Using ld

The GNU linker

`ld` version 2
January 1994

Steve Chamberlain and Roland Pesch
Cygnus Support

Cygnus Support
steve@cygnus.com, pesch@cygnus.com
Using LD, the GNU linker
Edited by Jeffrey Osier (jeffrey@cygnus.com)
and Roland Pesch (pesch@cygnus.com)

Copyright © 1991, 1992, 1993, 1994 Free Software Foundation, Inc.

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.

1 Overview

`ld` combines a number of object and archive files, relocates their data and ties up symbol references. Usually the last step in compiling a program is to run `ld`.

`ld` accepts Linker Command Language files written in a superset of AT&T's Link Editor Command Language syntax, to provide explicit and total control over the linking process.

This version of `ld` uses the general purpose BFD libraries to operate on object files. This allows `ld` to read, combine, and write object files in many different formats—for example, COFF or `a.out`. Different formats may be linked together to produce any available kind of object file. See Chapter 5 [BFD], page 29, for more information.

Aside from its flexibility, the GNU linker is more helpful than other linkers in providing diagnostic information. Many linkers abandon execution immediately upon encountering an error; whenever possible, `ld` continues executing, allowing you to identify other errors (or, in some cases, to get an output file in spite of the error).

2 Invocation

The GNU linker `ld` is meant to cover a broad range of situations, and to be as compatible as possible with other linkers. As a result, you have many choices to control its behavior.

2.1 Command Line Options

Here is a summary of the options you can use on the `ld` command line:

```
ld [ -o output ]  objfile...
  [ -Aarchitecture ]  [ -b input-format ]  [ -Bstatic ]
  [ -c MRI-commandfile ]  [ -d | -dc | -dp ]
  [ -defsym symbol=expression ]
  [ -dynamic-linker file ]
  [ -e entry ]  [ -F ]  [ -F format ]
  [ -format input-format ]  [ -g ]  [ -G size ]  [ -help ]
  [ -i ]  [ -larchive ]  [ -Lsearchdir ]  [ -M ]
  [ -Map mapfile ]  [ -m emulation ]  [ -N | -n ]
  [ -noinhibit-exec ]  [ -oformat output-format ]
  [ -R filename ]  [ -relax ]  [ -retain-symbols-file filename ]
  [ -r | -Ur ]  [ -rpath dir ] [ -S ]  [ -s ] [ -soname name ]
  [ -sort-common ] [ -stats ] [ -T commandfile ]
  [ -Ttext org ]  [ -Tdata org ]
  [ -Tbss org ]  [ -t ] [ -traditional-format ]
  [ -u symbol]  [-V]  [-v]  [ -verbose]  [ -version ]
  [ -warn-common ]  [ -warn-once ] [ -y symbol ]  [ -X ]  [-x ]
  [ -( [ archives ] -) ] [ --start-group [ archives ] --end-group ]
```

This plethora of command-line options may seem intimidating, but in actual practice few of them are used in any particular context. For instance, a frequent use of `ld` is to link standard Unix object files on a standard, supported Unix system. On such a system, to link a file `hello.o`:

```
ld -o output /lib/crt0.o hello.o -lc
```

This tells `ld` to produce a file called *output* as the result of linking the file `/lib/crt0.o` with `hello.o` and the library `libc.a`, which will come from the standard search directories. (See the discussion of the '`-l`' option below.)

The command-line options to `ld` may be specified in any order, and may be repeated at will. Repeating most options with a different argument will either have no further effect, or override prior occurrences (those further to the left on the command line) of that option.

The exceptions—which may meaningfully be used more than once—are '`-A`', '`-b`' (or its synonym '`-format`'), '`-defsym`', '`-L`', '`-l`', '`-R`', '`-u`', and '`-(`' (or its synonym '`--start-group`')..

The list of object files to be linked together, shown as *objfile*..., may follow, precede, or be mixed in with command-line options, except that an *objfile* argument may not be placed between an option and its argument.

Usually the linker is invoked with at least one object file, but you can specify other forms of binary input files using '`-l`', '`-R`', and the script command language. If *no* binary input files at all are specified, the linker does not produce any output, and issues the message '`No input files`'.

If the linker can not recognize the format of an object file, it will assume that it is a linker script. A script specified in this way augments the main linker script used for the link (either the default

linker script or the one specified by using '`-T`'). This feature permits the linker to link against a file which appears to be an object or an archive, but actually merely defines some symbol values, or uses `INPUT` or `GROUP` to load other objects. See Chapter 3 [Commands], page 11.

For options whose names are a single letter, option arguments must either follow the option letter without intervening whitespace, or be given as separate arguments immediately following the option that requires them.

For options whose names are multiple letters, either one dash or two can precede the option name; for example, '`--oformat`' and '`-oformat`' are equivalent. Arguments to multiple-letter options must either be separated from the option name by an equals sign, or be given as separate arguments immediately following the option that requires them. For example, '`--oformat srec`' and '`--oformat=srec`' are equivalent. Unique abbreviations of the names of multiple-letter options are accepted.

`-A`*architecture*

In the current release of `ld`, this option is useful only for the Intel 960 family of architectures. In that `ld` configuration, the *architecture* argument identifies the particular architecture in the 960 family, enabling some safeguards and modifying the archive-library search path. See Section 4.2 [`ld` and the Intel 960 family], page 27, for details.

Future releases of `ld` may support similar functionality for other architecture families.

`-b` *input-format*

`ld` may be configured to support more than one kind of object file. If your `ld` is configured this way, you can use the '`-b`' option to specify the binary format for input object files that follow this option on the command line. Even when `ld` is configured to support alternative object formats, you don't usually need to specify this, as `ld` should be configured to expect as a default input format the most usual format on each machine. *input-format* is a text string, the name of a particular format supported by the BFD libraries. (You can list the available binary formats with '`objdump -i`'.) '`-format` *input-format*' has the same effect, as does the script command `TARGET`. See Chapter 5 [BFD], page 29.

You may want to use this option if you are linking files with an unusual binary format. You can also use '`-b`' to switch formats explicitly (when linking object files of different formats), by including '`-b` *input-format*' before each group of object files in a particular format.

The default format is taken from the environment variable `GNUTARGET`. See Section 2.2 [Environment], page 10. You can also define the input format from a script, using the command `TARGET`; see Section 3.6 [Option Commands], page 24.

`-Bstatic` Do not link against shared libraries. This option is accepted for command-line compatibility with the SunOS linker.

`-c` *MRI-commandfile*

For compatibility with linkers produced by MRI, `ld` accepts script files written in an alternate, restricted command language, described in Appendix A [MRI Compatible Script Files], page 33. Introduce MRI script files with the option '`-c`'; use the '`-T`' option to run linker scripts written in the general-purpose `ld` scripting language. If *MRI-cmdfile* does not exist, `ld` looks for it in the directories specified by any '`-L`' options.

`-d`
`-dc`
`-dp` These three options are equivalent; multiple forms are supported for compatibility with other linkers. They assign space to common symbols even if a relocatable output file is specified (with '`-r`'). The script command `FORCE_COMMON_ALLOCATION` has the same effect. See Section 3.6 [Option Commands], page 24.

`-defsym` *symbol*`=`*expression*
Create a global symbol in the output file, containing the absolute address given by *expression*. You may use this option as many times as necessary to define multiple symbols in the command line. A limited form of arithmetic is supported for the *expression* in this context: you may give a hexadecimal constant or the name of an existing symbol, or use `+` and `-` to add or subtract hexadecimal constants or symbols. If you need more elaborate expressions, consider using the linker command language from a script (see Section 3.2.6 [Assignment: Symbol Definitions], page 14). *Note:* there should be no white space between *symbol*, the equals sign ("`=`"), and *expression*.

`-dynamic-linker` *file*
Set the name of the dynamic linker. This is only meaningful when generating dynamically linked ELF executables. The default dynamic linker is normally correct; don't use this unless you know what you are doing.

`-e` *entry* Use *entry* as the explicit symbol for beginning execution of your program, rather than the default entry point. See Section 3.5 [Entry Point], page 24, for a discussion of defaults and other ways of specifying the entry point.

`-F`
`-F`*format* Ignored. Some older linkers used this option throughout a compilation toolchain for specifying object-file format for both input and output object files. The mechanisms `ld` uses for this purpose (the '`-b`' or '`-format`' options for input files, '`-oformat`' option or the `TARGET` command in linker scripts for output files, the `GNUTARGET` environment variable) are more flexible, but `ld` accepts the '`-F`' option for compatibility with scripts written to call the old linker.

`-format` *input-format*
Synonym for '`-b` *input-format*'.

`-g` Ignored. Provided for compatibility with other tools.

`-G`*value*
`-G` *value* Set the maximum size of objects to be optimized using the GP register to *size* under MIPS ECOFF. Ignored for other object file formats.

`-help` Print a summary of the command-line options on the standard output and exit.

`-i` Perform an incremental link (same as option '`-r`').

`-l`*ar* Add archive file *archive* to the list of files to link. This option may be used any number of times. `ld` will search its path-list for occurrences of `lib`*ar*`.a` for every *archive* specified.

`-L`*searchdir*
`-L` *searchdir*
: Add path *searchdir* to the list of paths that `ld` will search for archive libraries and `ld` control scripts. You may use this option any number of times.

  The default set of paths searched (without being specified with '`-L`') depends on which emulation mode `ld` is using, and in some cases also on how it was configured. See Section 2.2 [Environment], page 10.

  The paths can also be specified in a link script with the `SEARCH_DIR` command.

`-M`
: Print (to the standard output) a link map—diagnostic information about where symbols are mapped by `ld`, and information on global common storage allocation.

`-Map` *mapfile*
: Print to the file *mapfile* a link map—diagnostic information about where symbols are mapped by `ld`, and information on global common storage allocation.

`-m`*emulation*
`-m` *emulation*
: Emulate the *emulation* linker. You can list the available emulations with the '`--verbose`' option. The default depends on how your `ld` was configured.

`-N`
: Set the text and data sections to be readable and writable. Also, do not page-align the data segment. If the output format supports Unix style magic numbers, mark the output as `OMAGIC`.

`-n`
: Set the text segment to be read only, and mark the output as `NMAGIC` if possible.

`-noinhibit-exec`
: Retain the executable output file whenever it is still usable. Normally, the linker will not produce an output file if it encounters errors during the link process; it exits without writing an output file when it issues any error whatsoever.

`-o` *output*
: Use *output* as the name for the program produced by `ld`; if this option is not specified, the name '`a.out`' is used by default. The script command `OUTPUT` can also specify the output file name.

`-oformat` *output-format*
: `ld` may be configured to support more than one kind of object file. If your `ld` is configured this way, you can use the '`-oformat`' option to specify the binary format for the output object file. Even when `ld` is configured to support alternative object formats, you don't usually need to specify this, as `ld` should be configured to produce as a default output format the most usual format on each machine. *output-format* is a text string, the name of a particular format supported by the BFD libraries. (You can list the available binary formats with '`objdump -i`'.) The script command `OUTPUT_FORMAT` can also specify the output format, but this option overrides it. See Chapter 5 [BFD], page 29.

`-R` *filename*
: Read symbol names and their addresses from *filename*, but do not relocate it or include it in the output. This allows your output file to refer symbolically to absolute locations of memory defined in other programs.

`-relax` An option with machine dependent effects. Currently this option is only supported on the H8/300 and the Intel 960. See Section 4.1 [`ld` and the H8/300], page 27. See Section 4.2 [`ld` and the Intel 960 family], page 27.

On some platforms, the '`-relax`' option performs global optimizations that become possible when the linker resolves addressing in the program, such as relaxing address modes and synthesizing new instructions in the output object file.

On platforms where this is not supported, '`-relax`' is accepted, but ignored.

`-retain-symbols-file` *filename*

Retain *only* the symbols listed in the file *filename*, discarding all others. *filename* is simply a flat file, with one symbol name per line. This option is especially useful in environments (such as VxWorks) where a large global symbol table is accumulated gradually, to conserve run-time memory.

'`-retain-symbols-file`' does *not* discard undefined symbols, or symbols needed for relocations.

You may only specify '`-retain-symbols-file`' once in the command line. It overrides '`-s`' and '`-S`'.

`-rpath` *dir*

Add a directory to the runtime library search path. This is only meaningful when linking an ELF executable with shared objects. All -rpath arguments are concatenated and passed to the runtime linker, which uses them to locate shared objects at runtime.

`-r` Generate relocatable output—i.e., generate an output file that can in turn serve as input to `ld`. This is often called *partial linking*. As a side effect, in environments that support standard Unix magic numbers, this option also sets the output file's magic number to `OMAGIC`. If this option is not specified, an absolute file is produced. When linking C`++` programs, this option *will not* resolve references to constructors; to do that, use '`-Ur`'.

This option does the same thing as '`-i`'.

`-S` Omit debugger symbol information (but not all symbols) from the output file.

`-s` Omit all symbol information from the output file.

`-soname` *name*

When creating an ELF shared object, set the internal DT_SONAME field to the specified name. When an executable is linked with a shared object which has a DT_SONAME field, then when the executable is run the dynamic linker will attempt to load the shared object specified by the DT_SONAME field rather than the using the file name given to the linker.

`-sort-common`

Normally, when `ld` places the global common symbols in the appropriate output sections, it sorts them by size. First come all the one byte symbols, then all the two bytes, then all the four bytes, and then everything else. This is to prevent gaps between symbols due to alignment constraints. This option disables that sorting.

`-stats` Compute and display statistics about the operation of the linker, such as execution time and memory usage.

`-Tbss` *org*
`-Tdata` *org*
`-Ttext` *org*
: Use *org* as the starting address for—respectively—the `bss`, `data`, or the `text` segment of the output file. *org* must be a single hexadecimal integer; for compatibility with other linkers, you may omit the leading '`0x`' usually associated with hexadecimal values.

`-T` *commandfile*
`-T`*commandfile*
: Read link commands from the file *commandfile*. These commands replace `ld`'s default link script (rather than adding to it), so *commandfile* must specify everything necessary to describe the target format. See Chapter 3 [Commands], page 11. If *commandfile* does not exist, `ld` looks for it in the directories specified by any preceding '`-L`' options. Multiple '`-T`' options accumulate.

`-t`
: Print the names of the input files as `ld` processes them.

`-traditional-format`
: For some targets, the output of `ld` is different in some ways from the output of some existing linker. This switch requests `ld` to use the traditional format instead.

  For example, on SunOS, `ld` combines duplicate entries in the symbol string table. This can reduce the size of an output file with full debugging information by over 30 percent. Unfortunately, the SunOS `dbx` program can not read the resulting program (`gdb` has no trouble). The '`-traditional-format`' switch tells `ld` to not combine duplicate entries.

`-u` *symbol*
: Force *symbol* to be entered in the output file as an undefined symbol. Doing this may, for example, trigger linking of additional modules from standard libraries. '`-u`' may be repeated with different option arguments to enter additional undefined symbols.

`-Ur`
: For anything other than C`++` programs, this option is equivalent to '`-r`': it generates relocatable output—i.e., an output file that can in turn serve as input to `ld`. When linking C`++` programs, '`-Ur`' *does* resolve references to constructors, unlike '`-r`'. It does not work to use '`-Ur`' on files that were themselves linked with '`-Ur`'; once the constructor table has been built, it cannot be added to. Use '`-Ur`' only for the last partial link, and '`-r`' for the others.

`--verbose`
: Display the version number for `ld` and list the linker emulations supported. Display which input files can and cannot be opened.

`-v`
`-V`
: Display the version number for `ld`.

`-version`
: Display the version number for `ld` and exit.

`-warn-common`
: Warn when a common symbol is combined with another common symbol or with a symbol definition. Unix linkers allow this somewhat sloppy practice, but linkers on some other operating systems do not. This option allows you to find potential problems from combining global symbols. Unfortunately, some C libraries use this practice, so you may get some warnings about symbols in the libraries as well as in your programs.

  There are three kinds of global symbols, illustrated here by C examples:

'`int i = 1;`'
: A definition, which goes in the initialized data section of the output file.

'`extern int i;`'
: An undefined reference, which does not allocate space. There must be either a definition or a common symbol for the variable somewhere.

'`int i;`'
: A common symbol. If there are only (one or more) common symbols for a variable, it goes in the uninitialized data area of the output file. The linker merges multiple common symbols for the same variable into a single symbol. If they are of different sizes, it picks the largest size. The linker turns a common symbol into a declaration, if there is a definition of the same variable.

The '`-warn-common`' option can produce five kinds of warnings. Each warning consists of a pair of lines: the first describes the symbol just encountered, and the second describes the previous symbol encountered with the same name. One or both of the two symbols will be a common symbol.

1. Turning a common symbol into a reference, because there is already a definition for the symbol.

   file(*section*): warning: common of '*symbol*'
   overridden by definition
   file(*section*): warning: defined here

2. Turning a common symbol into a reference, because a later definition for the symbol is encountered. This is the same as the previous case, except that the symbols are encountered in a different order.

   file(*section*): warning: definition of '*symbol*'
   overriding common
   file(*section*): warning: common is here

3. Merging a common symbol with a previous same-sized common symbol.

   file(*section*): warning: multiple common
   of '*symbol*'
   file(*section*): warning: previous common is here

4. Merging a common symbol with a previous larger common symbol.

   file(*section*): warning: common of '*symbol*'
   overridden by larger common
   file(*section*): warning: larger common is here

5. Merging a common symbol with a previous smaller common symbol. This is the same as the previous case, except that the symbols are encountered in a different order.

   file(*section*): warning: common of '*symbol*'
   overriding smaller common
   file(*section*): warning: smaller common is here

`-warn-once`
: Only warn once for each undefined symbol, rather than once per module which refers to it.

`-X`
: If '`-s`' or '`-S`' is also specified, delete only local symbols beginning with '`L`'.

`-x` If '`-s`' or '`-S`' is also specified, delete all local symbols, not just those beginning with '`L`'.

`-y` *symbol* Print the name of each linked file in which *symbol* appears. This option may be given any number of times. On many systems it is necessary to prepend an underscore.

This option is useful when you have an undefined symbol in your link but don't know where the reference is coming from.

`-(` *archives* `-)`
`--start-group` *archives* `--end-group`

The *archives* should be a list of archive files. They may be either explicit file names, or '`-l`' options.

The specified archives are searched repeatedly until no new undefined references are created. Normally, an archive is searched only once in the order that it is specified on the command line. If a symbol in that archive is needed to resolve an undefined symbol referred to by an object in an archive that appears later on the command line, the linker would not be able to resolve that reference. By grouping the archives, they all be searched repeatedly until all possible references are resolved.

Using this option has a significant performance cost. It is best to use it only when there are unavoidable circular references between two or more archives.

2.2 Environment Variables

You can change the behavior of `ld` with the environment variable `GNUTARGET`.

`GNUTARGET` determines the input-file object format if you don't use '`-b`' (or its synonym '`-format`'). Its value should be one of the BFD names for an input format (see Chapter 5 [BFD], page 29). If there is no `GNUTARGET` in the environment, `ld` uses the natural format of the target. If `GNUTARGET` is set to `default` then BFD attempts to discover the input format by examining binary input files; this method often succeeds, but there are potential ambiguities, since there is no method of ensuring that the magic number used to specify object-file formats is unique. However, the configuration procedure for BFD on each system places the conventional format for that system first in the search-list, so ambiguities are resolved in favor of convention.

3 Command Language

The command language provides explicit control over the link process, allowing complete specification of the mapping between the linker's input files and its output. It controls:

- input files
- file formats
- output file layout
- addresses of sections
- placement of common blocks

You may supply a command file (also known as a link script) to the linker either explicitly through the '`-T`' option, or implicitly as an ordinary file. If the linker opens a file which it cannot recognize as a supported object or archive format, it reports an error.

3.1 Linker Scripts

The `ld` command language is a collection of statements; some are simple keywords setting a particular option, some are used to select and group input files or name output files; and two statement types have a fundamental and pervasive impact on the linking process.

The most fundamental command of the `ld` command language is the `SECTIONS` command (see Section 3.4 [SECTIONS], page 17). Every meaningful command script must have a `SECTIONS` command: it specifies a "picture" of the output file's layout, in varying degrees of detail. No other command is required in all cases.

The `MEMORY` command complements `SECTIONS` by describing the available memory in the target architecture. This command is optional; if you don't use a `MEMORY` command, `ld` assumes sufficient memory is available in a contiguous block for all output. See Section 3.3 [MEMORY], page 17.

You may include comments in linker scripts just as in C: delimited by '`/*`' and '`*/`'. As in C, comments are syntactically equivalent to whitespace.

3.2 Expressions

Many useful commands involve arithmetic expressions. The syntax for expressions in the command language is identical to that of C expressions, with the following features:

- All expressions evaluated as integers and are of "long" or "unsigned long" type.
- All constants are integers.
- All of the C arithmetic operators are provided.
- You may reference, define, and create global variables.
- You may call special purpose built-in functions.

3.2.1 Integers

An octal integer is '`0`' followed by zero or more of the octal digits ('`01234567`').

```
_as_octal = 0157255;
```

A decimal integer starts with a non-zero digit followed by zero or more digits ('`0123456789`').

```
_as_decimal = 57005;
```

A hexadecimal integer is '`0x`' or '`0X`' followed by one or more hexadecimal digits chosen from '`0123456789abcdefABCDEF`'.

```
_as_hex = 0xdead;
```

To write a negative integer, use the prefix operator '`-`'; see Section 3.2.4 [Operators], page 13.

```
_as_neg = -57005;
```

Additionally the suffixes `K` and `M` may be used to scale a constant by 1024 or 1024^2 respectively. For example, the following all refer to the same quantity:

```
_fourk_1 = 4K;
_fourk_2 = 4096;
_fourk_3 = 0x1000;
```

3.2.2 Symbol Names

Unless quoted, symbol names start with a letter, underscore, or point and may include any letters, underscores, digits, points, and hyphens. Unquoted symbol names must not conflict with any keywords. You can specify a symbol which contains odd characters or has the same name as a keyword, by surrounding the symbol name in double quotes:

```
"SECTION" = 9;
"with a space" = "also with a space" + 10;
```

Since symbols can contain many non-alphabetic characters, it is safest to delimit symbols with spaces. For example, '`A-B`' is one symbol, whereas '`A - B`' is an expression involving subtraction.

3.2.3 The Location Counter

The special linker variable *dot* '`.`' always contains the current output location counter. Since the `.` always refers to a location in an output section, it must always appear in an expression within a `SECTIONS` command. The `.` symbol may appear anywhere that an ordinary symbol is allowed in an expression, but its assignments have a side effect. Assigning a value to the `.` symbol will cause the location counter to be moved. This may be used to create holes in the output section. The location counter may never be moved backwards.

```
SECTIONS
{
  output :
  {
  file1(.text)
  . = . + 1000;
  file2(.text)
  . += 1000;
  file3(.text)
  } = 0x1234;
```

```
}
```

In the previous example, `file1` is located at the beginning of the output section, then there is a 1000 byte gap. Then `file2` appears, also with a 1000 byte gap following before `file3` is loaded. The notation '`= 0x1234`' specifies what data to write in the gaps (see Section 3.4.4 [Section Options], page 22).

3.2.4 Operators

The linker recognizes the standard C set of arithmetic operators, with the standard bindings and precedence levels:

Precedence	Associativity	Operators
highest		
1	left	`- ~ !` †
2	left	`* / %`
3	left	`+ -`
4	left	`>> <<`
5	left	`== != > < <= >=`
6	left	`&`
7	left	`|`
8	left	`&&`
9	left	`||`
10	right	`? :`
11	right	`&= += -= *= /=` ‡
lowest		

† Prefix operators.
‡ See Section 3.2.6 [Assignment], page 14.

3.2.5 Evaluation

The linker uses "lazy evaluation" for expressions; it only calculates an expression when absolutely necessary. The linker needs the value of the start address, and the lengths of memory regions, in order to do any linking at all; these values are computed as soon as possible when the linker reads in the command file. However, other values (such as symbol values) are not known or needed until after storage allocation. Such values are evaluated later, when other information (such as the sizes of output sections) is available for use in the symbol assignment expression.

3.2.6 Assignment: Defining Symbols

You may create global symbols, and assign values (addresses) to global symbols, using any of the C assignment operators:

symbol `=` *expression* `;`
symbol `&=` *expression* `;`
symbol `+=` *expression* `;`
symbol `-=` *expression* `;`
symbol `*=` *expression* `;`
symbol `/=` *expression* `;`

Two things distinguish assignment from other operators in `ld` expressions.

- Assignment may only be used at the root of an expression; '`a=b+3;`' is allowed, but '`a+b=3;`' is an error.
- You must place a trailing semicolon ("`;`") at the end of an assignment statement.

Assignment statements may appear:

- as commands in their own right in an `ld` script; or
- as independent statements within a `SECTIONS` command; or
- as part of the contents of a section definition in a `SECTIONS` command.

The first two cases are equivalent in effect—both define a symbol with an absolute address. The last case defines a symbol whose address is relative to a particular section (see Section 3.4 [SECTIONS], page 17).

When a linker expression is evaluated and assigned to a variable, it is given either an absolute or a relocatable type. An absolute expression type is one in which the symbol contains the value that it will have in the output file; a relocatable expression type is one in which the value is expressed as a fixed offset from the base of a section.

The type of the expression is controlled by its position in the script file. A symbol assigned within a section definition is created relative to the base of the section; a symbol assigned in any other place is created as an absolute symbol. Since a symbol created within a section definition is relative to the base of the section, it will remain relocatable if relocatable output is requested. A symbol may be created with an absolute value even when assigned to within a section definition by using the absolute assignment function `ABSOLUTE`. For example, to create an absolute symbol whose address is the last byte of an output section named `.data`:

```
SECTIONS{ . . .
  .data :
    {
      *(.data)
      _edata = ABSOLUTE(.) ;
    }
. . . }
```

The linker tries to put off the evaluation of an assignment until all the terms in the source expression are known (see Section 3.2.5 [Evaluation], page 13). For instance, the sizes of sections cannot be known until after allocation, so assignments dependent upon these are not performed until after allocation. Some expressions, such as those depending upon the location counter *dot*, '`.`' must be evaluated during allocation. If the result of an expression is required, but the value is not available, then an error results. For example, a script like the following

```
SECTIONS { . . .
  text 9+this_isnt_constant :
    { . . .
    }
. . . }
```

will cause the error message "`Non constant expression for initial address`".

In some cases, it is desirable for a linker script to define a symbol only if it is referenced, and only if it is not defined by any object included in the link. For example, traditional linkers defined the symbol '`etext`'. However, ANSI C requires that the user be able to use '`etext`' as a function name without encountering an error. The `PROVIDE` keyword may be used to define a symbol, such as '`etext`', only if it is referenced but not defined. The syntax is `PROVIDE(`*symbol* `=` *expression*`)`.

3.2.7 Arithmetic Functions

The command language includes a number of built-in functions for use in link script expressions.

`ABSOLUTE(`*exp*`)`
: Return the absolute (non-relocatable, as opposed to non-negative) value of the expression *exp*. Primarily useful to assign an absolute value to a symbol within a section definition, where symbol values are normally section-relative.

`ADDR(`*section*`)`
: Return the absolute address of the named *section*. Your script must previously have defined the location of that section. In the following example, `symbol_1` and `symbol_2` are assigned identical values:

```
SECTIONS{ . . .
  .output1 :
    {
    start_of_output_1 = ABSOLUTE(.);
    . . .
    }
  .output :
    {
    symbol_1 = ADDR(.output1);
    symbol_2 = start_of_output_1;
    }
. . . }
```

`ALIGN(`*exp*`)`
: Return the result of the current location counter (`.`) aligned to the next *exp* boundary. *exp* must be an expression whose value is a power of two. This is equivalent to

```
(. + exp - 1) & ~(exp - 1)
```

`ALIGN` doesn't change the value of the location counter—it just does arithmetic on it. As an example, to align the output `.data` section to the next `0x2000` byte boundary after the preceding section and to set a variable within the section to the next `0x8000` boundary after the input sections:

```
SECTIONS{ . . .
  .data ALIGN(0x2000): {
    *(.data)
    variable = ALIGN(0x8000);
  }
. . . }
```

The first use of `ALIGN` in this example specifies the location of a section because it is used as the optional *start* attribute of a section definition (see Section 3.4.4 [Section Options], page 22). The second use simply defines the value of a variable.

The built-in `NEXT` is closely related to `ALIGN`.

`DEFINED`(*symbol*)

Return 1 if *symbol* is in the linker global symbol table and is defined, otherwise return 0. You can use this function to provide default values for symbols. For example, the following command-file fragment shows how to set a global symbol `begin` to the first location in the `.text` section—but if a symbol called `begin` already existed, its value is preserved:

```
SECTIONS{ . . .
  .text : {
    begin = DEFINED(begin) ? begin : . ;
    . . .
  }
. . . }
```

`NEXT`(*exp*)

Return the next unallocated address that is a multiple of *exp*. This function is closely related to `ALIGN`(*exp*); unless you use the `MEMORY` command to define discontinuous memory for the output file, the two functions are equivalent.

`SIZEOF`(*section*)

Return the size in bytes of the named *section*, if that section has been allocated. In the following example, `symbol_1` and `symbol_2` are assigned identical values:

```
SECTIONS{ . . .
  .output {
    .start = . ;
    . . .
    .end = . ;
    }
  symbol_1 = .end - .start ;
  symbol_2 = SIZEOF(.output);
. . . }
```

`SIZEOF_HEADERS`
`sizeof_headers`

Return the size in bytes of the output file's headers. You can use this number as the start address of the first section, if you choose, to facilitate paging.

3.3 Memory Layout

The linker's default configuration permits allocation of all available memory. You can override this configuration by using the `MEMORY` command. The `MEMORY` command describes the location and size of blocks of memory in the target. By using it carefully, you can describe which memory regions may be used by the linker, and which memory regions it must avoid. The linker does not shuffle sections to fit into the available regions, but does move the requested sections into the correct regions and issue errors when the regions become too full.

A command file may contain at most one use of the `MEMORY` command; however, you can define as many blocks of memory within it as you wish. The syntax is:

```
MEMORY
  {
    name (attr) : ORIGIN = origin, LENGTH = len
    . . .
  }
```

name is a name used internally by the linker to refer to the region. Any symbol name may be used. The region names are stored in a separate name space, and will not conflict with symbols, file names or section names. Use distinct names to specify multiple regions.

(*attr*) is an optional list of attributes, permitted for compatibility with the AT&T linker but not used by `ld` beyond checking that the attribute list is valid. Valid attribute lists must be made up of the characters "`LIRWX`". If you omit the attribute list, you may omit the parentheses around it as well.

origin is the start address of the region in physical memory. It is an expression that must evaluate to a constant before memory allocation is performed. The keyword `ORIGIN` may be abbreviated to `org` or `o` (but not, for example, '`ORG`').

len is the size in bytes of the region (an expression). The keyword `LENGTH` may be abbreviated to `len` or `l`.

For example, to specify that memory has two regions available for allocation—one starting at 0 for 256 kilobytes, and the other starting at `0x40000000` for four megabytes:

```
MEMORY
  {
  rom : ORIGIN = 0, LENGTH = 256K
  ram : org = 0x40000000, l = 4M
  }
```

Once you have defined a region of memory named *mem*, you can direct specific output sections there by using a command ending in '`>`*mem*' within the `SECTIONS` command (see Section 3.4.4 [Section Options], page 22). If the combined output sections directed to a region are too big for the region, the linker will issue an error message.

3.4 Specifying Output Sections

The `SECTIONS` command controls exactly where input sections are placed into output sections, their order in the output file, and to which output sections they are allocated.

You may use at most one `SECTIONS` command in a script file, but you can have as many statements within it as you wish. Statements within the `SECTIONS` command can do one of three things:

- define the entry point;
- assign a value to a symbol;
- describe the placement of a named output section, and which input sections go into it.

You can also use the first two operations—defining the entry point and defining symbols—outside the `SECTIONS` command: see Section 3.5 [Entry Point], page 24, and see Section 3.2.6 [Assignment], page 14. They are permitted here as well for your convenience in reading the script, so that symbols and the entry point can be defined at meaningful points in your output-file layout.

If you do not use a `SECTIONS` command, the linker places each input section into an identically named output section in the order that the sections are first encountered in the input files. If all input sections are present in the first file, for example, the order of sections in the output file will match the order in the first input file.

3.4.1 Section Definitions

The most frequently used statement in the `SECTIONS` command is the *section definition*, which specifies the properties of an output section: its location, alignment, contents, fill pattern, and target memory region. Most of these specifications are optional; the simplest form of a section definition is

```
SECTIONS { . . .
  secname : {
    contents
  }
. . . }
```

secname is the name of the output section, and *contents* a specification of what goes there—for example, a list of input files or sections of input files (see Section 3.4.2 [Section Placement], page 18). As you might assume, the whitespace shown is optional. You do need the colon ':' and the braces '`{}`', however.

secname must meet the constraints of your output format. In formats which only support a limited number of sections, such as `a.out`, the name must be one of the names supported by the format (`a.out`, for example, allows only `.text`, `.data` or `.bss`). If the output format supports any number of sections, but with numbers and not names (as is the case for Oasys), the name should be supplied as a quoted numeric string. A section name may consist of any sequence of characters, but any name which does not conform to the standard `ld` symbol name syntax must be quoted. See Section 3.2.2 [Symbol Names], page 12.

The linker will not create output sections which do not have any contents. This is for convenience when referring to input sections that may or may not exist. For example,

```
.foo { *(.foo }
```

will only create a '`.foo`' section in the output file if there is a '`.foo`' section in at least one input file.

3.4.2 Section Placement

In a section definition, you can specify the contents of an output section by listing particular input files, by listing particular input-file sections, or by a combination of the two. You can also place arbitrary data in the section, and define symbols relative to the beginning of the section.

The *contents* of a section definition may include any of the following kinds of statement. You can include as many of these as you like in a single section definition, separated from one another by whitespace.

filename You may simply name a particular input file to be placed in the current output section; *all* sections from that file are placed in the current section definition. If the file name has already been mentioned in another section definition, with an explicit section name list, then only those sections which have not yet been allocated are used.

To specify a list of particular files by name:

```
.data : { afile.o bfile.o cfile.o }
```

The example also illustrates that multiple statements can be included in the contents of a section definition, since each file name is a separate statement.

filename`(` *section* `)`
filename`(` *section*`,` *section*`, . . . )`
filename`(` *section section . . .)*

You can name one or more sections from your input files, for insertion in the current output section. If you wish to specify a list of input-file sections inside the parentheses, you may separate the section names by either commas or whitespace.

`*` `(`*section*`)`
`*` `(`*section*`,` *section*`, . . .)`
`*` `(`*section section . . .)`

Instead of explicitly naming particular input files in a link control script, you can refer to *all* files from the `ld` command line: use '`*`' instead of a particular file name before the parenthesized input-file section list.

If you have already explicitly included some files by name, '`*`' refers to all *remaining* files—those whose places in the output file have not yet been defined.

For example, to copy sections `1` through `4` from an Oasys file into the `.text` section of an `a.out` file, and sections `13` and `14` into the `.data` section:

```
SECTIONS {
  .text :{
    *("1" "2" "3" "4")
  }

  .data :{
    *("13" "14")
  }
}
```

'`[` *section . . .* `]`' used to be accepted as an alternate way to specify named sections from all unallocated input files. Because some operating systems (VMS) allow brackets in file names, that notation is no longer supported.

filename`( COMMON )`
`*( COMMON )`

Specify where in your output file to place uninitialized data with this notation. `*(COMMON)` by itself refers to all uninitialized data from all input files (so far as it is not yet allocated); *filename*`(COMMON)` refers to uninitialized data from a particular

file. Both are special cases of the general mechanisms for specifying where to place input-file sections: `ld` permits you to refer to uninitialized data as if it were in an input-file section named `COMMON`, regardless of the input file's format.

For example, the following command script arranges the output file into three consecutive sections, named `.text`, `.data`, and `.bss`, taking the input for each from the correspondingly named sections of all the input files:

```
SECTIONS {
  .text : { *(.text) }
  .data : { *(.data) }
  .bss :  { *(.bss)  *(COMMON) }
}
```

The following example reads all of the sections from file `all.o` and places them at the start of output section `outputa` which starts at location `0x10000`. All of section `.input1` from file `foo.o` follows immediately, in the same output section. All of section `.input2` from `foo.o` goes into output section `outputb`, followed by section `.input1` from `foo1.o`. All of the remaining `.input1` and `.input2` sections from any files are written to output section `outputc`.

```
SECTIONS {
  outputa 0x10000 :
    {
    all.o
    foo.o (.input1)
    }
  outputb :
    {
    foo.o (.input2)
    foo1.o (.input1)
    }
  outputc :
    {
    *(.input1)
    *(.input2)
    }
}
```

3.4.3 Section Data Expressions

The foregoing statements arrange, in your output file, data originating from your input files. You can also place data directly in an output section from the link command script. Most of these additional statements involve expressions; see Section 3.2 [Expressions], page 11. Although these statements are shown separately here for ease of presentation, no such segregation is needed within a section definition in the `SECTIONS` command; you can intermix them freely with any of the statements we've just described.

`CREATE_OBJECT_SYMBOLS`

Create a symbol for each input file in the current section, set to the address of the first byte of data written from that input file. For instance, with `a.out` files it is conventional to have a symbol for each input file. You can accomplish this by defining the output `.text` section as follows:

```
SECTIONS {
  .text 0x2020 :
     {
    CREATE_OBJECT_SYMBOLS
    *(.text)
    _etext = ALIGN(0x2000);
    }
  . . .
}
```

If `sample.ld` is a file containing this script, and `a.o`, `b.o`, `c.o`, and `d.o` are four input files with contents like the following—

```
/* a.c */

afunction() { }
int adata=1;
int abss;
```

'`ld -M -T sample.ld a.o b.o c.o d.o`' would create a map like this, containing symbols matching the object file names:

```
00000000 A __DYNAMIC
00004020 B _abss
00004000 D _adata
00002020 T _afunction
00004024 B _bbss
00004008 D _bdata
00002038 T _bfunction
00004028 B _cbss
00004010 D _cdata
00002050 T _cfunction
0000402c B _dbss
00004018 D _ddata
00002068 T _dfunction
00004020 D _edata
00004030 B _end
00004000 T _etext
00002020 t a.o
00002038 t b.o
00002050 t c.o
00002068 t d.o
```

symbol `=` *expression* `;`
symbol *f*`=` *expression* `;`

symbol is any symbol name (see Section 3.2.2 [Symbols], page 12). "*f*=" refers to any of the operators `&= += -= *= /=` which combine arithmetic and assignment.

When you assign a value to a symbol within a particular section definition, the value is relative to the beginning of the section (see Section 3.2.6 [Assignment], page 14). If you write

```
SECTIONS {
  abs = 14 ;
  . . .
  .data : { . . . rel = 14 ; . . . }
  abs2 = 14 + ADDR(.data);
  . . .
}
```

`abs` and `rel` do not have the same value; `rel` has the same value as `abs2`.

`BYTE(`*expression*`)`
`SHORT(`*expression*`)`
`LONG(`*expression*`)`
`QUAD(`*expression*`)`

By including one of these four statements in a section definition, you can explicitly place one, two, four, or eight bytes (respectively) at the current address of that section. `QUAD` is only supported when using a 64 bit host or target.

Multiple-byte quantities are represented in whatever byte order is appropriate for the output file format (see Chapter 5 [BFD], page 29).

`FILL(`*expression*`)`

Specify the "fill pattern" for the current section. Any otherwise unspecified regions of memory within the section (for example, regions you skip over by assigning a new value to the location counter '.') are filled with the two least significant bytes from the *expression* argument. A `FILL` statement covers memory locations *after* the point it occurs in the section definition; by including more than one `FILL` statement, you can have different fill patterns in different parts of an output section.

3.4.4 Optional Section Attributes

Here is the full syntax of a section definition, including all the optional portions:

```
SECTIONS {
. . .
secname start BLOCK(align) (NOLOAD) : AT ( ldadr )
  { contents } >region =fill
. . .
}
```

secname and *contents* are required. See Section 3.4.1 [Section Definition], page 18, and see Section 3.4.2 [Section Placement], page 18 for details on *contents*. The remaining elements—*start*, `BLOCK(`*align*`)`, `(NOLOAD)`, `AT ( `*ldadr*` )`, `>`*region*, and `=`*fill*—are all optional.

start You can force the output section to be loaded at a specified address by specifying *start* immediately following the section name. *start* can be represented as any expression. The following example generates section *output* at location `0x40000000`:

```
SECTIONS {
  . . .
  output 0x40000000: {
    . . .
    }
  . . .
```

```
}
```

`BLOCK(`*align*`)`
: You can include `BLOCK()` specification to advance the location counter `.` prior to the beginning of the section, so that the section will begin at the specified alignment. *align* is an expression.

`(NOLOAD)`
: Use '`(NOLOAD)`' to prevent a section from being loaded into memory each time it is accessed. For example, in the script sample below, the `ROM` segment is addressed at memory location '`0`' and does not need to be loaded into each object file:

```
SECTIONS {
  ROM  0  (NOLOAD)  : { . . . }
  . . .
}
```

`AT ( `*ldadr*` )`
: The expression *ldadr* that follows the `AT` keyword specifies the load address of the section. The default (if you do not use the `AT` keyword) is to make the load address the same as the relocation address. This feature is designed to make it easy to build a ROM image. For example, this `SECTIONS` definition creates two output sections: one called '`.text`', which starts at `0x1000`, and one called '`.mdata`', which is loaded at the end of the '`.text`' section even though its relocation address is `0x2000`. The symbol `_data` is defined with the value `0x2000`:

```
SECTIONS
  {
  .text 0x1000 : { *(.text) _etext = . ; }
  .mdata 0x2000 :
    AT ( ADDR(.text) + SIZEOF ( .text ) )
    { _data = . ; *(.data); _edata = . ;  }
  .bss 0x3000 :
    { _bstart = . ;  *(.bss) *(COMMON) ; _bend = . ;}
}
```

The run-time initialization code (for C programs, usually `crt0`) for use with a ROM generated this way has to include something like the following, to copy the initialized data from the ROM image to its runtime address:

```
char *src = _etext;
char *dst = _data;

/* ROM has data at end of text; copy it. */
while (dst < _edata) {
  *dst++ = *src++;
}

/* Zero bss */
for (dst = _bstart; dst< _bend; dst++)
  *dst = 0;
```

`>`*region*
: Assign this section to a previously defined region of memory. See Section 3.3 [MEMORY], page 17.

`=`*fill* Including `=`*fill* in a section definition specifies the initial fill value for that section. You may use any expression to specify *fill*. Any unallocated holes in the current output section when written to the output file will be filled with the two least significant bytes of the value, repeated as necessary. You can also change the fill value with a `FILL` statement in the *contents* of a section definition.

3.5 The Entry Point

The linker command language includes a command specifically for defining the first executable instruction in an output file (its *entry point*). Its argument is a symbol name:

`ENTRY(`*symbol*`)`

Like symbol assignments, the `ENTRY` command may be placed either as an independent command in the command file, or among the section definitions within the `SECTIONS` command—whatever makes the most sense for your layout.

`ENTRY` is only one of several ways of choosing the entry point. You may indicate it in any of the following ways (shown in descending order of priority: methods higher in the list override methods lower down).

- the '`-e`' *entry* command-line option;
- the `ENTRY(`*symbol*`)` command in a linker control script;
- the value of the symbol `start`, if present;
- the value of the symbol `_main`, if present;
- the address of the first byte of the `.text` section, if present;
- The address `0`.

For example, you can use these rules to generate an entry point with an assignment statement: if no symbol `start` is defined within your input files, you can simply define it, assigning it an appropriate value—

```
start = 0x2020;
```

The example shows an absolute address, but you can use any expression. For example, if your input object files use some other symbol-name convention for the entry point, you can just assign the value of whatever symbol contains the start address to `start`:

```
start = other_symbol ;
```

3.6 Option Commands

The command language includes a number of other commands that you can use for specialized purposes. They are similar in purpose to command-line options.

`CONSTRUCTORS`

This command ties up C`++` style constructor and destructor records. The details of the constructor representation vary from one object format to another, but usually lists of constructors and destructors appear as special sections. The `CONSTRUCTORS` command specifies where the linker is to place the data from these sections, relative to the rest of the linked output. Constructor data is marked by the symbol `__CTOR_LIST__` at the

start, and `__CTOR_LIST_END` at the end; destructor data is bracketed similarly, between `__DTOR_LIST__` and `__DTOR_LIST_END`. (The compiler must arrange to actually run this code; GNU C++ calls constructors from a subroutine `__main`, which it inserts automatically into the startup code for `main`, and destructors from `_exit`.)

`FLOAT`
`NOFLOAT` These keywords were used in some older linkers to request a particular math subroutine library. `ld` doesn't use the keywords, assuming instead that any necessary subroutines are in libraries specified using the general mechanisms for linking to archives; but to permit the use of scripts that were written for the older linkers, the keywords `FLOAT` and `NOFLOAT` are accepted and ignored.

`FORCE_COMMON_ALLOCATION`
This command has the same effect as the '`-d`' command-line option: to make `ld` assign space to common symbols even if a relocatable output file is specified ('`-r`').

`INPUT (` *file*`,` *file*`, . . . )`
`INPUT (` *file* *file* `. . . )`
Use this command to include binary input files in the link, without including them in a particular section definition. Specify the full name for each *file*, including '`.a`' if required.

`ld` searches for each *file* through the archive-library search path, just as for files you specify on the command line. See the description of '`-L`' in Section 2.1 [Command Line Options], page 3.

`GROUP (` *file*`,` *file*`, . . . )`
`GROUP (` *file* *file* `. . . )`
This command is like `INPUT`, except that the named files should all be archives, and they are searched repeatedly until no new undefined references are created. See the description of '`-(`' in Section 2.1 [Command Line Options], page 3.

`OUTPUT (` *filename* `)`
Use this command to name the link output file *filename*. The effect of `OUTPUT(`*filename*`)` is identical to the effect of '`-o` *filename*', which overrides it. You can use this command to supply a default output-file name other than `a.out`.

`OUTPUT_ARCH (` *bfdname* `)`
Specify a particular output machine architecture, with one of the names used by the BFD back-end routines (see Chapter 5 [BFD], page 29). This command is often unnecessary; the architecture is most often set implicitly by either the system BFD configuration or as a side effect of the `OUTPUT_FORMAT` command.

`OUTPUT_FORMAT (` *bfdname* `)`
When `ld` is configured to support multiple object code formats, you can use this command to specify a particular output format. *bfdname* is one of the names used by the BFD back-end routines (see Chapter 5 [BFD], page 29). The effect is identical to the effect of the '`-oformat`' command-line option. This selection affects only the output file; the related command `TARGET` affects primarily input files.

`SEARCH_DIR (` *path* `)`
Add *path* to the list of paths where `ld` looks for archive libraries. `SEARCH_DIR(`*path*`)` has the same effect as '`-L`*path*' on the command line.

`STARTUP ( filename )`
: Ensure that *filename* is the first input file used in the link process.

`TARGET ( format )`
: When `ld` is configured to support multiple object code formats, you can use this command to change the input-file object code format (like the command-line option '`-b`' or its synonym '`-format`'). The argument *format* is one of the strings used by BFD to name binary formats. If `TARGET` is specified but `OUTPUT_FORMAT` is not, the last `TARGET` argument is also used as the default format for the `ld` output file. See Chapter 5 [BFD], page 29.

  If you don't use the `TARGET` command, `ld` uses the value of the environment variable `GNUTARGET`, if available, to select the output file format. If that variable is also absent, `ld` uses the default format configured for your machine in the BFD libraries.

4 Machine Dependent Features

`ld` has additional features on some platforms; the following sections describe them. Machines where `ld` has no additional functionality are not listed.

4.1 `ld` and the H8/300

For the H8/300, `ld` can perform these global optimizations when you specify the '`-relax`' command-line option.

relaxing address modes
: `ld` finds all `jsr` and `jmp` instructions whose targets are within eight bits, and turns them into eight-bit program-counter relative `bsr` and `bra` instructions, respectively.

synthesizing instructions
: `ld` finds all `mov.b` instructions which use the sixteen-bit absolute address form, but refer to the top page of memory, and changes them to use the eight-bit address form. (That is: the linker turns '`mov.b @`*aa*`:16`' into '`mov.b @`*aa*`:8`' whenever the address *aa* is in the top page of memory).

4.2 `ld` and the Intel 960 family

You can use the '`-A`*architecture*' command line option to specify one of the two-letter names identifying members of the 960 family; the option specifies the desired output target, and warns of any incompatible instructions in the input files. It also modifies the linker's search strategy for archive libraries, to support the use of libraries specific to each particular architecture, by including in the search loop names suffixed with the string identifying the architecture.

For example, if your `ld` command line included '`-ACA`' as well as '`-ltry`', the linker would look (in its built-in search paths, and in any paths you specify with '`-L`') for a library with the names

```
try
libtry.a
tryca
libtryca.a
```

The first two possibilities would be considered in any event; the last two are due to the use of '`-ACA`'.

You can meaningfully use '`-A`' more than once on a command line, since the 960 architecture family allows combination of target architectures; each use will add another pair of name variants to search for when '`-l`' specifies a library.

`ld` supports the '`-relax`' option for the i960 family. If you specify '`-relax`', `ld` finds all `balx` and `calx` instructions whose targets are within 24 bits, and turns them into 24-bit program-counter relative `bal` and `cal` instructions, respectively. `ld` also turns `cal` instructions into `bal` instructions when it determines that the target subroutine is a leaf routine (that is, the target subroutine does not itself call any subroutines).

5 BFD

The linker accesses object and archive files using the BFD libraries. These libraries allow the linker to use the same routines to operate on object files whatever the object file format. A different object file format can be supported simply by creating a new BFD back end and adding it to the library. To conserve runtime memory, however, the linker and associated tools are usually configured to support only a subset of the object file formats available. You can use `objdump -i` (see section "objdump" in *The GNU Binary Utilities*) to list all the formats available for your configuration.

As with most implementations, BFD is a compromise between several conflicting requirements. The major factor influencing BFD design was efficiency: any time used converting between formats is time which would not have been spent had BFD not been involved. This is partly offset by abstraction payback; since BFD simplifies applications and back ends, more time and care may be spent optimizing algorithms for a greater speed.

One minor artifact of the BFD solution which you should bear in mind is the potential for information loss. There are two places where useful information can be lost using the BFD mechanism: during conversion and during output. See Section 5.1.1 [BFD information loss], page 29.

5.1 How it works: an outline of BFD

When an object file is opened, BFD subroutines automatically determine the format of the input object file. They then build a descriptor in memory with pointers to routines that will be used to access elements of the object file's data structures.

As different information from the the object files is required, BFD reads from different sections of the file and processes them. For example, a very common operation for the linker is processing symbol tables. Each BFD back end provides a routine for converting between the object file's representation of symbols and an internal canonical format. When the linker asks for the symbol table of an object file, it calls through a memory pointer to the routine from the relevant BFD back end which reads and converts the table into a canonical form. The linker then operates upon the canonical form. When the link is finished and the linker writes the output file's symbol table, another BFD back end routine is called to take the newly created symbol table and convert it into the chosen output format.

5.1.1 Information Loss

Information can be lost during output. The output formats supported by BFD do not provide identical facilities, and information which can be described in one form has nowhere to go in another format. One example of this is alignment information in `b.out`. There is nowhere in an `a.out` format file to store alignment information on the contained data, so when a file is linked from `b.out` and an `a.out` image is produced, alignment information will not propagate to the output file. (The linker will still use the alignment information internally, so the link is performed correctly).

Another example is COFF section names. COFF files may contain an unlimited number of sections, each one with a textual section name. If the target of the link is a format which does not have many sections (e.g., `a.out`) or has sections without names (e.g., the Oasys format), the link cannot be done simply. You can circumvent this problem by describing the desired input-to-output section mapping with the linker command language.

Information can be lost during canonicalization. The BFD internal canonical form of the external formats is not exhaustive; there are structures in input formats for which there is no direct representation internally. This means that the BFD back ends cannot maintain all possible data richness through the transformation between external to internal and back to external formats.

This limitation is only a problem when an application reads one format and writes another. Each BFD back end is responsible for maintaining as much data as possible, and the internal BFD canonical form has structures which are opaque to the BFD core, and exported only to the back ends. When a file is read in one format, the canonical form is generated for BFD and the application. At the same time, the back end saves away any information which may otherwise be lost. If the data is then written back in the same format, the back end routine will be able to use the canonical form provided by the BFD core as well as the information it prepared earlier. Since there is a great deal of commonality between back ends, there is no information lost when linking or copying big endian COFF to little endian COFF, or `a.out` to `b.out`. When a mixture of formats is linked, the information is only lost from the files whose format differs from the destination.

5.1.2 The BFD canonical object-file format

The greatest potential for loss of information occurs when there is the least overlap between the information provided by the source format, that stored by the canonical format, and that needed by the destination format. A brief description of the canonical form may help you understand which kinds of data you can count on preserving across conversions.

files Information stored on a per-file basis includes target machine architecture, particular implementation format type, a demand pageable bit, and a write protected bit. Information like Unix magic numbers is not stored here—only the magic numbers' meaning, so a `ZMAGIC` file would have both the demand pageable bit and the write protected text bit set. The byte order of the target is stored on a per-file basis, so that big- and little-endian object files may be used with one another.

sections Each section in the input file contains the name of the section, the section's original address in the object file, size and alignment information, various flags, and pointers into other BFD data structures.

symbols Each symbol contains a pointer to the information for the object file which originally defined it, its name, its value, and various flag bits. When a BFD back end reads in a symbol table, it relocates all symbols to make them relative to the base of the section where they were defined. Doing this ensures that each symbol points to its containing section. Each symbol also has a varying amount of hidden private data for the BFD back end. Since the symbol points to the original file, the private data format for that symbol is accessible. `ld` can operate on a collection of symbols of wildly different formats without problems.

Normal global and simple local symbols are maintained on output, so an output file (no matter its format) will retain symbols pointing to functions and to global, static, and common variables. Some symbol information is not worth retaining; in `a.out`, type information is stored in the symbol table as long symbol names. This information would be useless to most COFF debuggers; the linker has command line switches to allow users to throw it away.

There is one word of type information within the symbol, so if the format supports symbol type information within symbols (for example, COFF, IEEE, Oasys) and the type is simple enough to fit within one word (nearly everything but aggregates), the information will be preserved.

relocation level

Each canonical BFD relocation record contains a pointer to the symbol to relocate to, the offset of the data to relocate, the section the data is in, and a pointer to a relocation type descriptor. Relocation is performed by passing messages through the relocation type descriptor and the symbol pointer. Therefore, relocations can be performed on output data using a relocation method that is only available in one of the input formats. For instance, Oasys provides a byte relocation format. A relocation record requesting this relocation type would point indirectly to a routine to perform this, so the relocation may be performed on a byte being written to a 68k COFF file, even though 68k COFF has no such relocation type.

line numbers

Object formats can contain, for debugging purposes, some form of mapping between symbols, source line numbers, and addresses in the output file. These addresses have to be relocated along with the symbol information. Each symbol with an associated list of line number records points to the first record of the list. The head of a line number list consists of a pointer to the symbol, which allows finding out the address of the function whose line number is being described. The rest of the list is made up of pairs: offsets into the section and line numbers. Any format which can simply derive this information can pass it successfully between formats (COFF, IEEE and Oasys).

Appendix A MRI Compatible Script Files

To aid users making the transition to GNU `ld` from the MRI linker, `ld` can use MRI compatible linker scripts as an alternative to the more general-purpose linker scripting language described in Chapter 3 [Command Language], page 11. MRI compatible linker scripts have a much simpler command set than the scripting language otherwise used with `ld`. GNU `ld` supports the most commonly used MRI linker commands; these commands are described here.

In general, MRI scripts aren't of much use with the `a.out` object file format, since it only has three sections and MRI scripts lack some features to make use of them.

You can specify a file containing an MRI-compatible script using the '`-c`' command-line option.

Each command in an MRI-compatible script occupies its own line; each command line starts with the keyword that identifies the command (though blank lines are also allowed for punctuation). If a line of an MRI-compatible script begins with an unrecognized keyword, `ld` issues a warning message, but continues processing the script.

Lines beginning with '`*`' are comments.

You can write these commands using all upper-case letters, or all lower case; for example, '`chip`' is the same as '`CHIP`'. The following list shows only the upper-case form of each command.

`ABSOLUTE` *secname*

`ABSOLUTE` *secname*`,` *secname*`, . . .` *secname*
: Normally, `ld` includes in the output file all sections from all the input files. However, in an MRI-compatible script, you can use the `ABSOLUTE` command to restrict the sections that will be present in your output program. If the `ABSOLUTE` command is used at all in a script, then only the sections named explicitly in `ABSOLUTE` commands will appear in the linker output. You can still use other input sections (whatever you select on the command line, or using `LOAD`) to resolve addresses in the output file.

`ALIAS` *out-secname*`,` *in-secname*
: Use this command to place the data from input section *in-secname* in a section called *out-secname* in the linker output file.

  in-secname may be an integer.

`BASE` *expression*
: Use the value of *expression* as the lowest address (other than absolute addresses) in the output file.

`CHIP` *expression*

`CHIP` *expression*`,` *expression*
: This command does nothing; it is accepted only for compatibility.

`END`
: This command does nothing whatever; it's only accepted for compatibility.

`FORMAT` *output-format*
: Similar to the `OUTPUT_FORMAT` command in the more general linker language, but restricted to one of these output formats:

  1. S-records, if *output-format* is '`S`'
  2. IEEE, if *output-format* is '`IEEE`'
  3. COFF (the '`coff-m68k`' variant in BFD), if *output-format* is '`COFF`'

`LIST` *anything . . .*

Print (to the standard output file) a link map, as produced by the `ld` command-line option '`-M`'.

The keyword `LIST` may be followed by anything on the same line, with no change in its effect.

`LOAD` *filename*

`LOAD` *filename, filename, . . . filename*

Include one or more object file *filename* in the link; this has the same effect as specifying *filename* directly on the `ld` command line.

`NAME` *output-name*

output-name is the name for the program produced by `ld`; the MRI-compatible command `NAME` is equivalent to the command-line option '`-o`' or the general script language command `OUTPUT`.

`ORDER` *secname, secname, . . . secname*

`ORDER` *secname secname secname*

Normally, `ld` orders the sections in its output file in the order in which they first appear in the input files. In an MRI-compatible script, you can override this ordering with the `ORDER` command. The sections you list with `ORDER` will appear first in your output file, in the order specified.

`PUBLIC` *name=expression*

`PUBLIC` *name,expression*

`PUBLIC` *name expression*

Supply a value (*expression*) for external symbol *name* used in the linker input files.

`SECT` *secname, expression*

`SECT` *secname=expression*

`SECT` *secname expression*

You can use any of these three forms of the `SECT` command to specify the start address (*expression*) for section *secname*. If you have more than one `SECT` statement for the same *secname*, only the *first* sets the start address.

Index

*

*(COMMON) 19
*(*section*) 19

-

-(..... 10
--verbose 8
-A*arch* 4
-b *format* 4
-Bstatic 4
-c *MRI-cmdfile* 4
-d 4
-dc 5
-defsym *symbol=exp* 5
-dp 5
-dynamic-linker *file* 5
-e *entry* 5
-F 5
-format 5
-g 5
-G 5
-help 5
-i 5
-l*archive* 5
-L*dir* 5
-M 6
-m *emulation* 6
-Map 6
-n 6
-N 6
-noinhibit-exec 6
-o *output* 6
-oformat 6
-r 7
-R *file* 6
-relax 6
-relax on i960 27
-rpath 7
-s 7
-S 7
-soname 7
-t 8
-T *script* 8
-Tbss *org* 8
-Tdata *org* 8
-traditional-format 8
-Ttext *org* 8
-u *symbol* 8
-Ur 8
-v 8
-V 8
-version 8
-warn-comon 8
-warn-once 9
-x 9
-X 9
-y *symbol* 10

.

. 12

;

; 14

=

=*fill* 23

[

[*section*...], not supported 19

"

" 12

>

>*region* 23

0

0x 12

A

ABSOLUTE (MRI) 33
absolute and relocatable symbols 14
ABSOLUTE(*exp*) 15
ADDR(*section*) 15
ALIAS (MRI) 33
ALIGN(*exp*) 15
aligning sections 23
allocating memory 17
architectures 4
archive files, from cmd line 5
arithmetic 11

arithmetic operators 13
assignment in scripts 14
assignment, in section defn 21
AT (*ldadr*) 23

B

back end 29
BASE (MRI) 33
BFD canonical format 30
BFD requirements 29
binary input files 25
binary input format 4
BLOCK(*align*) 23
BYTE(*expression*) 22

C

C++ constructors, arranging in link 24
CHIP (MRI) 33
combining symbols, warnings on 8
command files 11
command line 3
commands, fundamental 11
comments 11
common allocation 4, 25
commons in output 19
compatibility, MRI 4
constructors 8
CONSTRUCTORS 24
constructors, arranging in link 24
contents of a section 18
CREATE_OBJECT_SYMBOLS 20
current output location 12

D

dbx 8
decimal integers 12
default input format 10
DEFINED(*symbol*) 16
deleting local symbols 9
direct output 22
discontinuous memory 17
dot 12
dynamic linker, from command line 5

E

emulation 6
END (MRI) 33
entry point, defaults 24
entry point, from command line 5
ENTRY(*symbol*) 24
expression evaluation order 13
expression syntax 11
expression, absolute 15
expressions in a section 20

F

filename 19
filename symbols 20
filename(*section*) 19
files and sections, section defn 19
files, including in output sections 19
fill pattern, entire section 23
FILL(*expression*) 22
first input file 26
first instruction 24
FLOAT 25
FORCE_COMMON_ALLOCATION 25
FORMAT (MRI) 33
format, output file 25
functions in expression language 15
fundamental script commands 11

G

GNU linker 1
GNUTARGET 10, 26
GROUP (*files*) 25
grouping input files 25
groups of archives 10

H

H8/300 support 27
header size 16
help 5
hexadecimal integers 12
holes 12
holes, filling 22

I

i960 support 27
incremental link 5
INPUT (*files*) 25
input file format 26
input filename symbols 20
input files, displaying 8
input files, section defn 19
input format 4
input sections to output section 19
integer notation 12
integer suffixes 12
internal object-file format 30

K

K and M integer suffixes ... 12

L

`l =` ... 17
L, deleting symbols beginning ... 9
layout of output file ... 11
lazy evaluation ... 13
`len =` ... 17
`LENGTH =` ... 17
link map ... 6
`LIST` (MRI) ... 34
`LOAD` (MRI) ... 34
load address, specifying ... 23
loading, preventing ... 23
local symbols, deleting ... 9
location counter ... 12
`LONG(`*expression*`)` ... 22

M

M and K integer suffixes ... 12
machine architecture, output ... 25
machine dependencies ... 27
`MEMORY` ... 17
memory region attributes ... 17
memory regions and sections ... 23
MRI compatibility ... 33

N

`NAME` (MRI) ... 34
names ... 12
naming memory regions ... 17
naming output sections ... 18
naming the output file ... 6, 25
negative integers ... 12
`NEXT(`*exp*`)` ... 16
`NMAGIC` ... 6
`NOFLOAT` ... 25
`NOLOAD` ... 23
`Non constant expression` ... 15

O

`o =` ... 17
`objdump -i` ... 29
object file management ... 29
object files ... 3
object formats available ... 29
object size ... 5
octal integers ... 12
`OMAGIC` ... 6
opening object files ... 29
Operators for arithmetic ... 13
options ... 3
`ORDER` (MRI) ... 34
`org =` ... 17
`ORIGIN =` ... 17
`OUTPUT (` *filename* `)` ... 25
output file after errors ... 6
output file layout ... 11
`OUTPUT_ARCH (` *bfdname* `)` ... 25
`OUTPUT_FORMAT (` *bfdname* `)` ... 25

P

partial link ... 7
path for libraries ... 25
precedence in expressions ... 13
prevent unnecessary loading ... 23
provide ... 15
`PUBLIC` (MRI) ... 34

Q

`QUAD(`*expression*`)` ... 22
quoted symbol names ... 12

R

read-only text ... 6
read/write from cmd line ... 6
regions of memory ... 17
relaxing addressing modes ... 6
relaxing on H8/300 ... 27
relaxing on i960 ... 27
relocatable and absolute symbols ... 14
relocatable output ... 7
requirements for BFD ... 29
retaining specified symbols ... 7
rounding up location counter ... 15
runtime library name ... 7
runtime library search path ... 7

S

scaled integers ... 12
script files ... 8
search directory, from cmd line ... 5
search path, libraries ... 25
`SEARCH_DIR (` *path* `)` ... 25
`SECT` (MRI) ... 34
section address ... 15, 22
section alignment ... 23
section definition ... 18
section defn, full syntax ... 22

section fill pattern 23
section size 16
section start 22
section, assigning to memory region 23
SECTIONS 17
segment origins, cmd line 8
semicolon 14
SHORT(*expression*) 22
SIZEOF(*section*) 16
sizeof_headers 16
SIZEOF_HEADERS 16
specify load address 23
standard Unix system 3
start address, section 22
start of execution 24
STARTUP (*filename*) 26
strip all symbols 7
strip debugger symbols 7
stripping all but some symbols 7
suffixes for integers 12
symbol = *expression* ; 21
symbol defaults 16
symbol definition, scripts 14
symbol f= *expression* ; 21
symbol names 12
symbol tracing 10
symbol-only input 6
symbols, from command line 5
symbols, relocatable and absolute 14
symbols, retaining selectively 7
synthesizing linker 6
synthesizing on H8/300 27

T

TARGET (*format*) 26
traditional format 8

U

unallocated address, next 16
undefined symbol 8
undefined symbols, warnings on 9
uninitialized data 19
unspecified memory 22
usage 5

V

variables, defining 14
verbose 8
version 8

W

warnings, on combining symbols 8
warnings, on undefined symbols 9
what is this? 1

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 and
cmsl10 at 10.95pt
are used for emphasis.

Table of Contents

1 Overview 1

2 Invocation 3

2.1 Command Line Options 3
2.2 Environment Variables 10

3 Command Language 11

3.1 Linker Scripts 11
3.2 Expressions 11
3.2.1 Integers 12
3.2.2 Symbol Names 12
3.2.3 The Location Counter 12
3.2.4 Operators 13
3.2.5 Evaluation 13
3.2.6 Assignment: Defining Symbols 14
3.2.7 Arithmetic Functions 15
3.3 Memory Layout 17
3.4 Specifying Output Sections 17
3.4.1 Section Definitions 18
3.4.2 Section Placement 18
3.4.3 Section Data Expressions 20
3.4.4 Optional Section Attributes 22
3.5 The Entry Point 24
3.6 Option Commands 24

4 Machine Dependent Features 27

4.1 `ld` and the H8/300 27
4.2 `ld` and the Intel 960 family 27

5 BFD 29

5.1 How it works: an outline of BFD 29
5.1.1 Information Loss 29
5.1.2 The BFD canonical object-file format 30

Appendix A MRI Compatible Script Files 33

Index 35