

This command is like the `thread threadno` (see [\(undefined\)](#) [Thread Stops], page [\(undefined\)](#)). Note that this command is not implemented on all the targets. Currently supported targets are DEC Unix and Linux.

```
(gdb) info tasks
  ID      TID P-ID Pri State           Name
   1    8077870   0  15 Child Activation Wait  main_task
*  2    807c458   1  15 Running                t
(gdb) task 1
[Switching to task 1]
#0  0x8067726 in pthread_cond_wait ()
(gdb) bt
#0  0x8067726 in pthread_cond_wait ()
#1  0x8056714 in system.os_interface.pthread_cond_wait ()
#2  0x805cb63 in system.task_primitives.operations.sleep ()
#3  0x806153e in system.tasking.stages.activate_tasks ()
#4  0x804aacc in un () at un.adb:5
#5  0x8049d56 in main (argc=1, argv=3221223804, envp=3221223812)
    at b~un.adb:163
```

`break linespec task taskno`

`break linespec task taskno if ...`

These commands are like the `break ... thread ...` (see [\(undefined\)](#) [Thread Stops], page [\(undefined\)](#)). `linespec` specifies source lines, as described in [Section 4.1.1 \[Setting breakpoints\], page 17](#).

Use the qualifier ‘`task taskno`’ with a breakpoint command to specify that you only want GDB to stop the program when a particular Ada task reaches this breakpoint. `taskno` is one of the numeric task identifiers assigned by GDB, shown in the first column of the ‘`info tasks`’ display.

If you do not specify ‘`task taskno`’ when you set a breakpoint, the breakpoint applies to *all* tasks of your program.

You can use the `task` qualifier on conditional breakpoints as well; in this case, place `'task taskno'` before the breakpoint condition (before the `if`).

For example,

```
(gdb) info tasks
  ID      TID P-ID Pri State           Name
  1 140022020 0 15 Child Activation Wait main_task
  2 140045060 1 15 Accept/Select Wait t2
  3 140044840 1 15 Runnable t1
* 4 140056040 1 15 Running t3
(gdb) b 15 task 2
Breakpoint 5 at 0x120044cb0: file test_task_debug.adb, line 15.
(gdb) cont
Continuing.
task # 1 running
task # 2 running

Breakpoint 5, test_task_debug.tB.1 (_task=0x11ffffc60)
  at test_task_debug.adb:15
15 flush;
(gdb) info tasks
  ID      TID P-ID Pri State           Name
  1 140022020 0 15 Child Activation Wait main_task
* 2 140045060 1 15 Running t2
  3 140044840 1 15 Runnable t1
  4 140056040 1 15 Delay Sleep t3
```

### 8.3.2.7 Debugging Generic Units

GNAT always uses code expansion for generic instantiation. This means that each time an instantiation occurs, a complete copy of the original code is made with appropriate substitutions.

It is not possible to refer to the original generic entities in GDB, but it is certainly possible to debug a particular instance of a generic, simply by using the appropriate expanded names. For example, if we have

```
procedure g is

  generic package k is
    procedure kp (v1 : in out integer);
  end k;

  package body k is
    procedure kp (v1 : in out integer) is
    begin
      v1 := v1 + 1;
    end kp;
  end k;
```

```

package k1 is new k;
package k2 is new k;

begin
  k1.kp;
  k2.kp;
  k1.kp;
  k2.kp;
end;
```

Then to break on a call to procedure `kp` in the `k2` instance, simply use the command:

```
break g.k2.kp
```

When the breakpoint occurs, you can step through the code of the generic instance in the normal manner. You can also examine values of data in the normal manner, providing the appropriate generic package qualification to refer to non-local entities.

### 8.3.2.8 Set commands for Ada

Ada introduces one new `set` command.

```
set varsize-limit size
```

Limit the size of the types of objects when those sizes are computed from run-time quantities to *size* bytes. When this is set to 0, there is no limit. By default, it is about 65K.

```
show varsize-limit
```

Show the limit on types whose size is determined by run-time quantities.

### 8.3.2.9 Known Peculiarities of Ada Mode

Besides the omissions listed previously (see [Section 8.3.2.2 \[Omissions from Ada\], page 64](#)), we know of several problems with and limitations of Ada mode in GDB, some of which will be fixed with planned future releases of the debugger and the GNU Ada compiler.

- Currently, the debugger has insufficient information to determine whether certain pointers represent pointers to objects or the objects themselves. Thus, the user may have to tack an extra `.all` after an expression to get it printed properly.
- Static constants that the compiler chooses not to materialize as objects in storage are invisible to the debugger.
- Renaming declarations are invisible to the debugger.
- Named parameter associations in function argument lists are ignored (the argument lists are treated as positional).
- Many useful library packages are currently invisible to the debugger.
- Fixed-point arithmetic, conversions, input, and output is carried out using floating-point arithmetic, and may give results that only approximate those on the host machine.
- The type of the `'Address` attribute may not be `System.Address`.



## 9 Examining the Symbol Table

The commands described in this section allow you to inquire about the symbols (names of variables, functions and types) defined in your program. This information is inherent in the text of your program and does not change as your program executes. GDB finds it in your program's symbol table, in the file indicated when you started GDB (see [Section 1.1.1 \[Choosing files\]](#), page 5), or by one of the file-management commands (see [Section 11.1 \[Commands to specify files\]](#), page 79).

Occasionally, you may need to refer to symbols that contain unusual characters, which GDB ordinarily treats as word delimiters. The most frequent case is in referring to static variables in other source files (see [Section 7.2 \[Program variables\]](#), page 42). File names are recorded in object files as debugging symbols, but GDB would ordinarily parse a typical file name, like `'foo.c'`, as the three words `'foo'` `'.'` `'c'`. To allow GDB to recognize `'foo.c'` as a single symbol, enclose it in single quotes; for example,

```
p 'foo.c'::x
```

looks up the value of `x` in the scope of the file `'foo.c'`.

**info address *symbol***

Describe where the data for *symbol* is stored. For a register variable, this says which register it is kept in. For a non-register local variable, this prints the stack-frame offset at which the variable is always stored.

Note the contrast with `'print &symbol'`, which does not work at all for a register variable, and for a stack local variable prints the exact address of the current instantiation of the variable.

**whatis *exp***

Print the data type of expression *exp*. *exp* is not actually evaluated, and any side-effecting operations (such as assignments or function calls) inside it do not take place. See [Section 7.1 \[Expressions\]](#), page 41.

**whatis**     Print the data type of `$`, the last value in the value history.

**ptype *typename***

Print a description of data type *typename*. *typename* may be the name of a type, or for C code it may have the form `'class class-name'`, `'struct struct-tag'`, `'union union-tag'` or `'enum enum-tag'`.

**ptype *exp***

**ptype**     Print a description of the type of expression *exp*. **ptype** differs from **whatis** by printing a detailed description, instead of just the name of the type.

For example, for this variable declaration:

```
struct complex {double real; double imag;} v;
```

the two commands give this output:

```
( gdb) whatis v
type = struct complex
( gdb) ptype v
type = struct complex {
    double real;
    double imag;
}
```

As with `whatis`, using `ptype` without an argument refers to the type of `$`, the last value in the value history.

`info types regexp`

`info types`

Print a brief description of all types whose name matches *regexp* (or all types in your program, if you supply no argument). Each complete typename is matched as though it were a complete line; thus, `'i type value'` gives information on all types in your program whose name includes the string *value*, but `'i type ^value$'` gives information only on types whose complete name is *value*.

This command differs from `ptype` in two ways: first, like `whatis`, it does not print a detailed description; second, it lists all source files where a type is defined.

`info source`

Show the name of the current source file—that is, the source file for the function containing the current point of execution—and the language it was written in.

`info sources`

Print the names of all source files in your program for which there is debugging information, organized into two lists: files whose symbols have already been read, and files whose symbols will be read when needed.

`info functions`

Print the names and data types of all defined functions.

`info functions regexp`

Print the names and data types of all defined functions whose names contain a match for regular expression *regexp*. Thus, `'info fun step'` finds all functions whose names include *step*; `'info fun ^step'` finds those whose names start with *step*.

`info variables`

Print the names and data types of all variables that are declared outside of functions (i.e., excluding local variables).

`info variables regexp`

Print the names and data types of all variables (except for local variables) whose names contain a match for regular expression *regexp*.

Some systems allow individual object files that make up your program to be replaced without stopping and restarting your program. For example, in Vx-Works you can simply recompile a defective object file and keep on running. If you are running on one of these systems, you can allow GDB to reload the symbols for automatically relinked modules:

`set symbol-reloading on`

Replace symbol definitions for the corresponding source file when an object file with a particular name is seen again.

`set symbol-reloading off`

Do not replace symbol definitions when re-encountering object files of the same name. This is the default state; if you are not running on a system that permits automatically relinking modules, you should leave `symbol-reloading` off, since otherwise GDB may discard symbols when linking large programs, that may contain several modules (from different directories or libraries) with the same name.

`show symbol-reloading`

Show the current `on` or `off` setting.

`maint print symbols filename`

`maint print psymbols filename`

`maint print msymbols filename`

Write a dump of debugging symbol data into the file *filename*. These commands are used to debug the GDB symbol-reading code. Only symbols with debugging data are included. If you use ‘`maint print symbols`’, GDB includes all the symbols for which it has already collected full details: that is, *filename* reflects symbols for only those files whose symbols GDB has read. You can use the command `info sources` to find out which files these are. If you use ‘`maint print psymbols`’ instead, the dump shows information about symbols that GDB only knows partially—that is, symbols defined in files that GDB has skimmed, but not yet read completely. Finally, ‘`maint print msymbols`’ dumps just the minimal symbol information required for each object file from which GDB has read some symbols. See [Section 11.1 \[Commands to specify files\], page 79](#), for a discussion of how GDB reads symbols (in the description of `symbol-file`).



## 10 Altering Execution

Once you think you have found an error in your program, you might want to find out for certain whether correcting the apparent error would lead to correct results in the rest of the run. You can find the answer by experiment, using the GDB features for altering execution of the program.

For example, you can store new values into variables or memory locations, restart your program at a different address, or even return prematurely from a function.

### 10.1 Assignment to variables

To alter the value of a variable, evaluate an assignment expression. See [Section 7.1 \[Expressions\], page 41](#). For example,

```
print x=4
```

stores the value 4 into the variable `x`, and then prints the value of the assignment expression (which is 4). See [Chapter 8 \[Using GDB with Different Languages\], page 57](#), for more information on operators in supported languages.

If you are not interested in seeing the value of the assignment, use the `set` command instead of the `print` command. `set` is really the same as `print` except that the expression's value is not printed and is not put in the value history (see [Section 7.8 \[Value history\], page 52](#)). The expression is evaluated only for its effects.

If the beginning of the argument string of the `set` command appears identical to a `set` subcommand, use the `set variable` command instead of just `set`. This command is identical to `set` except for its lack of subcommands. For example, if your program has a variable `width`, you get an error if you try to set a new value with just `'set width=13'`, because GDB has the command `set width`:

```
( gdb) whatis width
type = double
( gdb) p width
$4 = 13
( gdb) set width=47
Invalid syntax in expression.
```

The invalid expression, of course, is `'=47'`. In order to actually set the program's variable `width`, use

```
( gdb) set var width=47
```

GDB allows more implicit conversions in assignments than C; you can freely store an integer value into a pointer variable or vice versa, and you can convert any structure to any other structure that is the same length or shorter.

To store values into arbitrary places in memory, use the `'{...}'` construct to generate a value of specified type at a specified address (see [Section 7.1 \[Expressions\], page 41](#)). For example, `{int}0x83040` refers to memory location `0x83040` as an integer (which implies a certain size and representation in memory), and

```
set {int}0x83040 = 4
```

stores the value 4 into that memory location.

## 10.2 Continuing at a different address

Ordinarily, when you continue your program, you do so at the place where it stopped, with the `continue` command. You can instead continue at an address of your own choosing, with the following commands:

`jump linespec`

Resume execution at line *linespec*. Execution stops again immediately if there is a breakpoint there. See [Section 6.1 \[Printing source lines\], page 37](#), for a description of the different forms of *linespec*.

The `jump` command does not change the current stack frame, or the stack pointer, or the contents of any memory location or any register other than the program counter. If line *linespec* is in a different function from the one currently executing, the results may be bizarre if the two functions expect different patterns of arguments or of local variables. For this reason, the `jump` command requests confirmation if the specified line is not in the function currently executing. However, even bizarre results are predictable if you are well acquainted with the machine-language code of your program.

`jump *address`

Resume execution at the instruction at address *address*.

You can get much the same effect as the `jump` command by storing a new value into the register `$pc`. The difference is that this does not start your program running; it only changes the address of where it *will* run when you continue. For example,

```
set $pc = 0x485
```

makes the next `continue` command or stepping command execute at address `0x485`, rather than at the address where your program stopped. See [Section 4.2 \[Continuing and stepping\], page 27](#).

The most common occasion to use the `jump` command is to back up—perhaps with more breakpoints set—over a portion of a program that has already executed, in order to examine its execution in more detail.

## 10.3 Returning from a function

`return`

`return expression`

You can cancel execution of a function call with the `return` command. If you give an *expression* argument, its value is used as the function's return value.

When you use `return`, GDB discards the selected stack frame (and all frames within it). You can think of this as making the discarded frame return prematurely. If you wish to specify a value to be returned, give that value as the argument to `return`.

This pops the selected stack frame (see [Section 5.3 \[Selecting a frame\], page 35](#)), and any other frames inside of it, leaving its caller as the innermost remaining frame. That frame becomes selected. The specified value is stored in the registers used for returning values of functions.

The `return` command does not resume execution; it leaves the program stopped in the state that would exist if the function had just returned. In contrast, the `finish` command (see [Section 4.2 \[Continuing and stepping\], page 27](#)) resumes execution until the selected stack frame returns naturally.

## 10.4 Calling program functions

`call expr` Evaluate the expression `expr` without displaying `void` returned values.

You can use this variant of the `print` command if you want to execute a function from your program, but without cluttering the output with `void` returned values. If the result is not `void`, it is printed and saved in the value history.

A new user-controlled variable, `call_scratch_address`, specifies the location of a scratch area to be used when GDB calls a function in the target. This is necessary because the usual method of putting the scratch area on the stack does not work in systems that have separate instruction and data spaces.

## 10.5 Patching programs

By default, GDB opens the file containing your program's executable code read-only. This prevents accidental alterations to machine code; but it also prevents you from intentionally patching your program's binary.

If you'd like to be able to patch the binary, you can specify that explicitly with the `set write` command. For example, you might want to turn on internal debugging flags, or even to make emergency repairs.

```
set write on
```

```
set write off
```

If you specify '`set write on`', GDB opens executable files for both reading and writing; if you specify '`set write off`' (the default), GDB opens them read-only.

If you have already loaded a file, you must load it again (using the `exec-file` command) after changing `set write`, for your new setting to take effect.

```
show write
```

Display whether executable files are opened for writing as well as reading.





## 11 GDB Files

GDB needs to know the file name of the program to be debugged, both in order to read its symbol table and in order to start your program.

### 11.1 Commands to specify files

The usual way to specify an executable file name is with the command argument given when you start GDB, (see [Chapter 1 \[Getting In and Out of GDB\], page 5](#)).

Occasionally it is necessary to change to a different file during a GDB session. Or you may run GDB and forget to specify a file you want to use. In these situations the GDB commands to specify new files are useful.

**file** *filename*

Use *filename* as the program to be debugged. It is read for its symbols and for the contents of pure memory. It is also the program executed when you use the **run** command. If you do not specify a directory and the file is not found in the GDB working directory, GDB uses the environment variable `PATH` as a list of directories to search, just as the shell does when looking for a program to run. You can change the value of this variable, for both GDB and your program, using the **path** command.

On systems with memory-mapped files, an auxiliary file '*filename.syms*' may hold symbol table information for *filename*. If so, GDB maps in the symbol table from '*filename.syms*', starting up more quickly. See the descriptions of the file options '**-mapped**' and '**-readnow**' (available on the command line, and with the commands **file**, **symbol-file**, or **add-symbol-file**, described below), for more information.

**file** *file* with no argument makes GDB discard any information it has on both executable file and the symbol table.

**exec-file** [ *filename* ]

Specify that the program to be run (but not the symbol table) is found in *filename*. GDB searches the environment variable `PATH` if necessary to locate your program. Omitting *filename* means to discard information on the executable file.

**symbol-file** [ *filename* ]

Read symbol table information from file *filename*. `PATH` is searched when necessary. Use the **file** command to get both symbol table and program to run from the same file.

**symbol-file** with no argument clears out GDB information on your program's symbol table.

The **symbol-file** command causes GDB to forget the contents of its convenience variables, the value history, and all breakpoints and auto-display expressions. This is because they may contain pointers to the internal data recording symbols and data types, which are part of the old symbol table data being discarded inside GDB.

`symbol-file` does not repeat if you press `(RET)` again after executing it once.

When GDB is configured for a particular environment, it understands debugging information in whatever format is the standard generated for that environment; you may use either a GNU compiler, or other compilers that adhere to the local conventions. Best results are usually obtained from GNU compilers; for example, using `gcc` you can generate debugging information for optimized code.

On some kinds of object files, the `symbol-file` command does not normally read the symbol table in full right away. Instead, it scans the symbol table quickly to find which source files and which symbols are present. The details are read later, one source file at a time, as they are needed.

The purpose of this two-stage reading strategy is to make GDB start up faster. For the most part, it is invisible except for occasional pauses while the symbol table details for a particular source file are being read. (The `set verbose` command can turn these pauses into messages if desired. See [Section 13.6 \[Optional warnings and messages\]](#), page 100.)

We have not implemented the two-stage strategy for COFF yet. When the symbol table is stored in COFF format, `symbol-file` reads the symbol table data in full right away.

```
symbol-file filename [ -readnow ] [ -mapped ]
file filename [ -readnow ] [ -mapped ]
```

You can override the GDB two-stage strategy for reading symbol tables by using the `-readnow` option with any of the commands that load symbol table information, if you want to be sure GDB has the entire symbol table available.

`load filename`

Depending on what remote debugging facilities are configured into GDB, the `load` command may be available. Where it exists, it is meant to make *filename* (an executable) available for debugging on the remote system—by downloading, or dynamic linking, for example. `load` also records the *filename* symbol table in GDB, like the `add-symbol-file` command.

If your GDB does not have a `load` command, attempting to execute it gets the error message “You can’t do that when your target is ...”

The file is loaded at whatever address is specified in the executable. For some object file formats, you can specify the load address when you link the program; for other formats, like `a.out`, the object file format specifies a fixed address.

On VxWorks, `load` links *filename* dynamically on the current target system as well as adding its symbols in GDB.

`load` does not repeat if you press `(RET)` again after using it.

`section`

The `section` command changes the base address of section `SECTION` of the exec file to `ADDR`. This can be used if the exec file does not contain section addresses, (such as in the `a.out` format), or when the addresses specified in the file itself are wrong. Each section must be changed separately. The “info files” command lists all the sections and their addresses.

`info files`  
`info target`

`info files` and `info target` are synonymous; both print the current target (see [Chapter 12 \[Specifying a Debugging Target\]](#), page 83), including the name of the executable file currently in use by GDB, and the files from which symbols were loaded. The command `help target` lists all possible targets rather than current ones.

All file-specifying commands allow both absolute and relative file names as arguments. GDB always converts the file name to an absolute file name and remembers it that way.

## 11.2 Errors reading symbol files

While reading a symbol file, GDB occasionally encounters problems, such as symbol types it does not recognize, or known bugs in compiler output. By default, GDB does not notify you of such problems, since they are relatively common and primarily of interest to people debugging compilers. If you are interested in seeing information about ill-constructed symbol tables, you can either ask GDB to print only one message about each such type of problem, no matter how many times the problem occurs; or you can ask GDB to print more messages, to see how many times the problems occur, with the `set complaints` command (see [Section 13.6 \[Optional warnings and messages\]](#), page 100).

The messages currently printed, and their meanings, include:

`inner block not inside outer block in symbol`

The symbol information shows where symbol scopes begin and end (such as at the start of a function or a block of statements). This error indicates that an inner scope block is not fully contained in its outer scope blocks.

GDB circumvents the problem by treating the inner block as if it had the same scope as the outer block. In the error message, *symbol* may be shown as “(don't know)” if the outer block is not a function.

`block at address out of order`

The symbol information for symbol scope blocks should occur in order of increasing addresses. This error indicates that it does not do so.

GDB does not circumvent this problem, and has trouble locating symbols in the source file whose symbols it is reading. (You can often determine what source file is affected by specifying `set verbose on`. See [Section 13.6 \[Optional warnings and messages\]](#), page 100.)

`bad block start address patched`

The symbol information for a symbol scope block has a start address smaller than the address of the preceding source line. This is known to occur in the SunOS 4.1.1 (and earlier) C compiler.

GDB circumvents the problem by treating the symbol scope block as starting on the previous source line.

`bad string table offset in symbol n`

Symbol number *n* contains a pointer into the string table which is larger than the size of the string table.

GDB circumvents the problem by considering the symbol to have the name `foo`, which may cause other problems if many symbols end up with this name.

**unknown symbol type `0xnn`**

The symbol information contains new data types that GDB does not yet know how to read. `0xnn` is the symbol type of the misunderstood information, in hexadecimal.

GDB circumvents the error by ignoring this symbol information. This usually allows you to debug your program, though certain symbols are not accessible. If you encounter such a problem and feel like debugging it, you can debug `gdb` with itself, breakpoint on `complain`, then go up to the function `read_dbx_syntab` and examine `*bufp` to see the symbol.

**stub type has NULL name**

GDB could not find the full definition for a struct or class.

**const/volatile indicator missing (ok if using `g++ v1.x`), got...**

The symbol information for a C++ member function is missing some information that recent versions of the compiler should have output for it.

**info mismatch between compiler and debugger**

GDB could not parse a type specification output by the compiler.

## 12 Specifying a Debugging Target

A *target* is the execution environment occupied by your program. You can use the `target` command to specify one of the target types configured for GDB (see [Section 12.2 \[Commands for managing targets\]](#), page 83).

### 12.1 Active targets

When you type `run`, your executable file becomes an active process target as well. When a process target is active, all GDB commands requesting memory addresses refer to that target; addresses in an executable file target are obscured while the process target is active.

Use the `exec-file` command to select a new executable target (see [Section 11.1 \[Commands to specify files\]](#), page 79).

### 12.2 Commands for managing targets

`target type parameters`

Connects the GDB host environment to a target machine.

The `target` command does not repeat if you press `RET` again after executing the command.

`help target`

Displays the names of all targets available. To display targets currently selected, use either `info target` or `info files` (see [Section 11.1 \[Commands to specify files\]](#), page 79).

`help target name`

Describe a particular target, including any parameters necessary to select it.

`set gnutarget args`

GDB uses its own library BFD to read your files. GDB knows whether it is reading an *executable*, a *core*, or a *.o* file, however you can specify the file format with the `set gnutarget` command. Unlike most `target` commands, with `gnutarget` the `target` refers to a program, not a machine.

*Warning:* To specify a file format with `set gnutarget`, you must know the actual BFD name.

See [Section 11.1 \[Commands to specify files\]](#), page 79.

`show gnutarget`

Use the `show gnutarget` command to display what file format `gnutarget` is set to read. If you have not set `gnutarget`, GDB will determine the file format for each file automatically and `show gnutarget` displays `The current BDF target is "auto"`.

Here are some common targets (available, or not, depending on the GDB configuration):

`target exec program`

An executable file. ‘`target exec program`’ is the same as ‘`exec-file program`’.

**target remote *dev***

Remote serial target in GDB-specific protocol. The argument *dev* specifies what serial device to use for the connection (e.g. `‘/dev/ttya’`). See [Section 12.4 \[Remote debugging\], page 85](#). `target remote` now supports the `load` command. This is only useful if you have some other way of getting the stub to the target system, and you can put it somewhere in memory where it won't get clobbered by the download.

**target sim**

CPU simulator. See [Section 12.4.4 \[Simulated CPU Target\], page 95](#).

**target vxworks *machinename***

A VxWorks system, attached via TCP/IP. The argument *machinename* is the target system's machine name or IP address. See [Section 12.4.2 \[GDB and VxWorks\], page 92](#).

**target bug *dev***

BUG monitor, running on a MVME187 (m88k) board.

**target cpu32bug *dev***

CPU32BUG monitor, running on a CPU32 (M68K) board.

**target op50n *dev***

OP50N monitor, running on an OKI HPPA board.

**target w89k *dev***

W89K monitor, running on a Winbond HPPA board.

**target est *dev***

EST-300 ICE monitor, running on a CPU32 (M68K) board.

**target rom68k *dev***

ROM 68K monitor, running on an IDP board.

**target array *dev***

Array Tech LSI33K RAID controller board.

**target sparclite *dev***

Fujitsu sparclite boards, used only for the purpose of loading. You must use an additional command to debug the program. For example: `target remote dev` using GDB standard remote protocol.

Different targets are available on different configurations of GDB; your configuration may have more or fewer targets.

## 12.3 Choosing target byte order

You can now choose which byte order to use with a target system. Use the `set endian big` and `set endian little` commands. Use the `set endian auto` command to instruct GDB to use the byte order associated with the executable. You can see the current setting for byte order with the `show endian` command.

*Warning:* Currently, only embedded MIPS configurations support dynamic selection of target byte order.

## 12.4 Remote debugging

If you are trying to debug a program running on a machine that cannot run GDB in the usual way, it is often useful to use remote debugging. For example, you might use remote debugging on an operating system kernel, or on a small system which does not have a general purpose operating system powerful enough to run a full-featured debugger.

Some configurations of GDB have special serial or TCP/IP interfaces to make this work with particular debugging targets. In addition, GDB comes with a generic serial protocol (specific to GDB, but not specific to any particular target system) which you can use if you write the remote stubs—the code that runs on the remote system to communicate with GDB.

Other remote targets may be available in your configuration of GDB; use `help target` to list them.

### 12.4.1 The GDB remote serial protocol

To debug a program running on another machine (the debugging *target* machine), you must first arrange for all the usual prerequisites for the program to run by itself. For example, for a C program, you need:

1. A startup routine to set up the C runtime environment; these usually have a name like `'crt0'`. The startup routine may be supplied by your hardware supplier, or you may have to write your own.
2. You probably need a C subroutine library to support your program's subroutine calls, notably managing input and output.
3. A way of getting your program to the other machine—for example, a download program. These are often supplied by the hardware manufacturer, but you may have to write your own from hardware documentation.

The next step is to arrange for your program to use a serial port to communicate with the machine where GDB is running (the *host* machine). In general terms, the scheme looks like this:

*On the host,*

GDB already understands how to use this protocol; when everything else is set up, you can simply use the `'target remote'` command (see [Chapter 12 \[Specifying a Debugging Target\], page 83](#)).

*On the target,*

you must link with your program a few special-purpose subroutines that implement the GDB remote serial protocol. The file containing these subroutines is called a *debugging stub*.

On certain remote targets, you can use an auxiliary program `gdbserver` instead of linking a stub into your program. See [Section 12.4.1.5 \[Using the `gdbserver` program\], page 90](#), for details.

The debugging stub is specific to the architecture of the remote machine; for example, use `'sparc-stub.c'` to debug programs on SPARC boards.

These working remote stubs are distributed with GDB:

`i386-stub.c`

For Intel 386 and compatible architectures.

`m68k-stub.c`

For Motorola 680x0 architectures.

`sh-stub.c`

For Hitachi SH architectures.

`sparc-stub.c`

For SPARC architectures.

`sparcl-stub.c`

For Fujitsu SPARCLITE architectures.

The ‘README’ file in the GDB distribution may list other recently added stubs.

### 12.4.1.1 What the stub can do for you

The debugging stub for your architecture supplies these three subroutines:

`set_debug_traps`

This routine arranges for `handle_exception` to run when your program stops. You must call this subroutine explicitly near the beginning of your program.

`handle_exception`

This is the central workhorse, but your program never calls it explicitly—the setup code arranges for `handle_exception` to run when a trap is triggered.

`handle_exception` takes control when your program stops during execution (for example, on a breakpoint), and mediates communications with GDB on the host machine. This is where the communications protocol is implemented; `handle_exception` acts as the GDB representative on the target machine; it begins by sending summary information on the state of your program, then continues to execute, retrieving and transmitting any information GDB needs, until you execute a GDB command that makes your program resume; at that point, `handle_exception` returns control to your own code on the target machine.

`breakpoint`

Use this auxiliary subroutine to make your program contain a breakpoint. Depending on the particular situation, this may be the only way for GDB to get control. For instance, if your target machine has some sort of interrupt button, you won’t need to call this; pressing the interrupt button transfers control to `handle_exception`—in effect, to GDB. On some machines, simply receiving characters on the serial port may also trigger a trap; again, in that situation, you don’t need to call `breakpoint` from your own program—simply running ‘target remote’ from the host GDB session gets control.

Call `breakpoint` if none of these is true, or if you simply want to make certain your program stops at a predetermined point for the start of your debugging session.



### 12.4.1.2 What you must do for the stub

The debugging stubs that come with GDB are set up for a particular chip architecture, but they have no information about the rest of your debugging target machine.

First of all you need to tell the stub how to communicate with the serial port.

`int getDebugChar()`

Write this subroutine to read a single character from the serial port. It may be identical to `getchar` for your target system; a different name is used to allow you to distinguish the two if you wish.

`void putDebugChar(int)`

Write this subroutine to write a single character to the serial port. It may be identical to `putchar` for your target system; a different name is used to allow you to distinguish the two if you wish.

If you want GDB to be able to stop your program while it is running, you need to use an interrupt-driven serial driver, and arrange for it to stop when it receives a `^C` (`'\003'`, the control-C character). That is the character which GDB uses to tell the remote system to stop.

Getting the debugging target to return the proper status to GDB probably requires changes to the standard stub; one quick and dirty way is to just execute a breakpoint instruction (the “dirty” part is that GDB reports a `SIGTRAP` instead of a `SIGINT`).

Other routines you need to supply are:

`void exceptionHandler (int exception_number, void *exception_address)`

Write this function to install `exception_address` in the exception handling tables. You need to do this because the stub does not have any way of knowing what the exception handling tables on your target system are like (for example, the processor’s table might be in ROM, containing entries which point to a table in RAM). `exception_number` is the exception number which should be changed; its meaning is architecture-dependent (for example, different numbers might represent divide by zero, misaligned access, etc). When this exception occurs, control should be transferred directly to `exception_address`, and the processor state (stack, registers, and so on) should be just as it is when a processor exception occurs. So if you want to use a jump instruction to reach `exception_address`, it should be a simple jump, not a jump to subroutine.

For the 386, `exception_address` should be installed as an interrupt gate so that interrupts are masked while the handler runs. The gate should be at privilege level 0 (the most privileged level). The SPARC and 68k stubs are able to mask interrupt themselves without help from `exceptionHandler`.

`void flush_i_cache()`

(sparc and sparclite only) Write this subroutine to flush the instruction cache, if any, on your target machine. If there is no instruction cache, this subroutine may be a no-op.

On target machines that have instruction caches, GDB requires this function to make certain that the state of your program is stable.

You must also make sure this library routine is available:

```
void *memset(void *, int, int)
```

This is the standard library function `memset` that sets an area of memory to a known value. If you have one of the free versions of `libc.a`, `memset` can be found there; otherwise, you must either obtain it from your hardware manufacturer, or write your own.

If you do not use the GNU C compiler, you may need other standard library subroutines as well; this varies from one stub to another, but in general the stubs are likely to use any of the common library subroutines which `gcc` generates as inline code.

### 12.4.1.3 Putting it all together

In summary, when your program is ready to debug, you must follow these steps.

1. Make sure you have the supporting low-level routines (see [Section 12.4.1.2 \[What you must do for the stub\]](#), page 87):

```
getDebugChar, putDebugChar,  
flush_i_cache, memset, exceptionHandler.
```

2. Insert these lines near the top of your program:

```
set_debug_traps();  
breakpoint();
```

3. For the 680x0 stub only, you need to provide a variable called `exceptionHook`. Normally you just use:

```
void (*exceptionHook)() = 0;
```

but if before calling `set_debug_traps`, you set it to point to a function in your program, that function is called when GDB continues after stopping on a trap (for example, bus error). The function indicated by `exceptionHook` is called with one parameter: an `int` which is the exception number.

4. Compile and link together: your program, the GDB debugging stub for your target architecture, and the supporting subroutines.
5. Make sure you have a serial connection between your target machine and the GDB host, and identify the serial port on the host.
6. Download your program to your target machine (or get it there by whatever means the manufacturer provides), and start it.
7. To start remote debugging, run GDB on the host machine, and specify as an executable file the program that is running in the remote machine. This tells GDB how to find your program's symbols and the contents of its pure text.

Then establish communication using the `target remote` command. Its argument specifies how to communicate with the target machine—either via a devicename attached to a direct serial line, or a TCP port (usually to a terminal server which in turn has a serial line to the target). For example, to use a serial line connected to the device named `‘/dev/ttyb’`:

```
target remote /dev/ttyb
```

To use a TCP connection, use an argument of the form `host:port`. For example, to connect to port 2828 on a terminal server named `manyfarms`:

```
target remote manyfarms:2828
```

Now you can use all the usual commands to examine and change data and to step and continue the remote program.

To resume the remote program and stop debugging it, use the `detach` command.

Whenever GDB is waiting for the remote program, if you type the interrupt character (often `^C`), GDB attempts to stop the program. This may or may not succeed, depending in part on the hardware and the serial drivers the remote system uses. If you type the interrupt character once again, GDB displays this prompt:

```
Interrupted while waiting for the program.
Give up (and stop debugging it)? (y or n)
```

If you type `y`, GDB abandons the remote debugging session. (If you decide you want to try again later, you can use ‘`target remote`’ again to connect once more.) If you type `n`, GDB goes back to waiting.

#### 12.4.1.4 Communication protocol

The stub files provided with GDB implement the target side of the communication protocol, and the GDB side is implemented in the GDB source file ‘`remote.c`’. Normally, you can simply allow these subroutines to communicate, and ignore the details. (If you’re implementing your own stub file, you can still ignore the details: start with one of the existing stub files. ‘`sparc-stub.c`’ is the best organized, and therefore the easiest to read.)

However, there may be occasions when you need to know something about the protocol—for example, if there is only one serial port to your target machine, you might want your program to do something special if it recognizes a packet meant for GDB.

All GDB commands and responses (other than acknowledgements, which are single characters) are sent as a packet which includes a checksum. A packet is introduced with the character ‘`$`’, and ends with the character ‘`#`’ followed by a two-digit checksum:

```
$packet info#checksum
```

*checksum* is computed as the modulo 256 sum of the *packet info* characters.

When either the host or the target machine receives a packet, the first response expected is an acknowledgement: a single character, either ‘`+`’ (to indicate the package was received correctly) or ‘`-`’ (to request retransmission).

The host (GDB) sends commands, and the target (the debugging stub incorporated in your program) sends data in response. The target also sends data when your program stops.

Command packets are distinguished by their first character, which identifies the kind of command.

These are some of the commands currently supported (for a complete list of commands, look in ‘`gdb/remote.c`’):

```
g          Requests the values of CPU registers.
G          Sets the values of CPU registers.
maddr,count
           Read count bytes at location addr.
```

*Maddr,count:...*

Write *count* bytes at location *addr*.

**c**

*caddr* Resume execution at the current address (or at *addr* if supplied).

**s**

*saddr* Step the target program for one instruction, from either the current program counter or from *addr* if supplied.

**k**

Kill the target program.

**?**

Report the most recent signal. To allow you to take advantage of the GDB signal handling commands, one of the functions of the debugging stub is to report CPU traps as the corresponding POSIX signal values.

**T**

Allows the remote stub to send only the registers that GDB needs to make a quick decision about single-stepping or conditional breakpoints. This eliminates the need to fetch the entire register set for each instruction being stepped through.

The GDB remote serial protocol now implements a write-through cache for registers. GDB only re-reads the registers if the target has run.

If you have trouble with the serial connection, you can use the command `set remotedebug`. This makes GDB report on all packets sent back and forth across the serial line to the remote machine. The packet-debugging information is printed on the GDB standard output stream. `set remotedebug off` turns it off, and `show remotedebug` shows you its current state.

### 12.4.1.5 Using the `gdbserver` program

`gdbserver` is a control program for Unix-like systems, which allows you to connect your program with a remote GDB via `target remote`—but without linking in the usual debugging stub.

`gdbserver` is not a complete replacement for the debugging stubs, because it requires essentially the same operating-system facilities that GDB itself does. In fact, a system that can run `gdbserver` to connect to a remote GDB could also run GDB locally! `gdbserver` is sometimes useful nevertheless, because it is a much smaller program than GDB itself. It is also easier to port than all of GDB, so you may be able to get started more quickly on a new system by using `gdbserver`. Finally, if you develop code for real-time systems, you may find that the tradeoffs involved in real-time operation make it more convenient to do as much development work as possible on another system, for example by cross-compiling. You can use `gdbserver` to make a similar choice for debugging.

GDB and `gdbserver` communicate via either a serial line or a TCP connection, using the standard GDB remote serial protocol.

*On the target machine,*

you need to have a copy of the program you want to debug. `gdbserver` does not need your program's symbol table, so you can strip the program if necessary to save space. GDB on the host system does all the symbol handling.

To use the server, you must tell it how to communicate with GDB; the name of your program; and the arguments for your program. The syntax is:

```
target> gdbserver comm program [ args ... ]
```

*comm* is either a device name (to use a serial line) or a TCP hostname and portnumber. For example, to debug Emacs with the argument 'foo.txt' and communicate with GDB over the serial port '/dev/com1':

```
target> gdbserver /dev/com1 emacs foo.txt
```

`gdbserver` waits passively for the host GDB to communicate with it.

To use a TCP connection instead of a serial line:

```
target> gdbserver host:2345 emacs foo.txt
```

The only difference from the previous example is the first argument, specifying that you are communicating with the host GDB via TCP. The 'host:2345' argument means that `gdbserver` is to expect a TCP connection from machine 'host' to local TCP port 2345. (Currently, the 'host' part is ignored.) You can choose any number you want for the port number as long as it does not conflict with any TCP ports already in use on the target system (for example, 23 is reserved for `telnet`).<sup>1</sup> You must use the same port number with the host GDB `target remote` command.

*On the GDB host machine,*

you need an unstripped copy of your program, since GDB needs symbols and debugging information. Start up GDB as usual, using the name of the local copy of your program as the first argument. (You may also need the '--baud' option if the serial line is running at anything other than 9600 bps.) After that, use `target remote` to establish communications with `gdbserver`. Its argument is either a device name (usually a serial device, like '/dev/ttyb'), or a TCP port descriptor in the form *host:PORT*. For example:

```
( gdb) target remote /dev/ttyb
```

communicates with the server via serial line '/dev/ttyb', and

```
( gdb) target remote the-target:2345
```

communicates via a TCP connection to port 2345 on host 'the-target'. For TCP connections, you must start up `gdbserver` prior to using the `target remote` command. Otherwise you may get an error whose text depends on the host system, but which usually looks something like 'Connection refused'.

#### 12.4.1.6 Using the `gdbserve.nlm` program

`gdbserve.nlm` is a control program for NetWare systems, which allows you to connect your program with a remote GDB via `target remote`.

GDB and `gdbserve.nlm` communicate via a serial line, using the standard GDB remote serial protocol.

---

<sup>1</sup> If you choose a port number that conflicts with another service, `gdbserver` prints an error message and exits.

*On the target machine,*

you need to have a copy of the program you want to debug. `gdbserve.nlm` does not need your program's symbol table, so you can strip the program if necessary to save space. GDB on the host system does all the symbol handling.

To use the server, you must tell it how to communicate with GDB; the name of your program; and the arguments for your program. The syntax is:

```
load gdbserve [ BOARD=board ] [ PORT=port ]
               [ BAUD=baud ] program [ args ... ]
```

*board* and *port* specify the serial line; *baud* specifies the baud rate used by the connection. *port* and *node* default to 0, *baud* defaults to 9600 bps.

For example, to debug Emacs with the argument '`foo.txt`' and communicate with GDB over serial port number 2 or board 1 using a 19200 bps connection:

```
load gdbserve BOARD=1 PORT=2 BAUD=19200 emacs foo.txt
```

*On the GDB host machine,*

you need an unstripped copy of your program, since GDB needs symbols and debugging information. Start up GDB as usual, using the name of the local copy of your program as the first argument. (You may also need the '`--baud`' option if the serial line is running at anything other than 9600 bps. After that, use `target remote` to establish communications with `gdbserve.nlm`. Its argument is a device name (usually a serial device, like '`/dev/ttyb`'). For example:

```
(gdb) target remote /dev/ttyb
```

communications with the server via serial line '`/dev/ttyb`'.

## 12.4.2 GDB and VxWorks

GDB enables developers to spawn and debug tasks running on networked VxWorks targets from a Unix host. Already-running tasks spawned from the VxWorks shell can also be debugged. GDB uses code that runs on both the Unix host and on the VxWorks target. The program `gdb` is installed and executed on the Unix host. (It may be installed with the name `vxgdb`, to distinguish it from a GDB for debugging programs on the host itself.)

### VxWorks-timeout *args*

All VxWorks-based targets now support the option `vxworks-timeout`. This option is set by the user, and *args* represents the number of seconds GDB waits for responses to `rpc`'s. You might use this if your VxWorks target is a slow software simulator or is on the far side of a thin network line.

The following information on connecting to VxWorks was current when this manual was produced; newer releases of VxWorks may use revised procedures.

To use GDB with VxWorks, you must rebuild your VxWorks kernel to include the remote debugging interface routines in the VxWorks library '`rdb.a`'. To do this, define `INCLUDE_RDB` in the VxWorks configuration file '`configAll.h`' and rebuild your VxWorks kernel. The resulting kernel contains '`rdb.a`', and spawns the source debugging task `tRdbTask` when VxWorks is booted. For more information on configuring and remaking VxWorks, see the manufacturer's manual.

Once you have included ‘`rdb.a`’ in your VxWorks system image and set your Unix execution search path to find GDB, you are ready to run GDB. From your Unix host, run `gdb` (or `vxgdb`, depending on your installation).

GDB comes up showing the prompt:

```
(vxgdb)
```

### 12.4.2.1 Connecting to VxWorks

The GDB command `target` lets you connect to a VxWorks target on the network. To connect to a target whose host name is “`tt`”, type:

```
(vxgdb) target vxworks tt
```

GDB displays messages like these:

```
Attaching remote machine across net...
Connected to tt.
```

GDB then attempts to read the symbol tables of any object modules loaded into the VxWorks target since it was last booted. GDB locates these files by searching the directories listed in the command search path (see [\[Your program’s environment\]](#), page [\[undefined\]](#)); if it fails to find an object file, it displays a message such as:

```
prog.o: No such file or directory.
```

When this happens, add the appropriate directory to the search path with the GDB command `path`, and execute the `target` command again.

### 12.4.2.2 VxWorks download

If you have connected to the VxWorks target and you want to debug an object that has not yet been loaded, you can use the GDB `load` command to download a file from Unix to VxWorks incrementally. The object file given as an argument to the `load` command is actually opened twice: first by the VxWorks target in order to download the code, then by GDB in order to read the symbol table. This can lead to problems if the current working directories on the two systems differ. If both systems have NFS mounted the same filesystems, you can avoid these problems by using absolute paths. Otherwise, it is simplest to set the working directory on both systems to the directory in which the object file resides, and then to reference the file by its name, without any path. For instance, a program ‘`prog.o`’ may reside in ‘`vxpath/vw/demo/rdb`’ in VxWorks and in ‘`hostpath/vw/demo/rdb`’ on the host. To load this program, type this on VxWorks:

```
-> cd "vxpath/vw/demo/rdb"
```

v Then, in GDB, type:

```
(vxgdb) cd hostpath/vw/demo/rdb
(vxgdb) load prog.o
```

GDB displays a response similar to this:

```
Reading symbol data from wherever/vw/demo/rdb/prog.o... done.
```

You can also use the `load` command to reload an object module after editing and recompiling the corresponding source file. Note that this makes GDB delete all currently-defined breakpoints, auto-displays, and convenience variables, and to clear the value history. (This is necessary in order to preserve the integrity of debugger data structures that reference the target system’s symbol table.)

### 12.4.2.3 Running tasks

You can also attach to an existing task using the `attach` command as follows:

```
(vxgdb) attach task
```

where `task` is the VxWorks hexadecimal task ID. The task can be running or suspended when you attach to it. Running tasks are suspended at the time of attachment.

### 12.4.3 GDB and Sparclet

GDB enables developers to debug tasks running on Sparclet targets from a Unix host. GDB uses code that runs on both the Unix host and on the Sparclet target. The program `gdb` is installed and executed on the Unix host.

`timeout args`

GDB now supports the option `remotetimeout`. This option is set by the user, and `args` represents the number of seconds GDB waits for responses.

When compiling for debugging, include the options `"-g"` to get debug information and `"-Ttext"` to relocate the program to where you wish to load it on the target. You may also want to add the options `"-n"` or `"-N"` in order to reduce the size of the sections.

```
sparclet-aout-gcc prog.c -Ttext 0x12010000 -g -o prog -N
```

You can use `objdump` to verify that the addresses are what you intended.

```
sparclet-aout-objdump --headers --syms prog
```

Once you have set your Unix execution search path to find GDB, you are ready to run GDB. From your Unix host, run `gdb` (or `sparclet-aout-gdb`, depending on your installation).

GDB comes up showing the prompt:

```
(gdb) (gdbslet)
```

#### 12.4.3.1 Setting file to debug

The GDB command `file` lets you choose with program to debug.

```
(gdbslet) file prog
```

GDB then attempts to read the symbol table of `'prog'`. GDB locates the file by searching the directories listed in the command search path. If the file was compiled with debug information (option `"-g"`), source files will be searched as well. GDB locates the source files by searching the directories listed in the directory search path (see [\[Your program's environment\]](#), page [\[undefined\]](#)). If it fails to find a file, it displays a message such as:

```
prog: No such file or directory.
```

When this happens, add the appropriate directories to the search paths with the GDB commands `path` and `dir`, and execute the `target` command again.



### 12.4.3.2 Connecting to Sparclet

The GDB command `target` lets you connect to a Sparclet target. To connect to a target on serial port “ttya”, type:

```
(gdb) target sparclet /dev/ttya
Remote target sparclet connected to /dev/ttya
main () at ../prog.c:3
```

GDB displays messages like these:

```
Connected to ttya.
```

### 12.4.3.3 Sparclet download

Once connected to the Sparclet target, you can use the GDB `load` command to download the file from the host to the target. The file name and load offset should be given as arguments to the `load` command. Since the file format is aout, the program must be loaded to the starting address. You can use `objdump` to find out what this value is. The load offset is an offset which is added to the VMA (virtual memory address) of each of the file’s sections. For instance, if the program ‘prog’ was linked to text address 0x1201000, with data at 0x12010160 and bss at 0x12010170, in GDB, type:

```
(gdb) load prog 0x12010000
Loading section .text, size 0xdb0 vma 0x12010000
```

If the code is loaded at a different address than what the program was linked to, you may need to use the `section` and `add-symbol-file` commands to tell GDB where to map the symbol table.

### 12.4.3.4 Running and debugging

You can now begin debugging the task using GDB’s execution control commands, `b`, `step`, `run`, etc. See the GDB manual for the list of commands.

```
(gdb) b main
Breakpoint 1 at 0x12010000: file prog.c, line 3.
(gdb) run
Starting program: prog
Breakpoint 1, main (argc=1, argv=0xeffff21c) at prog.c:3
3     char *symarg = 0;
(gdb) step
4     char *execarg = "hello!";
(gdb)
```

### 12.4.4 Simulated CPU target

For some configurations, GDB includes a CPU simulator that you can use instead of a hardware CPU to debug your programs. Currently, a simulator is available when GDB is configured to debug Zilog Z8000 or Hitachi microprocessor targets.

```
target sim
```

```
Debug programs on a simulated CPU (which CPU depends on the GDB configuration)
```

After specifying this target, you can debug programs for the simulated CPU in the same style as programs for your host computer; use the `file` command to load a new program image, the `run` command to run your program, and so on.

As well as making available all the usual machine registers (see `info reg`), this debugging target provides three additional items of information as specially named registers:

<code>cycles</code>	Counts clock-ticks in the simulator.
<code>insts</code>	Counts instructions run in the simulator.
<code>time</code>	Execution time in 60ths of a second.

You can refer to these values in GDB expressions with the usual conventions; for example, `'b fputc if $cycles>5000'` sets a conditional breakpoint that suspends only after at least 5000 simulated clock ticks.

## 13 Controlling GDB

You can alter the way GDB interacts with you by using the `set` command. For commands controlling how GDB displays data, see [Section 7.7 \[Print settings\], page 48](#); other settings are described here.

### 13.1 Prompt

GDB indicates its readiness to read a command by printing a string called the *prompt*. This string is normally `( gdb)`. You can change the prompt string with the `set prompt` command. For instance, when debugging GDB with GDB, it is useful to change the prompt in one of the GDB sessions so that you can always tell which one you are talking to.

*Note:* `set prompt` no longer adds a space for you after the prompt you set. This allows you to set a prompt which ends in a space or a prompt that does not.

```
set prompt newprompt
    Directs GDB to use newprompt as its prompt string henceforth.

show prompt
    Prints a line of the form: 'Gdb's prompt is: your-prompt'
```

### 13.2 Command editing

GDB reads its input commands via the *readline* interface. This GNU library provides consistent behavior for programs which provide a command line interface to the user. Advantages are GNU Emacs-style or *vi*-style inline editing of commands, *cs*h-like history substitution, and a storage and recall of command history across debugging sessions.

You may control the behavior of command line editing in GDB with the command `set`.

```
set editing
set editing on
    Enable command line editing (enabled by default).

set editing off
    Disable command line editing.

show editing
    Show whether command line editing is enabled.
```

### 13.3 Command history

GDB can keep track of the commands you type during your debugging sessions, so that you can be certain of precisely what happened. Use these commands to manage the GDB command history facility.

```
set history filename fname
    Set the name of the GDB command history file to fname. This is the file where GDB reads an initial command history list, and where it writes the command
```

history from this session when it exits. You can access this list through history expansion or through the history command editing characters listed below. This file defaults to the value of the environment variable `GDBHISTFILE`, or to `./.gdb_history` if this variable is not set.

`set history save`

`set history save on`

Record command history in a file, whose name may be specified with the `set history filename` command. By default, this option is disabled.

`set history save off`

Stop recording command history in a file.

`set history size size`

Set the number of commands which GDB keeps in its history list. This defaults to the value of the environment variable `HISTSIZE`, or to 256 if this variable is not set.

History expansion assigns special meaning to the character `!`.

Since `!` is also the logical not operator in C, history expansion is off by default. If you decide to enable history expansion with the `set history expansion on` command, you may sometimes need to follow `!` (when it is used as logical not, in an expression) with a space or a tab to prevent it from being expanded. The readline history facilities do not attempt substitution on the strings `!=` and `!()`, even when history expansion is enabled.

The commands to control history expansion are:

`set history expansion on`

`set history expansion`

Enable history expansion. History expansion is off by default.

`set history expansion off`

Disable history expansion.

The readline code comes with more complete documentation of editing and history expansion features. Users unfamiliar with GNU Emacs or vi may wish to read it.

`show history`

`show history filename`

`show history save`

`show history size`

`show history expansion`

These commands display the state of the GDB history parameters. `show history` by itself displays all four states.

`show commands`

Display the last ten commands in the command history.

`show commands n`

Print ten commands centered on command number *n*.

`show commands +`

Print ten commands just after the commands last printed.

## 13.4 Screen size

Certain commands to GDB may produce large amounts of information output to the screen. To help you read all of it, GDB pauses and asks you for input at the end of each page of output. Type `(RET)` when you want to continue the output, or `q` to discard the remaining output. Also, the screen width setting determines when to wrap lines of output. Depending on what is being printed, GDB tries to break the line at a readable place, rather than simply letting it overflow onto the following line.

Normally GDB knows the size of the screen from the termcap data base together with the value of the `TERM` environment variable and the `stty rows` and `stty cols` settings. If this is not correct, you can override it with the `set height` and `set width` commands:

```
set height lpp
show height
set width cpl
show width
```

These `set` commands specify a screen height of *lpp* lines and a screen width of *cpl* characters. The associated `show` commands display the current settings.

If you specify a height of zero lines, GDB does not pause during output no matter how long the output is. This is useful if output is to a file or to an editor buffer.

Likewise, you can specify `'set width 0'` to prevent GDB from wrapping its output.

## 13.5 Numbers

You can always enter numbers in octal, decimal, or hexadecimal in GDB by the usual conventions: octal numbers begin with `'0'`, decimal numbers end with `'.'`, and hexadecimal numbers begin with `'0x'`. Numbers that begin with none of these are, by default, entered in base 10; likewise, the default display for numbers—when no particular format is specified—is base 10. You can change the default base for both input and output with the `set radix` command.

```
set input-radix base
```

Set the default base for numeric input. Supported choices for *base* are decimal 8, 10, or 16. *base* must itself be specified either unambiguously or using the current default radix; for example, any of

```
set radix 012
set radix 10.
set radix 0xa
```

sets the base to decimal. On the other hand, `'set radix 10'` leaves the radix unchanged no matter what it was.

```
set output-radix base
```

Set the default base for numeric display. Supported choices for *base* are decimal 8, 10, or 16. *base* must itself be specified either unambiguously or using the current default radix.

`show input-radix`

Display the current default base for numeric input.

`show output-radix`

Display the current default base for numeric display.

## 13.6 Optional warnings and messages

By default, GDB is silent about its inner workings. If you are running on a slow machine, you may want to use the `set verbose` command. This makes GDB tell you when it does a lengthy internal operation, so you will not think it has crashed.

Currently, the messages controlled by `set verbose` are those which announce that the symbol table for a source file is being read; see `symbol-file` in [Section 11.1 \[Commands to specify files\]](#), page 79.

`set verbose on`

Enables GDB output of certain informational messages.

`set verbose off`

Disables GDB output of certain informational messages.

`show verbose`

Displays whether `set verbose` is on or off.

By default, if GDB encounters bugs in the symbol table of an object file, it is silent; but if you are debugging a compiler, you may find this information useful (see [Section 11.2 \[Errors reading symbol files\]](#), page 81).

`set complaints limit`

Permits GDB to output *limit* complaints about each type of unusual symbols before becoming silent about the problem. Set *limit* to zero to suppress all complaints; set it to a large number to prevent complaints from being suppressed.

`show complaints`

Displays how many symbol complaints GDB is permitted to produce.

By default, GDB is cautious, and asks what sometimes seems to be a lot of stupid questions to confirm certain commands. For example, if you try to run a program which is already running:

```
( gdb) run
```

```
The program being debugged has been started already.
```

```
Start it from the beginning? (y or n)
```

If you are willing to unflinchingly face the consequences of your own commands, you can disable this “feature”:

`set confirm off`

Disables confirmation requests.

`set confirm on`

Enables confirmation requests (the default).

`show confirm`

Displays state of confirmation requests.

## 14 Canned Sequences of Commands

Aside from breakpoint commands (see [Section 4.1.7 \[Breakpoint command lists\]](#), [page 25](#)), GDB provides two ways to store sequences of commands for execution as a unit: user-defined commands and command files.

### 14.1 User-defined commands

A *user-defined command* is a sequence of GDB commands to which you assign a new name as a command. This is done with the `define` command. User commands may accept up to 10 arguments separated by whitespace. Arguments are accessed within the user command via `$arg0...$arg9`. A trivial example:

```
define adder
  print $arg0 + $arg1 + $arg2
```

To execute the command use:

```
adder 1 2 3
```

This defines the command `adder`, which prints the sum of its three arguments. Note the arguments are text substitutions, so they may reference variables, use complex expressions, or even perform inferior functions calls.

#### `define commandname`

Define a command named *commandname*. If there is already a command by that name, you are asked to confirm that you want to redefine it.

The definition of the command is made up of other GDB command lines, which are given following the `define` command. The end of these commands is marked by a line containing `end`.

**if** Takes a single argument, which is an expression to evaluate. It is followed by a series of commands that are executed only if the expression is true (nonzero). There can then optionally be a line `else`, followed by a series of commands that are only executed if the expression was false. The end of the list is marked by a line containing `end`.

**while** The syntax is similar to `if`: the command takes a single argument, which is an expression to evaluate, and must be followed by the commands to execute, one per line, terminated by an `end`. The commands are executed repeatedly as long as the expression evaluates to true.

#### `document commandname`

Document the user-defined command *commandname*, so that it can be accessed by `help`. The command *commandname* must already be defined. This command reads lines of documentation just as `define` reads the lines of the command definition, ending with `end`. After the `document` command is finished, `help` on command *commandname* displays the documentation you have written.

You may use the `document` command again to change the documentation of a command. Redefining the command with `define` does not change the documentation.

**help user-defined**

List all user-defined commands, with the first line of the documentation (if any) for each.

**show user**

**show user** *commandname*

Display the GDB commands used to define *commandname* (but not its documentation). If no *commandname* is given, display the definitions for all user-defined commands.

When user-defined commands are executed, the commands of the definition are not printed. An error in any command stops execution of the user-defined command.

If used interactively, commands that would ask for confirmation proceed without asking when used inside a user-defined command. Many GDB commands that normally print messages to say what they are doing omit the messages when used in a user-defined command.

## 14.2 User-defined command hooks

You may define *hooks*, which are a special kind of user-defined command. Whenever you run the command ‘foo’, if the user-defined command ‘hook-foo’ exists, it is executed (with no arguments) before that command.

In addition, a pseudo-command, ‘stop’ exists. Defining (‘hook-stop’) makes the associated commands execute every time execution stops in your program: before breakpoint commands are run, displays are printed, or the stack frame is printed.

You can define a hook for any single-word command in GDB, but not for command aliases; you should define a hook for the basic command name, e.g. `backtrace` rather than `bt`. If an error occurs during the execution of your hook, execution of GDB commands stops and GDB issues a prompt (before the command that you actually typed had a chance to run).

If you try to define a hook which does not match any known command, you get a warning from the `define` command.

## 14.3 Command files

A command file for GDB is a file of lines that are GDB commands. Comments (lines starting with `#`) may also be included. An empty line in a command file does nothing; it does not mean to repeat the last command, as it would from the terminal.

When you start GDB, it automatically executes commands from its *init files*. These are files named ‘`.gdbinit`’. GDB reads the init file (if any) in your home directory, then processes command line options and operands, and then reads the init file (if any) in the current working directory. This is so the init file in your home directory can set options (such as `set complaints`) which affect the processing of the command line options and operands. The init files are not executed if you use the ‘`-nx`’ option; see [Section 1.1.2 \[Choosing modes\], page 6](#).

On some configurations of GDB, the init file is known by a different name (these are typically environments where a specialized form of GDB may need to coexist with other forms,



hence a different name for the specialized version's init file). These are the environments with special init file names:

- VxWorks (Wind River Systems real-time OS): `‘.vxgdbinit’`
- OS68K (Enea Data Systems real-time OS): `‘.os68gdbinit’`
- ES-1800 (Ericsson Telecom AB M68000 emulator): `‘.esgdbinit’`

You can also request the execution of a command file with the `source` command:

`source filename`

Execute the command file *filename*.

The lines in a command file are executed sequentially. They are not printed as they are executed. An error in any command terminates execution of the command file.

Commands that would ask for confirmation if used interactively proceed without asking when used in a command file. Many GDB commands that normally print messages to say what they are doing omit the messages when called from command files.

## 14.4 Commands for controlled output

During the execution of a command file or a user-defined command, normal GDB output is suppressed; the only output that appears is what is explicitly printed by the commands in the definition. This section describes three commands useful for generating exactly the output you want.

`echo text` Print *text*. Nonprinting characters can be included in *text* using C escape sequences, such as `‘\n’` to print a newline. **No newline is printed unless you specify one.** In addition to the standard C escape sequences, a backslash followed by a space stands for a space. This is useful for displaying a string with spaces at the beginning or the end, since leading and trailing spaces are otherwise trimmed from all arguments. To print `‘ and foo = ’`, use the command `‘echo \ and foo = \ ’`.

A backslash at the end of *text* can be used, as in C, to continue the command onto subsequent lines. For example,

```
echo This is some text\n\
which is continued\n\
onto several lines.\n
```

produces the same output as

```
echo This is some text\n
echo which is continued\n
echo onto several lines.\n
```

`output expression`

Print the value of *expression* and nothing but that value: no newlines, no `‘$nn = ’`. The value is not entered in the value history either. See [Section 7.1 \[Expressions\]](#), page 41, for more information on expressions.

`output/fmt expression`

Print the value of *expression* in format *fmt*. You can use the same formats as for `print`. See [Section 7.4 \[Output formats\]](#), page 44, for more information.

`printf` *string*, *expressions...*

Print the values of the *expressions* under the control of *string*. The *expressions* are separated by commas and may be either numbers or pointers. Their values are printed as specified by *string*, exactly as if your program were to execute the C subroutine

```
printf (string, expressions...);
```

For example, you can print two values in hex like this:

```
printf "foo, bar-foo = 0x%x, 0x%x\n", foo, bar-foo
```

The only backslash-escape sequences that you can use in the format string are the simple ones that consist of backslash followed by a letter.

## 15 Using GDB under Emacs

A special interface allows you to use GNU Emacs to view (and edit) the source files for the program you are debugging with GDB.

To use this interface, use the command `M-x gdb` in Emacs. Give the executable file you want to debug as an argument. This command starts GDB as a subprocess of Emacs, with input and output through a newly created Emacs buffer.

Using GDB under Emacs is just like using GDB normally except for two things:

- All “terminal” input and output goes through the Emacs buffer.

This applies both to GDB commands and their output, and to the input and output done by the program you are debugging.

This is useful because it means that you can copy the text of previous commands and input them again; you can even use parts of the output in this way.

All the facilities of Emacs’ Shell mode are available for interacting with your program. In particular, you can send signals the usual way—for example, `C-c C-c` for an interrupt, `C-c C-z` for a stop.

- GDB displays source code through Emacs.

Each time GDB displays a stack frame, Emacs automatically finds the source file for that frame and puts an arrow (`=>`) at the left margin of the current line. Emacs uses a separate buffer for source display, and splits the screen to show both your GDB session and the source.

Explicit GDB `list` or search commands still produce output as usual, but you probably have no reason to use them from Emacs.

*Warning:* If the directory where your program resides is not your current directory, it can be easy to confuse Emacs about the location of the source files, in which case the auxiliary display buffer does not appear to show your source. GDB can find programs by searching your environment’s `PATH` variable, so the GDB input and output session proceeds normally; but Emacs does not get enough information back from GDB to locate the source files in this situation. To avoid this problem, either start GDB mode from the directory where your program resides, or specify an absolute file name when prompted for the `M-x gdb` argument.

A similar confusion can result if you use the GDB `file` command to switch to debugging a program in some other location, from an existing GDB buffer in Emacs.

By default, `M-x gdb` calls the program called `gdb`. If you need to call GDB by a different name (for example, if you keep several configurations around, with different names) you can set the Emacs variable `gdb-command-name`; for example,

```
(setq gdb-command-name "mygdb")
```

(preceded by `ESC ESC`, or typed in the `*scratch*` buffer, or in your `.emacs` file) makes Emacs call the program named `“mygdb”` instead.

In the GDB I/O buffer, you can use these special Emacs commands in addition to the standard Shell mode commands:

<i>C-h m</i>	Describe the features of Emacs' GDB Mode.
<i>M-s</i>	Execute to another source line, like the GDB <code>step</code> command; also update the display window to show the current file and location.
<i>M-n</i>	Execute to next source line in this function, skipping all function calls, like the GDB <code>next</code> command. Then update the display window to show the current file and location.
<i>M-i</i>	Execute one instruction, like the GDB <code>stepi</code> command; update display window accordingly.
<i>M-x gdb-nexti</i>	Execute to next instruction, using the GDB <code>nexti</code> command; update display window accordingly.
<i>C-c C-f</i>	Execute until exit from the selected stack frame, like the GDB <code>finish</code> command.
<i>M-c</i>	Continue execution of your program, like the GDB <code>continue</code> command. <i>Warning:</i> In Emacs v19, this command is <i>C-c C-p</i> .
<i>M-u</i>	Go up the number of frames indicated by the numeric argument (see <a href="#">section "Numeric Arguments" in <i>The GNU Emacs Manual</i></a> ), like the GDB <code>up</code> command. <i>Warning:</i> In Emacs v19, this command is <i>C-c C-u</i> .
<i>M-d</i>	Go down the number of frames indicated by the numeric argument, like the GDB <code>down</code> command. <i>Warning:</i> In Emacs v19, this command is <i>C-c C-d</i> .
<i>C-x &amp;</i>	Read the number where the cursor is positioned, and insert it at the end of the GDB I/O buffer. For example, if you wish to disassemble code around an address that was displayed earlier, type <code>disassemble</code> ; then move the cursor to the address display, and pick up the argument for <code>disassemble</code> by typing <i>C-x &amp;</i> .  You can customize this further by defining elements of the list <code>gdb-print-command</code> ; once it is defined, you can format or otherwise process numbers picked up by <i>C-x &amp;</i> before they are inserted. A numeric argument to <i>C-x &amp;</i> indicates that you wish special formatting, and also acts as an index to pick an element of the list. If the list element is a string, the number to be inserted is formatted using the Emacs function <code>format</code> ; otherwise the number is passed as an argument to the corresponding list element.

In any source file, the Emacs command *C-x SPC* (`gdb-break`) tells GDB to set a breakpoint on the source line point is on.

If you accidentally delete the source-display buffer, an easy way to get it back is to type the command `f` in the GDB buffer, to request a frame display; when you run under Emacs, this recreates the source buffer if necessary to show you the context of the current frame.

The source files displayed in Emacs are in ordinary Emacs buffers which are visiting the source files in the usual way. You can edit the files with these buffers if you wish; but keep

in mind that GDB communicates with Emacs in terms of line numbers. If you add or delete lines from the text, the line numbers that GDB knows cease to correspond properly with the code.



## 16 Reporting Bugs in GDB

Your bug reports play an essential role in making GDB reliable.

Reporting a bug may help you by bringing a solution to your problem, or it may not. But in any case the principal function of a bug report is to help the entire community by making the next version of GDB work better. Bug reports are your contribution to the maintenance of GDB.

In order for a bug report to serve its purpose, you must include the information that enables us to fix the bug.

### 16.1 Have you found a bug?

If you are not sure whether you have found a bug, here are some guidelines:

- If the debugger gets a fatal signal, for any input whatever, that is a GDB bug. Reliable debuggers never crash.
- If GDB produces an error message for valid input, that is a bug.
- If GDB does not produce an error message for invalid input, that is a bug. However, you should note that your idea of “invalid input” might be our idea of “an extension” or “support for traditional practice”.
- If you are an experienced user of debugging tools, your suggestions for improvement of GDB are welcome in any case.

### 16.2 How to report bugs

A number of companies and individuals offer support for GNU products. If you obtained GDB from a support organization, we recommend you contact that organization first.

You can find contact information for many support companies and individuals in the file ‘etc/SERVICE’ in the GNU Emacs distribution.

In any event, we also recommend that you send bug reports for GDB to one of these addresses:

```
bug-gdb@prep.ai.mit.edu  
{ucbvax|mit-eddie|uunet}!prep.ai.mit.edu!bug-gdb
```

**Do not send bug reports to ‘info-gdb’, or to ‘help-gdb’, or to any newsgroups.** Most users of GDB do not want to receive bug reports. Those that do have arranged to receive ‘bug-gdb’.

The mailing list ‘bug-gdb’ has a newsgroup ‘gnu.gdb.bug’ which serves as a repeater. The mailing list and the newsgroup carry exactly the same messages. Often people think of posting bug reports to the newsgroup instead of mailing them. This appears to work, but it has one problem which can be crucial: a newsgroup posting often lacks a mail path back to the sender. Thus, if we need to ask for more information, we may be unable to reach you. For this reason, it is better to send bug reports to the mailing list.

As a last resort, send bug reports on paper to:

GNU Debugger Bugs  
 Free Software Foundation Inc.  
 59 Temple Place - Suite 330  
 Boston, MA 02111-1307  
 USA

The fundamental principle of reporting bugs usefully is this: **report all the facts**. If you are not sure whether to state a fact or leave it out, state it!

Often people omit facts because they think they know what causes the problem and assume that some details do not matter. Thus, you might assume that the name of the variable you use in an example does not matter. Well, probably it does not, but one cannot be sure. Perhaps the bug is a stray memory reference which happens to fetch from the location where that name is stored in memory; perhaps, if the name were different, the contents of that location would fool the debugger into doing the right thing despite the bug. Play it safe and give a specific, complete example. That is the easiest thing for you to do, and the most helpful.

Keep in mind that the purpose of a bug report is to enable us to fix the bug if it is new to us. Therefore, always write your bug reports on the assumption that the bug has not been reported previously.

Sometimes people give a few sketchy facts and ask, “Does this ring a bell?” Those bug reports are useless, and we urge everyone to *refuse to respond to them* except to chide the sender to report bugs properly.

To enable us to fix the bug, you should include all these things:

- The version of GDB. GDB announces it if you start with no arguments; you can also print it at any time using `show version`.

Without this, we will not know whether there is any point in looking for the bug in the current version of GDB.

- The type of machine you are using, and the operating system name and version number.
- What compiler (and its version) was used to compile GDB—e.g. “gcc-2.0”.
- What compiler (and its version) was used to compile the program you are debugging—e.g. “gcc-2.0”.
- The command arguments you gave the compiler to compile your example and observe the bug. For example, did you use ‘-O’? To guarantee you will not omit something important, list them all. A copy of the Makefile (or the output from make) is sufficient. If we were to try to guess the arguments, we would probably guess wrong and then we might not encounter the bug.
- A complete input script, and all necessary source files, that will reproduce the bug.
- A description of what behavior you observe that you believe is incorrect. For example, “It gets a fatal signal.”

Of course, if the bug is that GDB gets a fatal signal, then we will certainly notice it. But if the bug is incorrect output, we might not notice unless it is glaringly wrong. You might as well not give us a chance to make a mistake.

Even if the problem you experience is a fatal signal, you should still say so explicitly. Suppose something strange is going on, such as, your copy of GDB is out of synch, or



you have encountered a bug in the C library on your system. (This has happened!) Your copy might crash and ours would not. If you told us to expect a crash, then when ours fails to crash, we would know that the bug was not happening for us. If you had not told us to expect a crash, then we would not be able to draw any conclusion from our observations.

- If you wish to suggest changes to the GDB source, send us context diffs. If you even discuss something in the GDB source, refer to it by context, not by line number. The line numbers in our development sources will not match those in your sources. Your line numbers would convey no useful information to us.

Here are some things that are not necessary:

- A description of the envelope of the bug.

Often people who encounter a bug spend a lot of time investigating which changes to the input file will make the bug go away and which changes will not affect it. This is often time consuming and not very useful, because the way we will find the bug is by running a single example under the debugger with breakpoints, not by pure deduction from a series of examples. We recommend that you save your time for something else.

Of course, if you can find a simpler example to report *instead* of the original one, that is a convenience for us. Errors in the output will be easier to spot, running under the debugger will take less time, and so on.

However, simplification is not vital; if you do not want to do this, report the bug anyway and send us the entire test case you used.
- A patch for the bug.

A patch for the bug does help us if it is a good one. But do not omit the necessary information, such as the test case, on the assumption that a patch is all we need. We might see problems with your patch and decide to fix the problem another way, or we might not understand it at all.

Sometimes with a program as complicated as GDB it is very hard to construct an example that will make the program follow a certain path through the code. If you do not send us the example, we will not be able to construct one, so we will not be able to verify that the bug is fixed.

And if we cannot understand what bug you are trying to fix, or why your patch should be an improvement, we will not install it. A test case will help us to understand.
- A guess about what the bug is or what it depends on.

Such guesses are usually wrong. Even we cannot guess right about such things without first using the debugger to find the facts.



## Appendix A Command Line Editing

This text describes GNU's command line editing interface.

### A.1 Introduction to Line Editing

The following paragraphs describe the notation we use to represent keystrokes.

The text `C-k` is read as 'Control-K' and describes the character produced when the Control key is depressed and the `k` key is struck.

The text `M-k` is read as 'Meta-K' and describes the character produced when the meta key (if you have one) is depressed, and the `k` key is struck. If you do not have a meta key, the identical keystroke can be generated by typing `ESC` *first*, and then typing `k`. Either process is known as *metafying* the `k` key.

The text `M-C-k` is read as 'Meta-Control-k' and describes the character produced by *metafying* `C-k`.

In addition, several keys have their own names. Specifically, `DEL`, `ESC`, `LFD`, `SPC`, `RET`, and `TAB` all stand for themselves when seen in this text, or in an init file (see [Section A.3 \[Readline Init File\], page 115](#), for more info).

### A.2 Readline Interaction

Often during an interactive session you type in a long line of text, only to notice that the first word on the line is misspelled. The Readline library gives you a set of commands for manipulating the text as you type it in, allowing you to just fix your typo, and not forcing you to retype the majority of the line. Using these editing commands, you move the cursor to the place that needs correction, and delete or insert the text of the corrections. Then, when you are satisfied with the line, you simply press `RET`. You do not have to be at the end of the line to press `RET`; the entire line is accepted regardless of the location of the cursor within the line.

#### A.2.1 Readline Bare Essentials

In order to enter characters into the line, simply type them. The typed character appears where the cursor was, and then the cursor moves one space to the right. If you mistype a character, you can use `DEL` to back up, and delete the mistyped character.

Sometimes you may miss typing a character that you wanted to type, and not notice your error until you have typed several other characters. In that case, you can type `C-b` to move the cursor to the left, and then correct your mistake. Afterwards, you can move the cursor to the right with `C-f`.

When you add text in the middle of a line, you will notice that characters to the right of the cursor get 'pushed over' to make room for the text that you have inserted. Likewise, when you delete text behind the cursor, characters to the right of the cursor get 'pulled back' to fill in the blank space created by the removal of the text. A list of the basic bare essentials for editing the text of an input line follows.

`C-b`      Move back one character.

- `<C-f>` Move forward one character.
- `<DEL>` Delete the character to the left of the cursor.
- `<C-d>` Delete the character underneath the cursor.

#### Printing characters

- Insert itself into the line at the cursor.
- `<C-_>` Undo the last thing that you did. You can undo all the way back to an empty line.

### A.2.2 Readline Movement Commands

The above table describes the most basic possible keystrokes that you need in order to do editing of the input line. For your convenience, many other commands have been added in addition to `<C-b>`, `<C-f>`, `<C-d>`, and `<DEL>`. Here are some commands for moving more rapidly about the line.

- `<C-a>` Move to the start of the line.
- `<C-e>` Move to the end of the line.
- `<M-f>` Move forward a word.
- `<M-b>` Move backward a word.
- `<C-l>` Clear the screen, reprinting the current line at the top.

Notice how `<C-f>` moves forward a character, while `<M-f>` moves forward a word. It is a loose convention that control keystrokes operate on characters while meta keystrokes operate on words.

### A.2.3 Readline Killing Commands

*Killing* text means to delete the text from the line, but to save it away for later use, usually by *yanking* it back into the line. If the description for a command says that it ‘kills’ text, then you can be sure that you can get the text back in a different (or the same) place later.

Here is the list of commands for killing text.

- `<C-k>` Kill the text from the current cursor position to the end of the line.
- `<M-d>` Kill from the cursor to the end of the current word, or if between words, to the end of the next word.
- `<M-DEL>` Kill from the cursor to the start of the previous word, or if between words, to the start of the previous word.
- `<C-w>` Kill from the cursor to the previous whitespace. This is different than `<M-DEL>` because the word boundaries differ.

And, here is how to *yank* the text back into the line.

- `<C-y>` Yank the most recently killed text back into the buffer at the cursor.

**(M-y)** Rotate the kill-ring, and yank the new top. You can only do this if the prior command is **(C-y)** or **(M-y)**.

When you use a kill command, the text is saved in a *kill-ring*. Any number of consecutive kills save all of the killed text together, so that when you yank it back, you get it in one clean sweep. The kill ring is not line specific; the text that you killed on a previously typed line is available to be yanked back later, when you are typing another line.

## A.2.4 Readline Arguments

You can pass numeric arguments to Readline commands. Sometimes the argument acts as a repeat count, other times it is the *sign* of the argument that is significant. If you pass a negative argument to a command which normally acts in a forward direction, that command will act in a backward direction. For example, to kill text back to the start of the line, you might type **(M-) (C-k)**.

The general way to pass numeric arguments to a command is to type meta digits before the command. If the first ‘digit’ you type is a minus sign (**(-)**), then the sign of the argument will be negative. Once you have typed one meta digit to get the argument started, you can type the remainder of the digits, and then the command. For example, to give the **(C-d)** command an argument of 10, you could type **(M-1 0 C-d)**.

## A.3 Readline Init File

Although the Readline library comes with a set of GNU Emacs-like keybindings, it is possible that you would like to use a different set of keybindings. You can customize programs that use Readline by putting commands in an *init* file in your home directory. The name of this file is ‘`~/inputrc`’.

When a program which uses the Readline library starts up, the ‘`~/inputrc`’ file is read, and the keybindings are set.

In addition, the **(C-x C-r)** command re-reads this init file, thus incorporating any changes that you might have made to it.

### A.3.1 Readline Init Syntax

There are only four constructs allowed in the ‘`~/inputrc`’ file:

#### Variable Settings

You can change the state of a few variables in Readline. You do this by using the **set** command within the init file. Here is how you would specify that you wish to use **vi** line editing commands:

```
set editing-mode vi
```

Right now, there are only a few variables which can be set; so few in fact, that we just iterate them here:

```
editing-mode
```

The **editing-mode** variable controls which editing mode you are using. By default, GNU Readline starts up in Emacs editing mode,

where the keystrokes are most similar to Emacs. This variable can either be set to `emacs` or `vi`.

#### `horizontal-scroll-mode`

This variable can either be set to `On` or `Off`. Setting it to `On` means that the text of the lines that you edit will scroll horizontally on a single screen line when they are larger than the width of the screen, instead of wrapping onto a new screen line. By default, this variable is set to `Off`.

#### `mark-modified-lines`

This variable when set to `On`, says to display an asterisk (`*`) at the starts of history lines which have been modified. This variable is off by default.

#### `prefer-visible-bell`

If this variable is set to `On` it means to use a visible bell if one is available, rather than simply ringing the terminal bell. By default, the value is `Off`.

### Key Bindings

The syntax for controlling keybindings in the `~/inputrc` file is simple. First you have to know the *name* of the command that you want to change. The following pages contain tables of the command name, the default keybinding, and a short description of what the command does.

Once you know the name of the command, simply place the name of the key you wish to bind the command to, a colon, and then the name of the command on a line in the `~/inputrc` file. The name of the key can be expressed in different ways, depending on which is most comfortable for you.

*keyname*: *function-name* or *macro*

*keyname* is the name of a key spelled out in English. For example:

```
Control-u: universal-argument
Meta-Rubout: backward-kill-word
Control-o: ">&output"
```

In the above example, `C-u` is bound to the function `universal-argument`, and `C-o` is bound to run the macro expressed on the right hand side (that is, to insert the text `'>&output'` into the line).

*"keyseq"*: *function-name* or *macro*

*keyseq* differs from *keyname* above in that strings denoting an entire key sequence can be specified. Simply place the key sequence in double quotes. GNU Emacs style key escapes can be used, as in the following example:

```
"\C-u": universal-argument
"\C-x\C-r": re-read-init-file
"\e[11~": "Function Key 1"
```

In the above example, `C-u` is bound to the function `universal-argument` (just as it was in the first example), `C-x C-r` is bound

to the function `re-read-init-file`, and `(ESC [11~)` is bound to insert the text ‘Function Key 1’.

### A.3.1.1 Commands For Moving

`beginning-of-line` (`(C-a)`)

Move to the start of the current line.

`end-of-line` (`(C-e)`)

Move to the end of the line.

`forward-char` (`(C-f)`)

Move forward a character.

`backward-char` (`(C-b)`)

Move back a character.

`forward-word` (`(M-f)`)

Move forward to the end of the next word.

`backward-word` (`(M-b)`)

Move back to the start of this, or the previous, word.

`clear-screen` (`(C-l)`)

Clear the screen leaving the current line at the top of the screen.

### A.3.1.2 Commands For Manipulating The History

`accept-line` (Newline, Return)

Accept the line regardless of where the cursor is. If this line is non-empty, add it to the history list. If this line was a history line, then restore the history line to its original state.

`previous-history` (`(C-p)`)

Move ‘up’ through the history list.

`next-history` (`(C-n)`)

Move ‘down’ through the history list.

`beginning-of-history` (`(M- $\leftarrow$ )`)

Move to the first line in the history.

`end-of-history` (`(M- $\rightarrow$ )`)

Move to the end of the input history, i.e., the line you are entering.

`reverse-search-history` (`(C-r)`)

Search backward starting at the current line and moving ‘up’ through the history as necessary. This is an incremental search.

`forward-search-history` (`(C-s)`)

Search forward starting at the current line and moving ‘down’ through the the history as necessary.

### A.3.1.3 Commands For Changing Text

`delete-char` ( $\overline{\text{C-d}}$ )

Delete the character under the cursor. If the cursor is at the beginning of the line, and there are no characters in the line, and the last character typed was not  $\overline{\text{C-d}}$ , then return EOF.

`backward-delete-char` (Rubout)

Delete the character behind the cursor. A numeric argument says to kill the characters instead of deleting them.

`quoted-insert` ( $\overline{\text{C-q}}$ ,  $\overline{\text{C-v}}$ )

Add the next character that you type to the line verbatim. This is how to insert things like  $\overline{\text{C-q}}$  for example.

`tab-insert` ( $\overline{\text{M-TAB}}$ )

Insert a tab character.

`self-insert` (a, b, A, 1, !, ...)

Insert yourself.

`transpose-chars` ( $\overline{\text{C-t}}$ )

Drag the character before point forward over the character at point. Point moves forward as well. If point is at the end of the line, then transpose the two characters before point. Negative arguments don't work.

`transpose-words` ( $\overline{\text{M-t}}$ )

Drag the word behind the cursor past the word in front of the cursor moving the cursor over that word as well.

`upcase-word` ( $\overline{\text{M-u}}$ )

Uppercase all letters in the current (or following) word. With a negative argument, do the previous word, but do not move point.

`downcase-word` ( $\overline{\text{M-l}}$ )

Lowercase all letters in the current (or following) word. With a negative argument, do the previous word, but do not move point.

`capitalize-word` ( $\overline{\text{M-c}}$ )

Uppercase the first letter in the current (or following) word. With a negative argument, do the previous word, but do not move point.



### A.3.1.4 Killing And Yanking

kill-line ( $\overline{C-k}$ )

Kill the text from the current cursor position to the end of the line.

backward-kill-line ()

Kill backward to the beginning of the line. This is normally unbound.

kill-word ( $\overline{M-d}$ )

Kill from the cursor to the end of the current word, or if between words, to the end of the next word.

backward-kill-word ( $\overline{M-DEL}$ )

Kill the word behind the cursor.

unix-line-discard ( $\overline{C-u}$ )

Kill the whole line the way  $\overline{C-u}$  used to in Unix line input. The killed text is saved on the kill-ring.

unix-word-rubout ( $\overline{C-w}$ )

Kill the word the way  $\overline{C-w}$  used to in Unix line input. The killed text is saved on the kill-ring. This is different than backward-kill-word because the word boundaries differ.

yank ( $\overline{C-y}$ )

Yank the top of the kill ring into the buffer at point.

yank-pop ( $\overline{M-y}$ )

Rotate the kill-ring, and yank the new top. You can only do this if the prior command is yank or yank-pop.

### A.3.1.5 Specifying Numeric Arguments

digit-argument ( $\overline{M-0}$ ,  $\overline{M-1}$ , ...  $\overline{M-}$ )

Add this digit to the argument already accumulating, or start a new argument.  $\overline{M-}$  starts a negative argument.

universal-argument ()

Do what  $\overline{C-u}$  does in GNU Emacs. By default, this is not bound.

### A.3.1.6 Letting Readline Type For You

complete (TAB)

Attempt to do completion on the text before point. This is implementation defined. Generally, if you are typing a filename argument, you can do filename completion; if you are typing a command, you can do command completion, if you are typing in a symbol to GDB, you can do symbol name completion, if you are typing in a variable to Bash, you can do variable name completion.

possible-completions (**M-?**)

List the possible completions of the text before point.

### A.3.1.7 Some Miscellaneous Commands

re-read-init-file (**C-x C-r**)

Read in the contents of your ‘~/inputrc’ file, and incorporate any bindings found there.

abort (**C-g**)

Stop running the current editing command.

prefix-meta (**ESC**)

Make the next character that you type be metafied. This is for people without a meta key. Typing **ESC f** is equivalent to typing **M-f**.

undo (**C-)**)

Incremental undo, separately remembered for each line.

revert-line (**M-r**)

Undo all changes made to this line. This is like typing the ‘undo’ command enough times to get back to the beginning.

### A.3.2 Readline vi Mode

While the Readline library does not have a full set of vi editing functions, it does contain enough to allow simple editing of the line.

In order to switch interactively between GNU Emacs and vi editing modes, use the command **M-C-j** (toggle-editing-mode).

When you enter a line in vi mode, you are already placed in ‘insertion’ mode, as if you had typed an ‘i’. Pressing **ESC** switches you into ‘edit’ mode, where you can edit the text of the line with the standard vi movement keys, move to previous history lines with ‘k’, and following lines with ‘j’, and so forth.

## Appendix B Using History Interactively

This chapter describes how to use the GNU History Library interactively, from a user's standpoint.

### B.1 History Interaction

The History library provides a history expansion feature similar to the history expansion in `csh`. The following text describes the syntax you use to manipulate history information.

History expansion takes two parts. In the first part, determine which line from the previous history will be used for substitution. This line is called the *event*. In the second part, select portions of that line for inclusion into the current line. These portions are called *words*. GDB breaks the line into words in the same way that the Bash shell does, so that several English (or Unix) words surrounded by quotes are considered one word.

#### B.1.1 Event Designators

An *event designator* is a reference to a command line entry in the history list.

!	Start a history substitution, except when followed by a space, tab, or the end of the line... $\equiv$ or $\hat{\phantom{a}}$ .
!!	Refer to the previous command. This is a synonym for <code>!-1</code> .
!n	Refer to command line <i>n</i> .
!-n	Refer to the command line <i>n</i> lines back.
!string	Refer to the most recent command starting with <i>string</i> .
!?string[?]	Refer to the most recent command containing <i>string</i> .

#### B.1.2 Word Designators

A  $\hat{\phantom{a}}$  separates the event designator from the *word designator*. It can be omitted if the word designator begins with a  $\hat{\phantom{a}}$ ,  $\$$ ,  $\ast$  or  $\%$ . Words are numbered from the beginning of the line, with the first word being denoted by a 0 (zero).

0 (zero)	The zero'th word. For many applications, this is the command word.
n	The <i>n</i> 'th word.
~	The first argument. that is, word 1.
\$	The last argument.
%	The word matched by the most recent <code>?string?</code> search.
x-y	A range of words; <code>-y</code> Abbreviates <code>0-y</code> .
*	All of the words, excepting the zero'th. This is a synonym for <code>1-\$</code> . It is not an error to use $\ast$ if there is just one word in the event. The empty string is returned in that case.

### B.1.3 Modifiers

After the optional word designator, you can add a sequence of one or more of the following *modifiers*, each preceded by a  $\hat{}$ .

- # The entire command line typed so far. This means the current command, not the previous command.
- h Remove a trailing pathname component, leaving only the head.
- r Remove a trailing suffix of the form `.'suffix`, leaving the basename.
- e Remove all but the suffix.
- t Remove all leading pathname components, leaving the tail.
- p Print the new command but do not execute it.

## Appendix C Formatting Documentation

The GDB 4 release includes an already-formatted reference card, ready for printing with PostScript or Ghostscript, in the ‘gdb’ subdirectory of the main source directory<sup>1</sup>. If you can use PostScript or Ghostscript with your printer, you can print the reference card immediately with ‘refcard.ps’.

The release also includes the source for the reference card. You can format it, using T<sub>E</sub>X, by typing:

```
make refcard.dvi
```

The GDB reference card is designed to print in *landscape* mode on US “letter” size paper; that is, on a sheet 11 inches wide by 8.5 inches high. You will need to specify this form of printing as an option to your DVI output program.

All the documentation for GDB comes as part of the machine-readable distribution. The documentation is written in Texinfo format, which is a documentation system that uses a single source file to produce both on-line information and a printed manual. You can use one of the Info formatting commands to create the on-line version of the documentation and T<sub>E</sub>X (or `texi2roff`) to typeset the printed version.

GDB includes an already formatted copy of the on-line Info version of this manual in the ‘gdb’ subdirectory. The main Info file is ‘gdb-version-number/gdb/gdb.info’, and it refers to subordinate files matching ‘gdb.info\*’ in the same directory. If necessary, you can print out these files, or read them with any editor; but they are easier to read using the `info` subsystem in GNU Emacs or the standalone `info` program, available as part of the GNU Texinfo distribution.

If you want to format these Info files yourself, you need one of the Info formatting programs, such as `texinfo-format-buffer` or `makeinfo`.

If you have `makeinfo` installed, and are in the top level GDB source directory (‘gdb-4.17.gnat.3.11p’, in the case of version 4.17.gnat.3.11p), you can make the Info file by typing:

```
cd gdb
make gdb.info
```

If you want to typeset and print copies of this manual, you need T<sub>E</sub>X, a program to print its DVI output files, and ‘`texinfo.tex`’, the Texinfo definitions file.

T<sub>E</sub>X is a typesetting program; it does not print files directly, but produces output files called DVI files. To print a typeset document, you need a program to print DVI files. If your system has T<sub>E</sub>X installed, chances are it has such a program. The precise command to use depends on your system; `lpr -d` is common; another (for PostScript devices) is `dvips`. The DVI print command may require a file name without any extension or a ‘.dvi’ extension.

T<sub>E</sub>X also requires a macro definitions file called ‘`texinfo.tex`’. This file tells T<sub>E</sub>X how to typeset a document written in Texinfo format. On its own, T<sub>E</sub>X cannot either read or typeset a Texinfo file. ‘`texinfo.tex`’ is distributed with GDB and is located in the ‘gdb-version-number/texinfo’ directory.

---

<sup>1</sup> In ‘gdb-4.17.gnat.3.11p/gdb/refcard.ps’ of the version 4.17.gnat.3.11p release.

If you have T<sub>E</sub>X and a DVI printer program installed, you can typeset and print this manual. First switch to the the 'gdb' subdirectory of the main source directory (for example, to 'gdb- 4.17.gnat.3.11p/gdb') and then type:

```
make gdb.dvi
```

## Appendix D Installing GDB

GDB comes with a `configure` script that automates the process of preparing GDB for installation; you can then use `make` to build the `gdb` program.<sup>1</sup>

The GDB distribution includes all the source code you need for GDB in a single directory, whose name is usually composed by appending the version number to ‘`gdb`’.

For example, the GDB version 4.17.gnat.3.11p distribution is in the ‘`gdb- 4.17.gnat.3.11p`’ directory. That directory contains:

```
gdb- 4.17.gnat.3.11p/configure (and supporting files)
      script for configuring GDB and all its supporting libraries

gdb- 4.17.gnat.3.11p/gdb
      the source specific to GDB itself

gdb- 4.17.gnat.3.11p/bfd
      source for the Binary File Descriptor library

gdb- 4.17.gnat.3.11p/include
      GNU include files

gdb- 4.17.gnat.3.11p/libiberty
      source for the ‘-liberty’ free software library

gdb- 4.17.gnat.3.11p/opcodes
      source for the library of opcode tables and disassemblers

gdb- 4.17.gnat.3.11p/readline
      source for the GNU command-line interface

gdb- 4.17.gnat.3.11p/glob
      source for the GNU filename pattern-matching subroutine

gdb- 4.17.gnat.3.11p/mmalloc
      source for the GNU memory-mapped malloc package
```

The simplest way to configure and build GDB is to run `configure` from the ‘`gdb-version-number`’ source directory, which in this example is the ‘`gdb- 4.17.gnat.3.11p`’ directory.

First switch to the ‘`gdb-version-number`’ source directory if you are not already in it; then run `configure`. Pass the identifier for the platform on which GDB will run as an argument.

For example:

```
cd gdb- 4.17.gnat.3.11p
./configure host
make
```

where *host* is an identifier such as ‘`sun4`’ or ‘`decstation`’, that identifies the platform where GDB will run. (You can often leave off *host*; `configure` tries to guess the correct value by examining your system.)

---

<sup>1</sup> If you have a more recent version of GDB than 4.17.gnat.3.11p, look at the ‘`README`’ file in the sources; we may have improved the installation procedures since publishing this manual.

Running `configure host` and then running `make` builds the `'bfd'`, `'readline'`, `'mmalloc'`, and `'libiberty'` libraries, then `gdb` itself. The configured source files, and the binaries, are left in the corresponding source directories.

`configure` is a Bourne-shell (`/bin/sh`) script; if your system does not recognize this automatically when you run a different shell, you may need to run `sh` on it explicitly:

```
sh configure host
```

If you run `configure` from a directory that contains source directories for multiple libraries or programs, such as the `'gdb- 4.17.gnat.3.11p'` source directory for version 4.17.gnat.3.11p, `configure` creates configuration files for every directory level underneath (unless you tell it not to, with the `'--norecursion'` option).

You can run the `configure` script from any of the subordinate directories in the GDB distribution if you only want to configure that subdirectory, but be sure to specify a path to it.

For example, with version 4.17.gnat.3.11p, type the following to configure only the `bfd` subdirectory:

```
cd gdb- 4.17.gnat.3.11p/bfd
../configure host
```

You can install `gdb` anywhere; it has no hardwired paths. However, you should make sure that the shell on your path (named by the `'SHELL'` environment variable) is publicly readable. Remember that GDB uses the shell to start your program—some systems refuse to let GDB debug child processes whose programs are not readable.

## D.1 Compiling GDB in another directory

If you want to run GDB versions for several host or target machines, you need a different `gdb` compiled for each combination of host and target. `configure` is designed to make this easy by allowing you to generate each configuration in a separate subdirectory, rather than in the source directory. If your `make` program handles the `'VPATH'` feature (GNU `make` does), running `make` in each of these directories builds the `gdb` program specified there.

To build `gdb` in a separate directory, run `configure` with the `'--srcdir'` option to specify where to find the source. (You also need to specify a path to find `configure` itself from your working directory. If the path to `configure` would be the same as the argument to `'--srcdir'`, you can leave out the `'--srcdir'` option; it is assumed.)

For example, with version 4.17.gnat.3.11p, you can build GDB in a separate directory for a Sun 4 like this:

```
cd gdb- 4.17.gnat.3.11p
mkdir ../gdb-sun4
cd ../gdb-sun4
../gdb- 4.17.gnat.3.11p/configure sun4
make
```

When `configure` builds a configuration using a remote source directory, it creates a tree for the binaries with the same structure (and using the same names) as the tree under the source directory. In the example, you'd find the Sun 4 library `'libiberty.a'` in the directory `'gdb-sun4/libiberty'`, and GDB itself in `'gdb-sun4/gdb'`.



One popular reason to build several GDB configurations in separate directories is to configure GDB for cross-compiling (where GDB runs on one machine—the *host*—while debugging programs that run on another machine—the *target*). You specify a cross-debugging target by giving the ‘`--target=target`’ option to `configure`.

When you run `make` to build a program or library, you must run it in a configured directory—whatever directory you were in when you called `configure` (or one of its subdirectories).

The `Makefile` that `configure` generates in each source directory also runs recursively. If you type `make` in a source directory such as ‘`gdb- 4.17.gnat.3.11p`’ (or in a separate configured directory configured with ‘`--srcdir=dirname/gdb- 4.17.gnat.3.11p`’), you will build all the required libraries, and then build GDB.

When you have multiple hosts or targets configured in separate directories, you can run `make` on them in parallel (for example, if they are NFS-mounted on each of the hosts); they will not interfere with each other.

## D.2 Specifying names for hosts and targets

The specifications used for hosts and targets in the `configure` script are based on a three-part naming scheme, but some short predefined aliases are also supported. The full naming scheme encodes three pieces of information in the following pattern:

*architecture-vendor-os*

For example, you can use the alias `sun4` as a *host* argument, or as the value for *target* in a `--target=target` option. The equivalent full name is ‘`sparc-sun-sunos4`’.

The `configure` script accompanying GDB does not provide any query facility to list all supported host and target names or aliases. `configure` calls the Bourne shell script `config.sub` to map abbreviations to full names; you can read the script, if you wish, or you can use it to test your guesses on abbreviations—for example:

```
% sh config.sub sun4
sparc-sun-sunos4.1.1
% sh config.sub sun3
m68k-sun-sunos4.1.1
% sh config.sub decstation
mips-dec-ultrix4.2
% sh config.sub hp300bsd
m68k-hp-bsd
% sh config.sub i386v
i386-unknown-sysv
% sh config.sub i786v
Invalid configuration 'i786v': machine 'i786v' not recognized
```

`config.sub` is also distributed in the GDB source directory (‘`gdb- 4.17.gnat.3.11p`’, for version `4.17.gnat.3.11p`).

### D.3 configure options

Here is a summary of the `configure` options and arguments that are most often useful for building GDB. `configure` also has several other options not listed here. See Info file ‘`configure.info`’, node ‘What Configure Does’, for a full explanation of `configure`.

```
configure [--help]
          [--prefix=dir]
          [--srcdir=dirname]
          [--norecursion] [--rm]
          [--target=target] host
```

You may introduce options with a single ‘-’ rather than ‘--’ if you prefer; but you may abbreviate option names if you use ‘--’.

`--help`     Display a quick summary of how to invoke `configure`.

`--prefix=dir`

Configure the source to install programs and files under directory ‘*dir*’.

`--srcdir=dirname`

**Warning: using this option requires GNU make, or another make that implements the VPATH feature.**

Use this option to make configurations in directories separate from the GDB source directories. Among other things, you can use this to build (or maintain) several configurations simultaneously, in separate directories. `configure` writes configuration specific files in the current directory, but arranges for them to use the source in the directory *dirname*. `configure` creates directories under the working directory in parallel to the source directories below *dirname*.

`--norecursion`

Configure only the directory level where `configure` is executed; do not propagate configuration to subdirectories.

`--rm`

Remove files otherwise built during configuration.

`--target=target`

Configure GDB for cross-debugging programs running on the specified *target*. Without this option, GDB is configured to debug programs that run on the same machine (*host*) as GDB itself.

There is no convenient way to generate a list of all available targets.

*host* ...

Configure GDB to run on the specified *host*.

There is no convenient way to generate a list of all available hosts.

`configure` accepts other options, for compatibility with configuring other GNU tools recursively; but these are the only options that affect GDB or its supporting libraries.

## Index

- #**  
# ..... 9
- \$**  
\$ ..... 52  
\$\$ ..... 52  
\$\_ ..... 54  
\$\_ and info breakpoints ..... 19  
\$\_ and info line ..... 40  
\$, \$\_, and value history ..... 46  
\$\_\_ ..... 54  
\$\_exitcode ..... 54  
\$bpnum ..... 17  
\$cdir ..... 39  
\$cwd ..... 39
- .**  
.esgdbinit ..... 103  
'gdbinit' ..... 102  
.os68gdbinit ..... 103  
.vxgdbinit ..... 103
- :**  
:: ..... 42
- @**  
@ ..... 43
- {**  
{type} ..... 42
- A**  
a.out and C++ ..... 61  
abbreviation ..... 9  
active targets ..... 83  
Ada ..... 64  
Ada mode, general ..... 64  
Ada, deviations from ..... 65  
Ada, exceptions ..... 66  
Ada, generic units ..... 69  
Ada, omissions from ..... 64  
Ada, problems ..... 70  
Ada, tasking ..... 66  
artificial array ..... 43  
assembly instructions ..... 40  
assignment ..... 75  
automatic display ..... 46  
awatch ..... 21
- B**  
b ..... 17  
backtrace ..... 34  
break ..... 17  
break ... task *taskno* ..... 68  
break assert ..... 21  
break exception ..... 21  
break in overloaded functions ..... 63  
breakpoint commands ..... 25  
breakpoint conditions ..... 24  
breakpoint numbers ..... 17  
breakpoint on memory address ..... 17  
breakpoint on variable modification ..... 17  
breakpoint subroutine, remote ..... 86  
breakpoints ..... 17  
breakpoints and tasks ..... 68  
bt ..... 34  
bug criteria ..... 109  
bug reports ..... 109  
bugs in GDB ..... 109
- C**  
c ..... 27  
C and C++ ..... 59  
C and C++ constants ..... 61  
C and C++ defaults ..... 62  
C and C++ operators ..... 59  
C++ ..... 59  
C++ and object formats ..... 61  
C++ exception handling ..... 63  
C++ scope resolution ..... 42  
C++ support, not in COFF ..... 61  
C++ symbol decoding style ..... 51  
C++ symbol display ..... 63  
call ..... 77  
call overloaded functions ..... 62  
call stack ..... 33  
calling functions ..... 77  
calling make ..... 7  
casts, to view memory ..... 42  
catch ..... 21  
catch exceptions ..... 36  
cdir ..... 39  
checksum, for GDB remote ..... 89

choosing target byte order .....	84
<b>clear</b> .....	23
clearing breakpoints, watchpoints .....	22
COFF versus C++ .....	61
colon-colon .....	42
command files .....	102
command line editing .....	97
<b>commands</b> .....	25
commands for C++ .....	63
comment .....	9
compilation directory .....	39
<b>Compiling</b> .....	94
<b>complete</b> .....	12
completion .....	9
completion of quoted strings .....	10
<b>condition</b> .....	24
conditional breakpoints .....	24
configuring GDB .....	125
confirmation .....	100
<b>continue</b> .....	27
continuing .....	27
control C, and remote debugging .....	87
convenience variables .....	53
CPU simulator .....	95
crash of debugger .....	109
current directory .....	39
<b>cwd</b> .....	39
<b>D</b>	
<b>d</b> .....	23
debugger crash .....	109
debugging optimized code .....	15
debugging stub, example .....	89
debugging target .....	83
<b>define</b> .....	101
<b>delete</b> .....	23
delete breakpoints .....	23
<b>delete display</b> .....	47
deleting breakpoints, watchpoints .....	22
demangling .....	51
<b>dir</b> .....	39
directories for source files .....	39
<b>directory</b> .....	39
directory, compilation .....	39
directory, current .....	39
<b>dis</b> .....	23
<b>disable</b> .....	23
disable breakpoints .....	23
<b>disable display</b> .....	47
<b>disassemble</b> .....	40
<b>display</b> .....	46
display of expressions .....	46
<b>do</b> .....	35
<b>document</b> .....	101
documentation .....	123
<b>down</b> .....	35
<b>down-silently</b> .....	35
download to Sparclet .....	95
download to VxWorks .....	93
<b>E</b>	
<b>echo</b> .....	103
ECOFF and C++ .....	61
editing .....	97
<b>editing-mode</b> .....	115
ELF/DWARF and C++ .....	61
ELF/stabs and C++ .....	61
<b>else</b> .....	101
Emacs .....	105
<b>enable</b> .....	23
<b>enable breakpoints</b> .....	23
<b>enable display</b> .....	47
<b>end</b> .....	25
entering numbers .....	99
error on valid input .....	109
event designators .....	121
examining data .....	41
examining memory .....	45
exception handlers .....	21, 36
<b>exceptionHandler</b> .....	87
<b>exec-file</b> .....	79
executable file .....	79
exiting GDB .....	7
expansion .....	121
expressions .....	41
expressions in Ada .....	64
expressions in C or C++ .....	59
expressions in C++ .....	61
<b>F</b>	
<b>f</b> .....	35
fatal signal .....	109
fatal signals .....	30
<b>fg</b> .....	27
<b>file</b> .....	79
<b>finish</b> .....	28
flinching .....	100
floating point .....	56
floating point registers .....	54
<b>flush_i_cache</b> .....	87

- foo ..... 81
- fork, debugging programs which call ..... 16
- format options ..... 48
- formatted output ..... 44
- forward-search ..... 38
- frame ..... 33
- frame ..... 34, 35
- frame number ..... 33
- frame pointer ..... 33
- frameless execution ..... 33
- Fujitsu ..... 86
  
- G**
- g++ ..... 59
- GDB bugs, reporting ..... 109
- GDB reference card ..... 123
- GDBHISTFILE ..... 97
- gdbserve.nlm ..... 91
- gdbserver ..... 90
- getDebugChar ..... 87
- GNU C++ ..... 59
- GNU Emacs ..... 105
  
- H**
- h ..... 11
- H8/300 or H8/500 simulator ..... 95
- handle ..... 30
- handle\_exception ..... 86
- handling signals ..... 30
- hbreak ..... 19
- help ..... 11
- help target ..... 83
- help user-defined ..... 101
- history expansion ..... 98
- history file ..... 97
- history number ..... 52
- history save ..... 98
- history size ..... 98
- history substitution ..... 97
- Hitachi ..... 86
- Hitachi SH simulator ..... 95
- horizontal-scroll-mode ..... 116
  
- I**
- i ..... 12
- i386 ..... 86
- i386-stub.c ..... 86
- if ..... 101
- ignore ..... 25
- ignore count (of breakpoint) ..... 25
- INCLUDE\_RDB ..... 92
- info ..... 12
- info address ..... 71
- info all-registers ..... 54
- info args ..... 36
- info breakpoints ..... 19
- info catch ..... 36
- info display ..... 47
- info exceptions ..... 66
- info f ..... 36
- info files ..... 80
- info float ..... 56
- info frame ..... 36, 58
- info functions ..... 72
- info line ..... 40
- info locals ..... 36
- info program ..... 17
- info registers ..... 54
- info s ..... 34
- info set ..... 12
- info signals ..... 30
- info source ..... 58, 72
- info sources ..... 72
- info stack ..... 34
- info target ..... 80
- info task *taskno* ..... 68
- info tasks ..... 67
- info types ..... 72
- info variables ..... 72
- info watchpoints ..... 21
- inheritance ..... 63
- init file ..... 102
- init file name ..... 102
- initial frame ..... 33
- innermost frame ..... 33
- inspect ..... 41
- installation ..... 125
- instructions, assembly ..... 40
- Intel ..... 86
- interaction, readline ..... 113
- internal GDB breakpoints ..... 20
- interrupt ..... 7
- interrupting remote programs ..... 89
- interrupting remote targets ..... 87
- invalid input ..... 109
  
- J**
- jump ..... 76

**L**

<code>l</code> .....	37
languages .....	57
latest breakpoint .....	17
leaving GDB .....	7
<code>linespec</code> .....	37
<code>list</code> .....	37
listing machine instructions .....	40
<code>load filename</code> .....	80

**M**

<code>m680x0</code> .....	86
<code>m68k-stub.c</code> .....	86
machine instructions .....	40
<code>maint info breakpoints</code> .....	20
<code>maint print psymbols</code> .....	73
<code>maint print symbols</code> .....	73
<code>make</code> .....	7
<code>mapped</code> .....	80
<code>mark-modified-lines</code> .....	116
member functions .....	62
memory tracing .....	17
memory, viewing as typed object .....	42
memory-mapped symbol file .....	80
<code>memset</code> .....	88
Motorola 680x0 .....	86
multiple processes .....	16
multiple targets .....	83

**N**

<code>n</code> .....	28
names of symbols .....	71
namespace in C++ .....	62
negative breakpoint numbers .....	20
<code>next</code> .....	28
<code>nexti</code> .....	29
<code>ni</code> .....	29
number representation .....	99
numbers for breakpoints .....	17

**O**

object formats and C++ .....	61
online documentation .....	11
optimized code, debugging .....	15
outermost frame .....	33
<code>output</code> .....	103
output formats .....	44
overloading .....	26
overloading in C++ .....	63

**P**

packets, reporting on stdout .....	90
partial symbol dump .....	73
patching binaries .....	77
pauses in output .....	99
pointer, finding referent .....	49
<code>prefer-visible-bell</code> .....	116
<code>print</code> .....	41
print settings .....	48
<code>printf</code> .....	103
printing data .....	41
processes, multiple .....	16
prompt .....	97
protocol, GDB remote serial .....	89
<code>ptype</code> .....	71
<code>putDebugChar</code> .....	87

**Q**

<code>q</code> .....	7
<code>quit [expression]</code> .....	7
quotes in commands .....	10
quoting names .....	71

**R**

raise exceptions .....	22
<code>rbreak</code> .....	19
reading symbols immediately .....	80
<code>readline</code> .....	97
<code>readnow</code> .....	80
reference card .....	123
reference declarations .....	62
registers .....	54
regular expression .....	19
reloading symbols .....	72
remote connection without stubs .....	90
remote debugging .....	85
remote programs, interrupting .....	89
remote serial debugging summary .....	88
remote serial debugging, overview .....	85
remote serial protocol .....	89
remote serial stub .....	86
remote serial stub list .....	85
remote serial stub, initialization .....	86
remote serial stub, main routine .....	86
remote stub, example .....	89
remote stub, support routines .....	87
<code>remotetimeout</code> .....	94
repeating commands .....	9
reporting bugs in GDB .....	109
resuming execution .....	27

RET .....	9	set print sevenbit-strings .....	50
return .....	76	set print static-members .....	52
returning from a function .....	76	set print symbol-filename .....	48
reverse-search .....	38	set print union .....	50
run .....	15	set print vtbl .....	52
running .....	15	set prompt .....	97
Running .....	94	set remotedebug .....	90
running and debugging Sparclet programs .....	95	set symbol-reloading .....	73
running VxWorks tasks .....	94	set variable .....	75
rwatch .....	21	set varsize-limit .....	70
		set verbose .....	100
		set width .....	99
		set write .....	77
		set_debug_traps .....	86
		setting variables .....	75
		setting watchpoints .....	21
		SH .....	86
		sh-stub.c .....	86
		shell .....	7
		shell escape .....	7
		show .....	12
		show commands .....	98
		show complaints .....	100
		show confirm .....	100
		show convenience .....	54
		show copying .....	13
		show demangle-style .....	52
		show directories .....	39
		show editing .....	97
		show endian .....	84
		show gnutarget .....	83
		show height .....	99
		show history .....	98
		show input-radix .....	99
		show language .....	58
		show listsize .....	37
		show output-radix .....	100
		show print address .....	48
		show print array .....	49
		show print asm-demangle .....	51
		show print demangle .....	51
		show print elements .....	49
		show print max-symbolic-offset .....	49
		show print object .....	52
		show print pretty .....	50
		show print sevenbit-strings .....	50
		show print static-members .....	52
		show print symbol-filename .....	48
		show print union .....	50
		show print vtbl .....	52
<b>S</b> .....			
s .....	28		
saving symbol table .....	80		
search .....	38		
searching .....	38		
section .....	80		
select-frame .....	34		
selected frame .....	33		
serial connections, debugging .....	90		
serial line, target remote .....	88		
serial protocol, GDB remote .....	89		
set .....	12		
set assembly-language .....	40		
set complaints .....	100		
set confirm .....	100		
set demangle-style .....	51		
set editing .....	97		
set endian auto .....	84		
set endian big .....	84		
set endian little .....	84		
set gnutarget .....	83		
set height .....	99		
set history expansion .....	98		
set history filename .....	97		
set history save .....	98		
set history size .....	98		
set input-radix .....	99		
set language .....	58		
set listsize .....	37		
set output-radix .....	99		
set print address .....	48		
set print array .....	49		
set print asm-demangle .....	51		
set print demangle .....	51		
set print elements .....	49		
set print max-symbolic-offset .....	49		
set print null-stop .....	49		
set print object .....	52		
set print pretty .....	50		

show prompt	97
show remotedebug	90
show symbol-reloading	73
show user	102
show values	53
show varsize-limit	70
show verbose	100
show version	12
show warranty	13
show width	99
show write	77
si	29
signals	30
silent	26
sim	95
simulator	95
simulator, H8/300 or H8/500	95
simulator, Hitachi SH	95
simulator, Z8000	95
size of screen	99
source	103
source path	39
Sparc	86
sparc-stub.c	86
sparcl-stub.c	86
Sparclet	94
SparcLite	86
stack frame	33
stacking targets	83
starting	15
step	28
stepi	29
stepping	27
stub example, remote debugging	89
stupid questions	100
symbol decoding style, C++	51
symbol dump	73
symbol names	71
symbol overloading	26
symbol table	79
symbol-file	79
symbols, reading immediately	80
<b>T</b>	
target	83
target array	84
target bug	84
target byte order	84
target cpu32bug	84
target est	84
target exec	83
target op50n	84
target remote	83
target rom68k	84
target sim	84, 95
target sparclite	84
target vxworks	84
target w89k	84
task breakpoints	68
task switching	68
task taskno	68
tbreak	18
TCP port, target remote	88
thbreak	19
this	62
toggle-editing-mode	120
type casting memory	42
type conversions in C++	62
<b>U</b>	
u	28
undisplay	47
unknown address, locating	44
until	28
up	35
up-silently	35
user-defined command	101
<b>V</b>	
value history	52
variable name conflict	42
variable values, wrong	42
variables, setting	75
version number	12
vi style command editing	120
VxWorks	92
vxworks-timeout	92
<b>W</b>	
watch	21
watchpoints	17
whatis	71
where	34
while	101
wild pointer, interpreting	49
word completion	9
working directory	39
working language	57
writing into executables	77
wrong values	42



**X**

x .....	45
XCOFF and C++ .....	61

**Z**

Z8000 simulator .....	95
-----------------------	----

The body of this manual is set in  
cmr10 at 10.95pt,  
with headings in **cmb10 at 10.95pt**  
and examples in cmtt10 at 10.95pt.  
*cmti10 at 10.95pt*,  
**cmb10 at 10.95pt**, and  
*cmsl10 at 10.95pt*  
are used for emphasis.

# Table of Contents

<b>Summary of GDB</b> .....	<b>1</b>
Free software .....	1
Contributors to GDB .....	1
<b>1 Getting In and Out of GDB</b> .....	<b>5</b>
1.1 Invoking GDB .....	5
1.1.1 Choosing files .....	5
1.1.2 Choosing modes .....	6
1.2 Quitting GDB .....	7
1.3 Shell commands .....	7
<b>2 GDB Commands</b> .....	<b>9</b>
2.1 Command syntax .....	9
2.2 Command completion .....	9
2.3 Getting help .....	11
<b>3 Running Programs Under GDB</b> .....	<b>15</b>
3.1 Compiling for debugging .....	15
3.2 Starting your program .....	15
3.3 Debugging programs with multiple processes .....	16
<b>4 Stopping and Continuing</b> .....	<b>17</b>
4.1 Breakpoints, watchpoints, and exceptions .....	17
4.1.1 Setting breakpoints .....	17
4.1.2 Setting watchpoints .....	21
4.1.3 Breakpoints and exceptions .....	21
4.1.4 Deleting breakpoints .....	22
4.1.5 Disabling breakpoints .....	23
4.1.6 Break conditions .....	24
4.1.7 Breakpoint command lists .....	25
4.1.8 Breakpoint menus .....	26
4.2 Continuing and stepping .....	27
4.3 Signals .....	30
<b>5 Examining the Stack</b> .....	<b>33</b>
5.1 Stack frames .....	33
5.2 Backtraces .....	34
5.3 Selecting a frame .....	35
5.4 Information about a frame .....	36

<b>6</b>	<b>Examining Source Files . . . . .</b>	<b>37</b>
6.1	Printing source lines . . . . .	37
6.2	Searching source files . . . . .	38
6.3	Specifying source directories . . . . .	39
6.4	Source and machine code . . . . .	39
<b>7</b>	<b>Examining Data . . . . .</b>	<b>41</b>
7.1	Expressions . . . . .	41
7.2	Program variables . . . . .	42
7.3	Artificial arrays . . . . .	43
7.4	Output formats . . . . .	44
7.5	Examining memory . . . . .	45
7.6	Automatic display . . . . .	46
7.7	Print settings . . . . .	48
7.8	Value history . . . . .	52
7.9	Convenience variables . . . . .	53
7.10	Registers . . . . .	54
7.11	Floating point hardware . . . . .	56
<b>8</b>	<b>Using GDB with Different Languages . . . . .</b>	<b>57</b>
8.1	Switching between source languages . . . . .	57
8.1.1	List of filename extensions and languages . . . . .	57
8.1.2	Setting the working language . . . . .	58
8.1.3	Having GDB infer the source language . . . . .	58
8.2	Displaying the language . . . . .	58
8.3	Supported languages . . . . .	59
8.3.1	C and C++ . . . . .	59
8.3.1.1	C and C++ operators . . . . .	59
8.3.1.2	C and C++ constants . . . . .	61
8.3.1.3	C++ expressions . . . . .	61
8.3.1.4	C and C++ defaults . . . . .	62
8.3.1.5	GDB and C . . . . .	62
8.3.1.6	GDB features for C++ . . . . .	63
8.3.2	Ada . . . . .	64
8.3.2.1	Introduction . . . . .	64
8.3.2.2	Omissions from Ada . . . . .	64
8.3.2.3	Additions to Ada . . . . .	65
8.3.2.4	Stopping at the Beginning . . . . .	66
8.3.2.5	Breaking on Ada Exceptions . . . . .	66
8.3.2.6	Extensions for Ada Tasks . . . . .	66
8.3.2.7	Debugging Generic Units . . . . .	69
8.3.2.8	Set commands for Ada . . . . .	70
8.3.2.9	Known Peculiarities of Ada Mode . . . . .	70
<b>9</b>	<b>Examining the Symbol Table . . . . .</b>	<b>71</b>

<b>10</b>	<b>Altering Execution</b> . . . . .	<b>75</b>
10.1	Assignment to variables . . . . .	75
10.2	Continuing at a different address . . . . .	76
10.3	Returning from a function . . . . .	76
10.4	Calling program functions . . . . .	77
10.5	Patching programs . . . . .	77
<b>11</b>	<b>GDB Files</b> . . . . .	<b>79</b>
11.1	Commands to specify files . . . . .	79
11.2	Errors reading symbol files . . . . .	81
<b>12</b>	<b>Specifying a Debugging Target</b> . . . . .	<b>83</b>
12.1	Active targets . . . . .	83
12.2	Commands for managing targets . . . . .	83
12.3	Choosing target byte order . . . . .	84
12.4	Remote debugging . . . . .	85
12.4.1	The GDB remote serial protocol . . . . .	85
12.4.1.1	What the stub can do for you . . . . .	86
12.4.1.2	What you must do for the stub . . . . .	87
12.4.1.3	Putting it all together . . . . .	88
12.4.1.4	Communication protocol . . . . .	89
12.4.1.5	Using the <code>gdbserver</code> program . . . . .	90
12.4.1.6	Using the <code>gdbserve.nlm</code> program . . . . .	91
12.4.2	GDB and VxWorks . . . . .	92
12.4.2.1	Connecting to VxWorks . . . . .	93
12.4.2.2	VxWorks download . . . . .	93
12.4.2.3	Running tasks . . . . .	94
12.4.3	GDB and Sparclet . . . . .	94
12.4.3.1	Setting file to debug . . . . .	94
12.4.3.2	Connecting to Sparclet . . . . .	95
12.4.3.3	Sparclet download . . . . .	95
12.4.3.4	Running and debugging . . . . .	95
12.4.4	Simulated CPU target . . . . .	95
<b>13</b>	<b>Controlling GDB</b> . . . . .	<b>97</b>
13.1	Prompt . . . . .	97
13.2	Command editing . . . . .	97
13.3	Command history . . . . .	97
13.4	Screen size . . . . .	99
13.5	Numbers . . . . .	99
13.6	Optional warnings and messages . . . . .	100
<b>14</b>	<b>Canned Sequences of Commands</b> . . . . .	<b>101</b>
14.1	User-defined commands . . . . .	101
14.2	User-defined command hooks . . . . .	102
14.3	Command files . . . . .	102
14.4	Commands for controlled output . . . . .	103

<b>15</b>	<b>Using GDB under Emacs</b> . . . . .	<b>105</b>
<b>16</b>	<b>Reporting Bugs in GDB</b> . . . . .	<b>109</b>
16.1	Have you found a bug? . . . . .	109
16.2	How to report bugs . . . . .	109
<b>Appendix A</b>	<b>Command Line Editing</b> . . . . .	<b>113</b>
A.1	Introduction to Line Editing . . . . .	113
A.2	Readline Interaction . . . . .	113
A.2.1	Readline Bare Essentials . . . . .	113
A.2.2	Readline Movement Commands . . . . .	114
A.2.3	Readline Killing Commands . . . . .	114
A.2.4	Readline Arguments . . . . .	115
A.3	Readline Init File . . . . .	115
A.3.1	Readline Init Syntax . . . . .	115
A.3.1.1	Commands For Moving . . . . .	117
A.3.1.2	Commands For Manipulating The History . . . . .	117
A.3.1.3	Commands For Changing Text . . . . .	118
A.3.1.4	Killing And Yanking . . . . .	119
A.3.1.5	Specifying Numeric Arguments . . . . .	119
A.3.1.6	Letting Readline Type For You . . . . .	119
A.3.1.7	Some Miscellaneous Commands . . . . .	120
A.3.2	Readline vi Mode . . . . .	120
<b>Appendix B</b>	<b>Using History Interactively</b> . . . . .	<b>121</b>
B.1	History Interaction . . . . .	121
B.1.1	Event Designators . . . . .	121
B.1.2	Word Designators . . . . .	121
B.1.3	Modifiers . . . . .	122
<b>Appendix C</b>	<b>Formatting Documentation</b> . . . . .	<b>123</b>
<b>Appendix D</b>	<b>Installing GDB</b> . . . . .	<b>125</b>
D.1	Compiling GDB in another directory . . . . .	126
D.2	Specifying names for hosts and targets . . . . .	127
D.3	configure options . . . . .	128
<b>Index</b> . . . . .		<b>129</b>