

Using `as`

The GNU Assembler

Version 2.9.5

The Free Software Foundation Inc. thanks The Nice Computer Company of Australia for loaning Dean Elsner to write the first (Vax) version of `as` for Project GNU. The proprietors, management and staff of TNCCA thank FSF for distracting the boss while they got some work done.

Dean Elsner, Jay Fenlason & friends

Using as
Edited by Cygnus Support

Copyright © 1991, 92, 93, 94, 95, 96, 97, 1998 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 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

This manual is a user guide to the GNU assembler `as`.

Here is a brief summary of how to invoke `as`. For details, see [Chapter 2 \[Comand-Line Options\]](#), [page 9](#).

```
as [ -a[cdhlms] [=file] ] [ -D ] [ --defsym sym=val ]
[ -f ] [ --gstabs ] [ --gdwarf2 ] [ --help ] [ -I dir ] [ -J ] [ -K ] [ -L ]
[ --keep-locals ] [ -o objfile ] [ -R ] [ --statistics ] [ -v ]
[ -version ] [ --version ] [ -W ] [ --warn ] [ --fatal-warnings ]
[ -w ] [ -x ] [ -Z ]
[ -mbig-endian | -mlittle-endian ]
[ -m[arm]1 | -m[arm]2 | -m[arm]250 | -m[arm]3 | -m[arm]6 | -m[arm]60 |
  -m[arm]600 | -m[arm]610 | -m[arm]620 | -m[arm]7[t][[d]m[i]][fe] | -m[arm]70 |
  -m[arm]700 | -m[arm]710[c] | -m[arm]7100 | -m[arm]7500 | -m[arm]8 |
  -m[arm]810 | -m[arm]9 | -m[arm]920 | -m[arm]920t | -m[arm]9tdmi |
  -mstrongarm | -mstrongarm110 | -mstrongarm1100 ]
[ -m[arm]v2 | -m[arm]v2a | -m[arm]v3 | -m[arm]v3m | -m[arm]v4 | -m[arm]v4t |
  -m[arm]v5 | -[arm]v5t ]
[ -mthumb | -mall ]
[ -mfpa10 | -mfpa11 | -mfpe-old | -mno-fpu ]
[ -EB | -EL ]
[ -mapcs-32 | -mapcs-26 | -mapcs-float | -mapcs-reentrant ]
[ -mthumb-interwork ]
[ -moabi ]
[ -k ]
[ -O ]
[ -O | -n | -N ]
[ -mb | -me ]
[ -Av6 | -Av7 | -Av8 | -Asparclet | -Asparclite
  -Av8plus | -Av8plusa | -Av9 | -Av9a ]
[ -xarch=v8plus | -xarch=v8plusa ] [ -bump ] [ -32 | -64 ]
[ -ACA | -ACA_A | -ACB | -ACC | -AKA | -AKB | -AKC | -AMC ]
[ -b ] [ -no-relax ]
[ -l ] [ -m68000 | -m68010 | -m68020 | ... ]
[ -jsri2bsr ] [ -sifilter ] [ -relax ]
[ -nocpp ] [ -EL ] [ -EB ] [ -G num ] [ -mcpu=CPU ]
[ -mips1 ] [ -mips2 ] [ -mips3 ] [ -m4650 ] [ -no-m4650 ]
[ --trap ] [ --break ]
[ --emulation=name ]
[ -- | files ... ]
```

`-a[cdhlms]`

Turn on listings, in any of a variety of ways:

<code>-ac</code>	omit false conditionals
<code>-ad</code>	omit debugging directives
<code>-ah</code>	include high-level source
<code>-al</code>	include assembly

- `-am` include macro expansions
 - `-an` omit forms processing
 - `-as` include symbols
 - `=file` set the name of the listing file
- You may combine these options; for example, use `-aln` for assembly listing without forms processing. The `=file` option, if used, must be the last one. By itself, `-a` defaults to `-ahls`.
- `-D` Ignored. This option is accepted for script compatibility with calls to other assemblers.
 - `--defsym sym=value`
Define the symbol *sym* to be *value* before assembling the input file. *value* must be an integer constant. As in C, a leading `'0x'` indicates a hexadecimal value, and a leading `'0'` indicates an octal value.
 - `-f` “fast”—skip whitespace and comment preprocessing (assume source is compiler output).
 - `--gstabs` Generate stabs debugging information for each assembler line. This may help debugging assembler code, if the debugger can handle it.
 - `--gdwarf2`
Generate DWARF2 debugging information for each assembler line. This may help debugging assembler code, if the debugger can handle it.
 - `--help` Print a summary of the command line options and exit.
 - `-I dir` Add directory *dir* to the search list for `.include` directives.
 - `-J` Don't warn about signed overflow.
 - `-K` Issue warnings when difference tables altered for long displacements.
 - `-L`
 - `--keep-locals`
Keep (in the symbol table) local symbols. On traditional a.out systems these start with `'L'`, but different systems have different local label prefixes.
 - `-o objfile` Name the object-file output from `as` *objfile*.
 - `-R` Fold the data section into the text section.
 - `--statistics`
Print the maximum space (in bytes) and total time (in seconds) used by assembly.
 - `--strip-local-absolute`
Remove local absolute symbols from the outgoing symbol table.
 - `-v`
 - `-version` Print the `as` version.
 - `--version`
Print the `as` version and exit.

```

-W
--no-warn          Suppress warning messages.
--fatal-warnings  Treat warnings as errors.
--warn            Don't suppress warning messages or treat them as errors.
-w               Ignored.
-x               Ignored.
-Z               Generate an object file even after errors.
-- | files ...    Standard input, or source files to assemble.

```

The following options are available when as is configured for an ARC processor.

```

-mbig-endian      Generate "big endian" format output.
-mlittle-endian  Generate "little endian" format output.

```

The following options are available when as is configured for the ARM processor family.

```

-m[arm] [1|2|3|6|7|8|9] [...]
    Specify which ARM processor variant is the target.
-m[arm]v[2|2a|3|3m|4|4t|5|5t]
    Specify which ARM architecture variant is used by the target.
-mthumb | -mll   Enable or disable Thumb only instruction decoding.
-mfpa10 | -mfpa11 | -mfpe-old | -mno-fpu
    Select which Floating Point architecture is the target.
-mapcs-32 | -mapcs-26 | -mapcs-float | -mapcs-reentrant | -moabi
    Select which procedure calling convention is in use.
-EB | -EL       Select either big-endian (-EB) or little-endian (-EL) output.
-mthumb-interwork
    Specify that the code has been generated with interworking between Thumb
    and ARM code in mind.
-k              Specify that PIC code has been generated.

```

The following options are available when as is configured for a D10V processor.

```

-O              Optimize output by parallelizing instructions.

```

The following options are available when as is configured for a D30V processor.

```

-O              Optimize output by parallelizing instructions.

```

-n Warn when nops are generated.

-N Warn when a nop after a 32-bit multiply instruction is generated.

The following options are available when as is configured for the Intel 80960 processor.

-ACA | -ACA_A | -ACB | -ACC | -AKA | -AKB | -AKC | -AMC

Specify which variant of the 960 architecture is the target.

-b Add code to collect statistics about branches taken.

-no-relax

Do not alter compare-and-branch instructions for long displacements; error if necessary.

The following options are available when as is configured for the Motorola 68000 series.

-l Shorten references to undefined symbols, to one word instead of two.

**-m68000 | -m68008 | -m68010 | -m68020 | -m68030 | -m68040 | -m68060
| -m68302 | -m68331 | -m68332 | -m68333 | -m68340 | -mcpu32 | -m5200**

Specify what processor in the 68000 family is the target. The default is normally the 68020, but this can be changed at configuration time.

-m68881 | -m68882 | -mno-68881 | -mno-68882

The target machine does (or does not) have a floating-point coprocessor. The default is to assume a coprocessor for 68020, 68030, and cpu32. Although the basic 68000 is not compatible with the 68881, a combination of the two can be specified, since it's possible to do emulation of the coprocessor instructions with the main processor.

-m68851 | -mno-68851

The target machine does (or does not) have a memory-management unit coprocessor. The default is to assume an MMU for 68020 and up.

The following options are available when as is configured for a picoJava processor.

-mb Generate "big endian" format output.

-ml Generate "little endian" format output.

The following options are available when as is configured for the SPARC architecture:

-Av6 | -Av7 | -Av8 | -Asparclet | -Asparclite

-Av8plus | -Av8plusa | -Av9 | -Av9a

Explicitly select a variant of the SPARC architecture.

'-Av8plus' and '-Av8plusa' select a 32 bit environment. '-Av9' and '-Av9a' select a 64 bit environment.

'-Av8plusa' and '-Av9a' enable the SPARC V9 instruction set with Ultra-SPARC extensions.

-xarch=v8plus | -xarch=v8plusa

For compatibility with the Solaris v9 assembler. These options are equivalent to -Av8plus and -Av8plusa, respectively.

-bump Warn when the assembler switches to another architecture.

The following options are available when `as` is configured for a MIPS processor.

- `-G num` This option sets the largest size of an object that can be referenced implicitly with the `gp` register. It is only accepted for targets that use ECOFF format, such as a DECstation running Ultrix. The default value is 8.
- `-EB` Generate “big endian” format output.
- `-EL` Generate “little endian” format output.
- `-mips1`
- `-mips2`
- `-mips3` Generate code for a particular MIPS Instruction Set Architecture level. ‘`-mips1`’ corresponds to the R2000 and R3000 processors, ‘`-mips2`’ to the R6000 processor, and ‘`-mips3`’ to the R4000 processor.
- `-m4650`
- `-no-m4650` Generate code for the MIPS R4650 chip. This tells the assembler to accept the ‘`mad`’ and ‘`madu`’ instruction, and to not schedule ‘`nop`’ instructions around accesses to the ‘`HI`’ and ‘`LO`’ registers. ‘`-no-m4650`’ turns off this option.
- `-mcpu=CPU` Generate code for a particular MIPS cpu. This has little effect on the assembler, but it is passed by `gcc`.
- `--emulation=name` This option causes `as` to emulate `as` configured for some other target, in all respects, including output format (choosing between ELF and ECOFF only), handling of pseudo-opcodes which may generate debugging information or store symbol table information, and default endianness. The available configuration names are: ‘`mipsecoff`’, ‘`mipsel`’, ‘`mipslecoff`’, ‘`mipsbecoff`’, ‘`mipslelf`’, ‘`mipsbelf`’. The first two do not alter the default endianness from that of the primary target for which the assembler was configured; the others change the default to little- or big-endian as indicated by the ‘`b`’ or ‘`l`’ in the name. Using ‘`-EB`’ or ‘`-EL`’ will override the endianness selection in any case.

This option is currently supported only when the primary target `as` is configured for is a MIPS ELF or ECOFF target. Furthermore, the primary target or others specified with ‘`--enable-targets=...`’ at configuration time must include support for the other format, if both are to be available. For example, the Irix 5 configuration includes support for both.

Eventually, this option will support more configurations, with more fine-grained control over the assembler’s behavior, and will be supported for more processors.
- `-nocpp` `as` ignores this option. It is accepted for compatibility with the native tools.
- `--trap`
- `--no-trap`
- `--break`
- `--no-break` Control how to deal with multiplication overflow and division by zero. ‘`--trap`’ or ‘`--no-break`’ (which are synonyms) take a trap exception (and only work

for Instruction Set Architecture level 2 and higher); ‘`--break`’ or ‘`--no-trap`’ (also synonyms, and the default) take a break exception.

The following options are available when `as` is configured for an MCore processor.

- `-jsri2bsr`
- `-nojsri2bsr`
 - Enable or disable the JSRI to BSR transformation. By default this is enabled. The command line option ‘`-nojsri2bsr`’ can be used to disable it.
- `-sifilter`
- `-nosifilter`
 - Enable or disable the silicon filter behaviour. By default this is disabled. The default can be overridden by the ‘`-sifilter`’ command line option.
- `-relax`
 - Alter jump instructions for long displacements.

1.1 Structure of this Manual

This manual is intended to describe what you need to know to use GNU `as`. We cover the syntax expected in source files, including notation for symbols, constants, and expressions; the directives that `as` understands; and of course how to invoke `as`.

This manual also describes some of the machine-dependent features of various flavors of the assembler.

On the other hand, this manual is *not* intended as an introduction to programming in assembly language—let alone programming in general! In a similar vein, we make no attempt to introduce the machine architecture; we do *not* describe the instruction set, standard mnemonics, registers or addressing modes that are standard to a particular architecture. You may want to consult the manufacturer’s machine architecture manual for this information.

1.2 The GNU Assembler

GNU `as` is really a family of assemblers. If you use (or have used) the GNU assembler on one architecture, you should find a fairly similar environment when you use it on another architecture. Each version has much in common with the others, including object file formats, most assembler directives (often called *pseudo-ops*) and assembler syntax.

`as` is primarily intended to assemble the output of the GNU C compiler `gcc` for use by the linker `ld`. Nevertheless, we’ve tried to make `as` assemble correctly everything that other assemblers for the same machine would assemble. Any exceptions are documented explicitly (see [Chapter 8 \[Machine Dependencies\]](#), page 53). This doesn’t mean `as` always uses the same syntax as another assembler for the same architecture; for example, we know of several incompatible versions of 680x0 assembly language syntax.

Unlike older assemblers, `as` is designed to assemble a source program in one pass of the source file. This has a subtle impact on the `.org` directive (see [Section 7.49 \[.org\]](#), page 44).

1.3 Object File Formats

The GNU assembler can be configured to produce several alternative object file formats. For the most part, this does not affect how you write assembly language programs; but directives for debugging symbols are typically different in different file formats. See [Section 5.5 \[Symbol Attributes\]](#), page 28.

1.4 Command Line

After the program name `as`, the command line may contain options and file names. Options may appear in any order, and may be before, after, or between file names. The order of file names is significant.

`--` (two hyphens) by itself names the standard input file explicitly, as one of the files for `as` to assemble.

Except for `--` any command line argument that begins with a hyphen (`-`) is an option. Each option changes the behavior of `as`. No option changes the way another option works. An option is a `-` followed by one or more letters; the case of the letter is important. All options are optional.

Some options expect exactly one file name to follow them. The file name may either immediately follow the option's letter (compatible with older assemblers) or it may be the next command argument (GNU standard). These two command lines are equivalent:

```
as -o my-object-file.o mumble.s
as -omy-object-file.o mumble.s
```

1.5 Input Files

We use the phrase *source program*, abbreviated *source*, to describe the program input to one run of `as`. The program may be in one or more files; how the source is partitioned into files doesn't change the meaning of the source.

The source program is a concatenation of the text in all the files, in the order specified.

Each time you run `as` it assembles exactly one source program. The source program is made up of one or more files. (The standard input is also a file.)

You give `as` a command line that has zero or more input file names. The input files are read (from left file name to right). A command line argument (in any position) that has no special meaning is taken to be an input file name.

If you give `as` no file names it attempts to read one input file from the `as` standard input, which is normally your terminal. You may have to type `⌘-D` to tell `as` there is no more program to assemble.

Use `--` if you need to explicitly name the standard input file in your command line.

If the source is empty, `as` produces a small, empty object file.

Filename and Line-numbers

There are two ways of locating a line in the input file (or files) and either may be used in reporting error messages. One way refers to a line number in a physical file; the other

refers to a line number in a “logical” file. See [Section 1.7 \[Error and Warning Messages\]](#), page 8.

Physical files are those files named in the command line given to `as`.

Logical files are simply names declared explicitly by assembler directives; they bear no relation to physical files. Logical file names help error messages reflect the original source file, when `as` source is itself synthesized from other files. `as` understands the ‘#’ directives emitted by the `gcc` preprocessor. See also [Section 7.26 \[.file\]](#), page 37.

1.6 Output (Object) File

Every time you run `as` it produces an output file, which is your assembly language program translated into numbers. This file is the object file. Its default name is `a.out`, or `b.out` when `as` is configured for the Intel 80960. You can give it another name by using the `-o` option. Conventionally, object file names end with ‘.o’. The default name is used for historical reasons: older assemblers were capable of assembling self-contained programs directly into a runnable program. (For some formats, this isn’t currently possible, but it can be done for the `a.out` format.)

The object file is meant for input to the linker `ld`. It contains assembled program code, information to help `ld` integrate the assembled program into a runnable file, and (optionally) symbolic information for the debugger.

1.7 Error and Warning Messages

`as` may write warnings and error messages to the standard error file (usually your terminal). This should not happen when a compiler runs `as` automatically. Warnings report an assumption made so that `as` could keep assembling a flawed program; errors report a grave problem that stops the assembly.

Warning messages have the format

```
file_name:NNN:Warning Message Text
```

(where `NNN` is a line number). If a logical file name has been given (see [Section 7.26 \[.file\]](#), page 37) it is used for the filename, otherwise the name of the current input file is used. If a logical line number was given (see [Section 7.40 \[.line\]](#), page 41) then it is used to calculate the number printed, otherwise the actual line in the current source file is printed. The message text is intended to be self explanatory (in the grand Unix tradition).

Error messages have the format

```
file_name:NNN:FATAL:Error Message Text
```

The file name and line number are derived as for warning messages. The actual message text may be rather less explanatory because many of them aren’t supposed to happen.

2 Command-Line Options

This chapter describes command-line options available in *all* versions of the GNU assembler; see [Chapter 8 \[Machine Dependencies\]](#), page 53, for options specific to particular machine architectures.

If you are invoking `as` via the GNU C compiler (version 2), you can use the `-Wa` option to pass arguments through to the assembler. The assembler arguments must be separated from each other (and the `-Wa`) by commas. For example:

```
gcc -c -g -O -Wa,-alh,-L file.c
```

This passes two options to the assembler: `-alh` (emit a listing to standard output with high-level and assembly source) and `-L` (retain local symbols in the symbol table).

Usually you do not need to use this `-Wa` mechanism, since many compiler command-line options are automatically passed to the assembler by the compiler. (You can call the GNU compiler driver with the `-v` option to see precisely what options it passes to each compilation pass, including the assembler.)

2.1 Enable Listings: `-a[cdhlns]`

These options enable listing output from the assembler. By itself, `-a` requests high-level, assembly, and symbols listing. You can use other letters to select specific options for the list: `-ah` requests a high-level language listing, `-al` requests an output-program assembly listing, and `-as` requests a symbol table listing. High-level listings require that a compiler debugging option like `-g` be used, and that assembly listings (`-al`) be requested also.

Use the `-ac` option to omit false conditionals from a listing. Any lines which are not assembled because of a false `.if` (or `.ifdef`, or any other conditional), or a true `.if` followed by an `.else`, will be omitted from the listing.

Use the `-ad` option to omit debugging directives from the listing.

Once you have specified one of these options, you can further control listing output and its appearance using the directives `.list`, `.nolist`, `.psize`, `.eject`, `.title`, and `.sbttl`. The `-an` option turns off all forms processing. If you do not request listing output with one of the `-a` options, the listing-control directives have no effect.

The letters after `-a` may be combined into one option, *e.g.*, `-alns`.

2.2 `-D`

This option has no effect whatsoever, but it is accepted to make it more likely that scripts written for other assemblers also work with `as`.

2.3 Work Faster: `-f`

`-f` should only be used when assembling programs written by a (trusted) compiler. `-f` stops the assembler from doing whitespace and comment preprocessing on the input file(s) before assembling them. See [Section 3.1 \[Preprocessing\]](#), page 15.

Warning: if you use `-f` when the files actually need to be preprocessed (if they contain comments, for example), `as` does not work correctly.

2.4 .include search path: -I *path*

Use this option to add a *path* to the list of directories **as** searches for files specified in `.include` directives (see [Section 7.34 \[.include\]](#), page 40). You may use `-I` as many times as necessary to include a variety of paths. The current working directory is always searched first; after that, **as** searches any `-I` directories in the same order as they were specified (left to right) on the command line.

2.5 Difference Tables: -K

as sometimes alters the code emitted for directives of the form `‘.word sym1-sym2’`; see [Section 7.77 \[.word\]](#), page 52. You can use the `-K` option if you want a warning issued when this is done.

2.6 Include Local Labels: -L

Labels beginning with `‘L’` (upper case only) are called *local labels*. See [Section 5.3 \[Symbol Names\]](#), page 27. Normally you do not see such labels when debugging, because they are intended for the use of programs (like compilers) that compose assembler programs, not for your notice. Normally both **as** and **ld** discard such labels, so you do not normally debug with them.

This option tells **as** to retain those `‘L...’` symbols in the object file. Usually if you do this you also tell the linker **ld** to preserve symbols whose names begin with `‘L’`.

By default, a local label is any label beginning with `‘L’`, but each target is allowed to redefine the local label prefix. On the HPPA local labels begin with `‘L$’`. `‘;’` for the ARM family;

2.7 Assemble in MRI Compatibility Mode: -M

The `-M` or `--mri` option selects MRI compatibility mode. This changes the syntax and pseudo-op handling of **as** to make it compatible with the ASM68K or the ASM960 (depending upon the configured target) assembler from Microtec Research. The exact nature of the MRI syntax will not be documented here; see the MRI manuals for more information. Note in particular that the handling of macros and macro arguments is somewhat different. The purpose of this option is to permit assembling existing MRI assembler code using **as**.

The MRI compatibility is not complete. Certain operations of the MRI assembler depend upon its object file format, and can not be supported using other object file formats. Supporting these would require enhancing each object file format individually. These are:

- global symbols in common section

The m68k MRI assembler supports common sections which are merged by the linker. Other object file formats do not support this. **as** handles common sections by treating them as a single common symbol. It permits local symbols to be defined within a common section, but it can not support global symbols, since it has no way to describe them.

- complex relocations

The MRI assemblers support relocations against a negated section address, and relocations which combine the start addresses of two or more sections. These are not supported by other object file formats.
- END pseudo-op specifying start address

The MRI END pseudo-op permits the specification of a start address. This is not supported by other object file formats. The start address may instead be specified using the `-e` option to the linker, or in a linker script.
- IDNT, `.ident` and NAME pseudo-ops

The MRI IDNT, `.ident` and NAME pseudo-ops assign a module name to the output file. This is not supported by other object file formats.
- ORG pseudo-op

The m68k MRI ORG pseudo-op begins an absolute section at a given address. This differs from the usual `as .org` pseudo-op, which changes the location within the current section. Absolute sections are not supported by other object file formats. The address of a section may be assigned within a linker script.

There are some other features of the MRI assembler which are not supported by `as`, typically either because they are difficult or because they seem of little consequence. Some of these may be supported in future releases.

- EBCDIC strings

EBCDIC strings are not supported.
- packed binary coded decimal

Packed binary coded decimal is not supported. This means that the `DC.P` and `DCB.P` pseudo-ops are not supported.
- FEQU pseudo-op

The m68k FEQU pseudo-op is not supported.
- NOOBJ pseudo-op

The m68k NOOBJ pseudo-op is not supported.
- OPT branch control options

The m68k OPT branch control options—`B`, `BRS`, `BRB`, `BRL`, and `BRW`—are ignored. `as` automatically relaxes all branches, whether forward or backward, to an appropriate size, so these options serve no purpose.
- OPT list control options

The following m68k OPT list control options are ignored: `C`, `CEX`, `CL`, `CRE`, `E`, `G`, `I`, `M`, `MEX`, `MC`, `MD`, `X`.
- other OPT options

The following m68k OPT options are ignored: `NEST`, `O`, `OLD`, `OP`, `P`, `PCO`, `PCR`, `PCS`, `R`.
- OPT D option is default

The m68k OPT D option is the default, unlike the MRI assembler. OPT NOD may be used to turn it off.
- XREF pseudo-op.

The m68k XREF pseudo-op is ignored.

- `.debug` pseudo-op
The i960 `.debug` pseudo-op is not supported.
- `.extended` pseudo-op
The i960 `.extended` pseudo-op is not supported.
- `.list` pseudo-op.
The various options of the i960 `.list` pseudo-op are not supported.
- `.optimize` pseudo-op
The i960 `.optimize` pseudo-op is not supported.
- `.output` pseudo-op
The i960 `.output` pseudo-op is not supported.
- `.setreal` pseudo-op
The i960 `.setreal` pseudo-op is not supported.

2.8 Dependency tracking: `--MD`

`as` can generate a dependency file for the file it creates. This file consists of a single rule suitable for `make` describing the dependencies of the main source file.

The rule is written to the file named in its argument.

This feature is used in the automatic updating of makefiles.

2.9 Name the Object File: `-o`

There is always one object file output when you run `as`. By default it has the name `'a.out'` (or `'b.out'`, for Intel 960 targets only). You use this option (which takes exactly one filename) to give the object file a different name.

Whatever the object file is called, `as` overwrites any existing file of the same name.

2.10 Join Data and Text Sections: `-R`

`-R` tells `as` to write the object file as if all data-section data lives in the text section. This is only done at the very last moment: your binary data are the same, but data section parts are relocated differently. The data section part of your object file is zero bytes long because all its bytes are appended to the text section. (See [Chapter 4 \[Sections and Relocation\]](#), page 21.)

When you specify `-R` it would be possible to generate shorter address displacements (because we do not have to cross between text and data section). We refrain from doing this simply for compatibility with older versions of `as`. In future, `-R` may work this way.

When `as` is configured for COFF output, this option is only useful if you use sections named `' .text'` and `' .data'`.

`-R` is not supported for any of the HPPA targets. Using `-R` generates a warning from `as`.

2.11 Display Assembly Statistics: `--statistics`

Use `--statistics` to display two statistics about the resources used by `as`: the maximum amount of space allocated during the assembly (in bytes), and the total execution time taken for the assembly (in CPU seconds).

2.12 Compatible output: `--traditional-format`

For some targets, the output of `as` is different in some ways from the output of some existing assembler. This switch requests `as` to use the traditional format instead.

For example, it disables the exception frame optimizations which `as` normally does by default on `gcc` output.

2.13 Announce Version: `-v`

You can find out what version of `as` is running by including the option `-v` (which you can also spell as `-version`) on the command line.

2.14 Control Warnings: `-W`, `--warn`, `--no-warn`, `--fatal-warnings`

`as` should never give a warning or error message when assembling compiler output. But programs written by people often cause `as` to give a warning that a particular assumption was made. All such warnings are directed to the standard error file.

If you use the `-W` and `--no-warn` options, no warnings are issued. This only affects the warning messages: it does not change any particular of how `as` assembles your file. Errors, which stop the assembly, are still reported.

If you use the `--fatal-warnings` option, `as` considers files that generate warnings to be in error.

You can switch these options off again by specifying `--warn`, which causes warnings to be output as usual.

2.15 Generate Object File in Spite of Errors: `-Z`

After an error message, `as` normally produces no output. If for some reason you are interested in object file output even after `as` gives an error message on your program, use the `-Z` option. If there are any errors, `as` continues anyways, and writes an object file after a final warning message of the form `'n errors, m warnings, generating bad object file.'`

3 Syntax

This chapter describes the machine-independent syntax allowed in a source file. `as` syntax is similar to what many other assemblers use; it is inspired by the BSD 4.2 assembler, except that `as` does not assemble Vax bit-fields.

3.1 Preprocessing

The `as` internal preprocessor:

- adjusts and removes extra whitespace. It leaves one space or tab before the keywords on a line, and turns any other whitespace on the line into a single space.
- removes all comments, replacing them with a single space, or an appropriate number of newlines.
- converts character constants into the appropriate numeric values.

It does not do macro processing, include file handling, or anything else you may get from your C compiler's preprocessor. You can do include file processing with the `.include` directive (see [Section 7.34 \[.include\], page 40](#)). You can use the GNU C compiler driver to get other “CPP” style preprocessing, by giving the input file a `.S` suffix. See [section “Options Controlling the Kind of Output” in *Using GNU CC*](#).

Excess whitespace, comments, and character constants cannot be used in the portions of the input text that are not preprocessed.

If the first line of an input file is `#NO_APP` or if you use the `-f` option, whitespace and comments are not removed from the input file. Within an input file, you can ask for whitespace and comment removal in specific portions of the by putting a line that says `#APP` before the text that may contain whitespace or comments, and putting a line that says `#NO_APP` after this text. This feature is mainly intend to support `asm` statements in compilers whose output is otherwise free of comments and whitespace.

3.2 Whitespace

Whitespace is one or more blanks or tabs, in any order. Whitespace is used to separate symbols, and to make programs neater for people to read. Unless within character constants (see [Section 3.6.1 \[Character Constants\], page 17](#)), any whitespace means the same as exactly one space.

3.3 Comments

There are two ways of rendering comments to `as`. In both cases the comment is equivalent to one space.

Anything from `/*` through the next `*/` is a comment. This means you may not nest these comments.

```
/*
```

```
    The only way to include a newline ('\n') in a comment  
    is to use this sort of comment.
```

```
*/
/* This sort of comment does not nest. */
```

Anything from the *line comment* character to the next newline is considered a comment and is ignored. The line comment character is ‘;’ for the AMD 29K family; ‘;’ on the ARC; ‘@’ on the ARM; ‘;’ for the H8/300 family; ‘!’ for the H8/500 family; ‘;’ for the HPPA; ‘#’ on the i960; ‘;’ for picoJava; ‘!’ for the Hitachi SH; ‘!’ on the SPARC; ‘#’ on the m32r; ‘|’ on the 680x0; ‘#’ on the Vax; ‘!’ for the Z8000; ‘#’ on the V850; see [Chapter 8 \[Machine Dependencies\]](#), page 53.

On some machines there are two different line comment characters. One character only begins a comment if it is the first non-whitespace character on a line, while the other always begins a comment.

The V850 assembler also supports a double dash as starting a comment that extends to the end of the line.

```
‘--’;
```

To be compatible with past assemblers, lines that begin with ‘#’ have a special interpretation. Following the ‘#’ should be an absolute expression (see [Chapter 6 \[Expressions\]](#), page 31): the logical line number of the *next* line. Then a string (see [Section 3.6.1.1 \[Strings\]](#), page 17) is allowed: if present it is a new logical file name. The rest of the line, if any, should be whitespace.

If the first non-whitespace characters on the line are not numeric, the line is ignored. (Just like a comment.)

```

# 42-6 "new_file_name"    # This is an ordinary comment.
                        # New logical file name
                        # This is logical line # 36.
```

This feature is deprecated, and may disappear from future versions of `as`.

3.4 Symbols

A *symbol* is one or more characters chosen from the set of all letters (both upper and lower case), digits and the three characters ‘_.\$’. On most machines, you can also use \$ in symbol names; exceptions are noted in [Chapter 8 \[Machine Dependencies\]](#), page 53. No symbol may begin with a digit. Case is significant. There is no length limit: all characters are significant. Symbols are delimited by characters not in that set, or by the beginning of a file (since the source program must end with a newline, the end of a file is not a possible symbol delimiter). See [Chapter 5 \[Symbols\]](#), page 27.

3.5 Statements

A *statement* ends at a newline character (‘\n’) or line separator character. (The line separator is usually ‘;’, unless this conflicts with the comment character; see [Chapter 8 \[Machine Dependencies\]](#), page 53.) The newline or separator character is considered part of the preceding statement. Newlines and separators within character constants are an exception: they do not end statements.

It is an error to end any statement with end-of-file: the last character of any input file should be a newline.

An empty statement is allowed, and may include whitespace. It is ignored.

A statement begins with zero or more labels, optionally followed by a key symbol which determines what kind of statement it is. The key symbol determines the syntax of the rest of the statement. If the symbol begins with a dot ‘.’ then the statement is an assembler directive: typically valid for any computer. If the symbol begins with a letter the statement is an assembly language *instruction*: it assembles into a machine language instruction. Different versions of `as` for different computers recognize different instructions. In fact, the same symbol may represent a different instruction in a different computer’s assembly language.

A label is a symbol immediately followed by a colon (:). Whitespace before a label or after a colon is permitted, but you may not have whitespace between a label’s symbol and its colon. See [Section 5.1 \[Labels\], page 27](#).

For HPPA targets, labels need not be immediately followed by a colon, but the definition of a label must begin in column zero. This also implies that only one label may be defined on each line.

```
label:      .directive      followed by something
another_label:      # This is an empty statement.
                instruction  operand_1, operand_2, ...
```

3.6 Constants

A constant is a number, written so that its value is known by inspection, without knowing any context. Like this:

```
.byte  74, 0112, 092, 0x4A, 0X4a, 'J, '\J # All the same value.
.ascii "Ring the bell\7"                # A string constant.
.octa  0x123456789abcdef0123456789ABCDEF0 # A bignum.
.float 0f-314159265358979323846264338327\
95028841971.693993751E-40                # - pi, a flonum.
```

3.6.1 Character Constants

There are two kinds of character constants. A *character* stands for one character in one byte and its value may be used in numeric expressions. String constants (properly called string *literals*) are potentially many bytes and their values may not be used in arithmetic expressions.

3.6.1.1 Strings

A *string* is written between double-quotes. It may contain double-quotes or null characters. The way to get special characters into a string is to *escape* these characters: precede them with a backslash ‘\’ character. For example ‘\\’ represents one backslash: the first \ is an escape which tells `as` to interpret the second character literally as a backslash (which prevents `as` from recognizing the second \ as an escape character). The complete list of escapes follows.

<code>\b</code>	Mnemonic for backspace; for ASCII this is octal code 010.
<code>\f</code>	Mnemonic for FormFeed; for ASCII this is octal code 014.
<code>\n</code>	Mnemonic for newline; for ASCII this is octal code 012.
<code>\r</code>	Mnemonic for carriage-Return; for ASCII this is octal code 015.
<code>\t</code>	Mnemonic for horizontal Tab; for ASCII this is octal code 011.
<code>\ digit digit digit</code>	An octal character code. The numeric code is 3 octal digits. For compatibility with other Unix systems, 8 and 9 are accepted as digits: for example, <code>\008</code> has the value 010, and <code>\009</code> the value 011.
<code>\x hex-digits...</code>	A hex character code. All trailing hex digits are combined. Either upper or lower case <code>x</code> works.
<code>\\</code>	Represents one <code>'\'</code> character.
<code>\"</code>	Represents one <code>'\"'</code> character. Needed in strings to represent this character, because an unescaped <code>'\"'</code> would end the string.
<code>\ anything-else</code>	Any other character when escaped by <code>\</code> gives a warning, but assembles as if the <code>'\'</code> was not present. The idea is that if you used an escape sequence you clearly didn't want the literal interpretation of the following character. However <code>as</code> has no other interpretation, so <code>as</code> knows it is giving you the wrong code and warns you of the fact.

Which characters are escapable, and what those escapes represent, varies widely among assemblers. The current set is what we think the BSD 4.2 assembler recognizes, and is a subset of what most C compilers recognize. If you are in doubt, do not use an escape sequence.

3.6.1.2 Characters

A single character may be written as a single quote immediately followed by that character. The same escapes apply to characters as to strings. So if you want to write the character backslash, you must write `'\\'` where the first `\` escapes the second `\`. As you can see, the quote is an acute accent, not a grave accent. A newline immediately following an acute accent is taken as a literal character and does not count as the end of a statement. The value of a character constant in a numeric expression is the machine's byte-wide code for that character. `as` assumes your character code is ASCII: `'A'` means 65, `'B'` means 66, and so on.

3.6.2 Number Constants

`as` distinguishes three kinds of numbers according to how they are stored in the target machine. *Integers* are numbers that would fit into an `int` in the C language. *Bignums* are integers, but they are stored in more than 32 bits. *Flonums* are floating point numbers, described below.

3.6.2.1 Integers

A binary integer is ‘0b’ or ‘0B’ followed by zero or more of the binary digits ‘01’.

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

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

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

Integers have the usual values. To denote a negative integer, use the prefix operator ‘-’ discussed under expressions (see [Section 6.2.3 \[Prefix Operators\]](#), page 32).

3.6.2.2 Bignums

A *bignum* has the same syntax and semantics as an integer except that the number (or its negative) takes more than 32 bits to represent in binary. The distinction is made because in some places integers are permitted while bignums are not.

3.6.2.3 Flonums

A *flonum* represents a floating point number. The translation is indirect: a decimal floating point number from the text is converted by `as` to a generic binary floating point number of more than sufficient precision. This generic floating point number is converted to a particular computer’s floating point format (or formats) by a portion of `as` specialized to that computer.

A flonum is written by writing (in order)

- The digit ‘0’. (‘0’ is optional on the HPPA.)
- A letter, to tell `as` the rest of the number is a flonum. `e` is recommended. Case is not important.

On the H8/300, H8/500, Hitachi SH, and AMD 29K architectures, the letter must be one of the letters ‘DFPRSX’ (in upper or lower case).

On the ARC, the letter must be one of the letters ‘DFRS’ (in upper or lower case).

On the Intel 960 architecture, the letter must be one of the letters ‘DFT’ (in upper or lower case).

On the HPPA architecture, the letter must be ‘E’ (upper case only).

- An optional sign: either ‘+’ or ‘-’.
- An optional *integer part*: zero or more decimal digits.
- An optional *fractional part*: ‘.’ followed by zero or more decimal digits.
- An optional exponent, consisting of:
 - An ‘E’ or ‘e’.
 - Optional sign: either ‘+’ or ‘-’.
 - One or more decimal digits.

At least one of the integer part or the fractional part must be present. The floating point number has the usual base-10 value.

`as` does all processing using integers. Flonums are computed independently of any floating point hardware in the computer running `as`.

4 Sections and Relocation

4.1 Background

Roughly, a section is a range of addresses, with no gaps; all data “in” those addresses is treated the same for some particular purpose. For example there may be a “read only” section.

The linker `ld` reads many object files (partial programs) and combines their contents to form a runnable program. When `as` emits an object file, the partial program is assumed to start at address 0. `ld` assigns the final addresses for the partial program, so that different partial programs do not overlap. This is actually an oversimplification, but it suffices to explain how `as` uses sections.

`ld` moves blocks of bytes of your program to their run-time addresses. These blocks slide to their run-time addresses as rigid units; their length does not change and neither does the order of bytes within them. Such a rigid unit is called a *section*. Assigning run-time addresses to sections is called *relocation*. It includes the task of adjusting mentions of object-file addresses so they refer to the proper run-time addresses. For the H8/300 and H8/500, and for the Hitachi SH, `as` pads sections if needed to ensure they end on a word (sixteen bit) boundary.

An object file written by `as` has at least three sections, any of which may be empty. These are named *text*, *data* and *bss* sections.

When it generates COFF output, `as` can also generate whatever other named sections you specify using the `‘.section’` directive (see [Section 7.58 \[‘.section’\]](#), page 46). If you do not use any directives that place output in the `‘.text’` or `‘.data’` sections, these sections still exist, but are empty.

When `as` generates SOM or ELF output for the HPPA, `as` can also generate whatever other named sections you specify using the `‘.space’` and `‘.subspace’` directives. See *HP9000 Series 800 Assembly Language Reference Manual* (HP 92432-90001) for details on the `‘.space’` and `‘.subspace’` assembler directives.

Additionally, `as` uses different names for the standard text, data, and bss sections when generating SOM output. Program text is placed into the `‘$CODE$’` section, data into `‘$DATA$’`, and BSS into `‘BSS’`.

Within the object file, the text section starts at address 0, the data section follows, and the bss section follows the data section.

When generating either SOM or ELF output files on the HPPA, the text section starts at address 0, the data section at address 0x4000000, and the bss section follows the data section.

To let `ld` know which data changes when the sections are relocated, and how to change that data, `as` also writes to the object file details of the relocation needed. To perform relocation `ld` must know, each time an address in the object file is mentioned:

- Where in the object file is the beginning of this reference to an address?
- How long (in bytes) is this reference?
- Which section does the address refer to? What is the numeric value of

$(address) - (start\text{-}address\ of\ section)?$

- Is the reference to an address “Program-Counter relative”?

In fact, every address `as` ever uses is expressed as

$(section) + (offset\ into\ section)$

Further, most expressions `as` computes have this section-relative nature. (For some object formats, such as SOM for the HPPA, some expressions are symbol-relative instead.)

In this manual we use the notation `{secname N}` to mean “offset *N* into section *secname*.”

Apart from text, data and bss sections you need to know about the *absolute* section. When `ld` mixes partial programs, addresses in the absolute section remain unchanged. For example, address `{absolute 0}` is “relocated” to run-time address 0 by `ld`. Although the linker never arranges two partial programs’ data sections with overlapping addresses after linking, *by definition* their absolute sections must overlap. Address `{absolute 239}` in one part of a program is always the same address when the program is running as address `{absolute 239}` in any other part of the program.

The idea of sections is extended to the *undefined* section. Any address whose section is unknown at assembly time is by definition rendered `{undefined U}`—where *U* is filled in later. Since numbers are always defined, the only way to generate an undefined address is to mention an undefined symbol. A reference to a named common block would be such a symbol: its value is unknown at assembly time so it has section *undefined*.

By analogy the word *section* is used to describe groups of sections in the linked program. `ld` puts all partial programs’ text sections in contiguous addresses in the linked program. It is customary to refer to the *text section* of a program, meaning all the addresses of all partial programs’ text sections. Likewise for data and bss sections.

Some sections are manipulated by `ld`; others are invented for use of `as` and have no meaning except during assembly.

4.2 Linker Sections

`ld` deals with just four kinds of sections, summarized below.

named sections

text section

data section

These sections hold your program. `as` and `ld` treat them as separate but equal sections. Anything you can say of one section is true another. When the program is running, however, it is customary for the text section to be unalterable. The text section is often shared among processes: it contains instructions, constants and the like. The data section of a running program is usually alterable: for example, C variables would be stored in the data section.

bss section

This section contains zeroed bytes when your program begins running. It is used to hold uninitialized variables or common storage. The length of each partial program’s bss section is important, but because it starts out containing zeroed bytes there is no need to store explicit zero bytes in the object file. The bss section was invented to eliminate those explicit zeros from object files.

absolute section

Address 0 of this section is always “relocated” to runtime address 0. This is useful if you want to refer to an address that `ld` must not change when relocating. In this sense we speak of absolute addresses being “unrelocatable”: they do not change during relocation.

undefined section

This “section” is a catch-all for address references to objects not in the preceding sections.

An idealized example of three relocatable sections follows. The example uses the traditional section names `.text` and `.data`. Memory addresses are on the horizontal axis.

Partial program #1:

text	data	bss
ttttt	dddd	00

Partial program #2:

text	data	bss
TTT	DDDD	000

linked program:

text		data		bss		
	TTT	ttttt		dddd	DDDD	00000 ...

addresses:

0...

4.3 Assembler Internal Sections

These sections are meant only for the internal use of `as`. They have no meaning at run-time. You do not really need to know about these sections for most purposes; but they can be mentioned in `as` warning messages, so it might be helpful to have an idea of their meanings to `as`. These sections are used to permit the value of every expression in your assembly language program to be a section-relative address.

ASSEMBLER-INTERNAL-LOGIC-ERROR!

An internal assembler logic error has been found. This means there is a bug in the assembler.

expr section

The assembler stores complex expression internally as combinations of symbols. When it needs to represent an expression as a symbol, it puts it in the `expr` section.

4.4 Sub-Sections

Assembled bytes conventionally fall into two sections: `text` and `data`. You may have separate groups of data in named sections that you want to end up near to each other in the object file, even though they are not contiguous in the assembler source. `as` allows you to use *subsections* for this purpose. Within each section, there can be numbered subsections

with values from 0 to 8192. Objects assembled into the same subsection go into the object file together with other objects in the same subsection. For example, a compiler might want to store constants in the text section, but might not want to have them interspersed with the program being assembled. In this case, the compiler could issue a `.text 0` before each section of code being output, and a `.text 1` before each group of constants being output.

Subsections are optional. If you do not use subsections, everything goes in subsection number zero.

Each subsection is zero-padded up to a multiple of four bytes. (Subsections may be padded a different amount on different flavors of `as`.)

Subsections appear in your object file in numeric order, lowest numbered to highest. (All this to be compatible with other people's assemblers.) The object file contains no representation of subsections; `ld` and other programs that manipulate object files see no trace of them. They just see all your text subsections as a text section, and all your data subsections as a data section.

To specify which subsection you want subsequent statements assembled into, use a numeric argument to specify it, in a `.text expression` or a `.data expression` statement. When generating COFF output, you can also use an extra subsection argument with arbitrary named sections: `.section name, expression`. *Expression* should be an absolute expression. (See [Chapter 6 \[Expressions\]](#), page 31.) If you just say `.text` then `.text 0` is assumed. Likewise `.data` means `.data 0`. Assembly begins in text 0. For instance:

```
.text 0      # The default subsection is text 0 anyway.
.ascii "This lives in the first text subsection. *"
.text 1
.ascii "But this lives in the second text subsection."
.data 0
.ascii "This lives in the data section,"
.ascii "in the first data subsection."
.text 0
.ascii "This lives in the first text section,"
.ascii "immediately following the asterisk (*)."
```

Each section has a *location counter* incremented by one for every byte assembled into that section. Because subsections are merely a convenience restricted to `as` there is no concept of a subsection location counter. There is no way to directly manipulate a location counter—but the `.align` directive changes it, and any label definition captures its current value. The location counter of the section where statements are being assembled is said to be the *active* location counter.

4.5 bss Section

The bss section is used for local common variable storage. You may allocate address space in the bss section, but you may not dictate data to load into it before your program executes. When your program starts running, all the contents of the bss section are zeroed bytes.

The `.lcomm` pseudo-op defines a symbol in the bss section; see [Section 7.38 \[.lcomm\]](#), page 41.

The `.comm` pseudo-op may be used to declare a common symbol, which is another form of uninitialized symbol; see See [Section 7.8 \[.comm\]](#), page 35.

When assembling for a target which supports multiple sections, such as ELF or COFF, you may switch into the `.bss` section and define symbols as usual; see [Section 7.58 \[.section\]](#), page 46. You may only assemble zero values into the section. Typically the section will only contain symbol definitions and `.skip` directives (see [Section 7.64 \[.skip\]](#), page 48).

5 Symbols

Symbols are a central concept: the programmer uses symbols to name things, the linker uses symbols to link, and the debugger uses symbols to debug.

Warning: `as` does not place symbols in the object file in the same order they were declared. This may break some debuggers.

5.1 Labels

A *label* is written as a symbol immediately followed by a colon ‘:’. The symbol then represents the current value of the active location counter, and is, for example, a suitable instruction operand. You are warned if you use the same symbol to represent two different locations: the first definition overrides any other definitions.

On the HPPA, the usual form for a label need not be immediately followed by a colon, but instead must start in column zero. Only one label may be defined on a single line. To work around this, the HPPA version of `as` also provides a special directive `.label` for defining labels more flexibly.

5.2 Giving Symbols Other Values

A symbol can be given an arbitrary value by writing a symbol, followed by an equals sign ‘=’, followed by an expression (see [Chapter 6 \[Expressions\], page 31](#)). This is equivalent to using the `.set` directive. See [Section 7.59 \[.set\], page 47](#).

5.3 Symbol Names

Symbol names begin with a letter or with one of ‘. _’. On most machines, you can also use \$ in symbol names; exceptions are noted in [Chapter 8 \[Machine Dependencies\], page 53](#). That character may be followed by any string of digits, letters, dollar signs (unless otherwise noted in [Chapter 8 \[Machine Dependencies\], page 53](#)), and underscores. For the AMD 29K family, ‘?’ is also allowed in the body of a symbol name, though not at its beginning.

Case of letters is significant: `foo` is a different symbol name than `Foo`.

Each symbol has exactly one name. Each name in an assembly language program refers to exactly one symbol. You may use that symbol name any number of times in a program.

Local Symbol Names

Local symbols help compilers and programmers use names temporarily. There are ten local symbol names, which are re-used throughout the program. You may refer to them using the names ‘0’ ‘1’ . . . ‘9’. To define a local symbol, write a label of the form ‘**N**:’ (where **N** represents any digit). To refer to the most recent previous definition of that symbol write ‘**Nb**’, using the same digit as when you defined the label. To refer to the next definition of a local label, write ‘**Nf**’—where **N** gives you a choice of 10 forward references. The ‘b’ stands for “backwards” and the ‘f’ stands for “forwards”.

Local symbols are not emitted by the current GNU C compiler.

There is no restriction on how you can use these labels, but remember that at any point in the assembly you can refer to at most 10 prior local labels and to at most 10 forward local labels.

Local symbol names are only a notation device. They are immediately transformed into more conventional symbol names before the assembler uses them. The symbol names stored in the symbol table, appearing in error messages and optionally emitted to the object file have these parts:

- L** All local labels begin with ‘L’. Normally both `as` and `ld` forget symbols that start with ‘L’. These labels are used for symbols you are never intended to see. If you use the ‘-L’ option then `as` retains these symbols in the object file. If you also instruct `ld` to retain these symbols, you may use them in debugging.
- digit* If the label is written ‘0:’ then the digit is ‘0’. If the label is written ‘1:’ then the digit is ‘1’. And so on up through ‘9:’.
- C-A** This unusual character is included so you do not accidentally invent a symbol of the same name. The character has ASCII value ‘\001’.
- ordinal number*
This is a serial number to keep the labels distinct. The first ‘0:’ gets the number ‘1’; The 15th ‘0:’ gets the number ‘15’; *etc.*. Likewise for the other labels ‘1:’ through ‘9:’.

For instance, the first 1: is named `L1C-A1`, the 44th 3: is named `L3C-A44`.

5.4 The Special Dot Symbol

The special symbol ‘.’ refers to the current address that `as` is assembling into. Thus, the expression ‘`melvin: .long .`’ defines `melvin` to contain its own address. Assigning a value to `.` is treated the same as a `.org` directive. Thus, the expression ‘`.=. +4`’ is the same as saying ‘`.space 4`’.

5.5 Symbol Attributes

Every symbol has, as well as its name, the attributes “Value” and “Type”. Depending on output format, symbols can also have auxiliary attributes.

If you use a symbol without defining it, `as` assumes zero for all these attributes, and probably won’t warn you. This makes the symbol an externally defined symbol, which is generally what you would want.

5.5.1 Value

The value of a symbol is (usually) 32 bits. For a symbol which labels a location in the text, data, bss or absolute sections the value is the number of addresses from the start of that section to the label. Naturally for text, data and bss sections the value of a symbol changes as `ld` changes section base addresses during linking. Absolute symbols’ values do not change during linking: that is why they are called absolute.

The value of an undefined symbol is treated in a special way. If it is 0 then the symbol is not defined in this assembler source file, and `ld` tries to determine its value from other files linked into the same program. You make this kind of symbol simply by mentioning a symbol name without defining it. A non-zero value represents a `.comm` common declaration. The value is how much common storage to reserve, in bytes (addresses). The symbol refers to the first address of the allocated storage.

5.5.2 Type

The type attribute of a symbol contains relocation (section) information, any flag settings indicating that a symbol is external, and (optionally), other information for linkers and debuggers. The exact format depends on the object-code output format in use.

5.5.3 Symbol Attributes: `a.out`

5.5.3.1 Descriptor

This is an arbitrary 16-bit value. You may establish a symbol's descriptor value by using a `.desc` statement (see [Section 7.11 \[`.desc`\], page 35](#)). A descriptor value means nothing to `as`.

5.5.3.2 Other

This is an arbitrary 8-bit value. It means nothing to `as`.

5.5.4 Symbol Attributes for COFF

The COFF format supports a multitude of auxiliary symbol attributes; like the primary symbol attributes, they are set between `.def` and `.endef` directives.

5.5.4.1 Primary Attributes

The symbol name is set with `.def`; the value and type, respectively, with `.val` and `.type`.

5.5.4.2 Auxiliary Attributes

The `as` directives `.dim`, `.line`, `.scl`, `.size`, and `.tag` can generate auxiliary symbol table information for COFF.

5.5.5 Symbol Attributes for SOM

The SOM format for the HPPA supports a multitude of symbol attributes set with the `.EXPORT` and `.IMPORT` directives.

The attributes are described in *HP9000 Series 800 Assembly Language Reference Manual* (HP 92432-90001) under the `IMPORT` and `EXPORT` assembler directive documentation.

6 Expressions

An *expression* specifies an address or numeric value. Whitespace may precede and/or follow an expression.

The result of an expression must be an absolute number, or else an offset into a particular section. If an expression is not absolute, and there is not enough information when `as` sees the expression to know its section, a second pass over the source program might be necessary to interpret the expression—but the second pass is currently not implemented. `as` aborts with an error message in this situation.

6.1 Empty Expressions

An empty expression has no value: it is just whitespace or null. Wherever an absolute expression is required, you may omit the expression, and `as` assumes a value of (absolute) 0. This is compatible with other assemblers.

6.2 Integer Expressions

An *integer expression* is one or more *arguments* delimited by *operators*.

6.2.1 Arguments

Arguments are symbols, numbers or subexpressions. In other contexts arguments are sometimes called “arithmetic operands”. In this manual, to avoid confusing them with the “instruction operands” of the machine language, we use the term “argument” to refer to parts of expressions only, reserving the word “operand” to refer only to machine instruction operands.

Symbols are evaluated to yield `{section NNN}` where *section* is one of `text`, `data`, `bss`, `absolute`, or `undefined`. *NNN* is a signed, 2’s complement 32 bit integer.

Numbers are usually integers.

A number can be a flonum or bignum. In this case, you are warned that only the low order 32 bits are used, and `as` pretends these 32 bits are an integer. You may write integer-manipulating instructions that act on exotic constants, compatible with other assemblers.

Subexpressions are a left parenthesis ‘(’ followed by an integer expression, followed by a right parenthesis ‘)’; or a prefix operator followed by an argument.

6.2.2 Operators

Operators are arithmetic functions, like `+` or `%`. Prefix operators are followed by an argument. Infix operators appear between their arguments. Operators may be preceded and/or followed by whitespace.

6.2.3 Prefix Operator

as has the following *prefix operators*. They each take one argument, which must be absolute.

- *Negation*. Two's complement negation.
- ~ *Complementation*. Bitwise not.

6.2.4 Infix Operators

Infix operators take two arguments, one on either side. Operators have precedence, but operations with equal precedence are performed left to right. Apart from + or -, both arguments must be absolute, and the result is absolute.

1. Highest Precedence

- * *Multiplication*.
- / *Division*. Truncation is the same as the C operator '/'
- % *Remainder*.
- <
- << *Shift Left*. Same as the C operator '<<'
- >
- >> *Shift Right*. Same as the C operator '>>'

2. Intermediate precedence

- | *Bitwise Inclusive Or*.
- & *Bitwise And*.
- ^ *Bitwise Exclusive Or*.
- ! *Bitwise Or Not*.

3. Lowest Precedence

- + *Addition*. If either argument is absolute, the result has the section of the other argument. You may not add together arguments from different sections.
- *Subtraction*. If the right argument is absolute, the result has the section of the left argument. If both arguments are in the same section, the result is absolute. You may not subtract arguments from different sections.

In short, it's only meaningful to add or subtract the *offsets* in an address; you can only have a defined section in one of the two arguments.

7 Assembler Directives

All assembler directives have names that begin with a period (‘.’). The rest of the name is letters, usually in lower case.

This chapter discusses directives that are available regardless of the target machine configuration for the GNU assembler. Some machine configurations provide additional directives. See [Chapter 8 \[Machine Dependencies\]](#), page 53.

7.1 .abort

This directive stops the assembly immediately. It is for compatibility with other assemblers. The original idea was that the assembly language source would be piped into the assembler. If the sender of the source quit, it could use this directive tells `as` to quit also. One day `.abort` will not be supported.

7.2 .ABORT

When producing COFF output, `as` accepts this directive as a synonym for ‘`.abort`’.

When producing `b.out` output, `as` accepts this directive, but ignores it.

7.3 .align *abs-expr*, *abs-expr*, *abs-expr*

Pad the location counter (in the current subsection) to a particular storage boundary. The first expression (which must be absolute) is the alignment required, as described below.

The second expression (also absolute) gives the fill value to be stored in the padding bytes. It (and the comma) may be omitted. If it is omitted, the padding bytes are normally zero. However, on some systems, if the section is marked as containing code and the fill value is omitted, the space is filled with no-op instructions.

The third expression is also absolute, and is also optional. If it is present, it is the maximum number of bytes that should be skipped by this alignment directive. If doing the alignment would require skipping more bytes than the specified maximum, then the alignment is not done at all. You can omit the fill value (the second argument) entirely by simply using two commas after the required alignment; this can be useful if you want the alignment to be filled with no-op instructions when appropriate.

The way the required alignment is specified varies from system to system. For the `a29k`, `hppa`, `m68k`, `m88k`, `w65`, `sparc`, and Hitachi SH, and `i386` using ELF format, the first expression is the alignment request in bytes. For example ‘`.align 8`’ advances the location counter until it is a multiple of 8. If the location counter is already a multiple of 8, no change is needed.

For other systems, including the `i386` using `a.out` format, it is the number of low-order zero bits the location counter must have after advancement. For example ‘`.align 3`’ advances the location counter until it a multiple of 8. If the location counter is already a multiple of 8, no change is needed.

This inconsistency is due to the different behaviors of the various native assemblers for these systems which GAS must emulate. GAS also provides `.balign` and `.p2align`

directives, described later, which have a consistent behavior across all architectures (but are specific to GAS).

7.4 `.ascii "string"...`

`.ascii` expects zero or more string literals (see [Section 3.6.1.1 \[Strings\], page 17](#)) separated by commas. It assembles each string (with no automatic trailing zero byte) into consecutive addresses.

7.5 `.asciz "string"...`

`.asciz` is just like `.ascii`, but each string is followed by a zero byte. The “z” in ‘`.asciz`’ stands for “zero”.

7.6 `.balign[wl] abs-expr, abs-expr, abs-expr`

Pad the location counter (in the current subsection) to a particular storage boundary. The first expression (which must be absolute) is the alignment request in bytes. For example ‘`.balign 8`’ advances the location counter until it is a multiple of 8. If the location counter is already a multiple of 8, no change is needed.

The second expression (also absolute) gives the fill value to be stored in the padding bytes. It (and the comma) may be omitted. If it is omitted, the padding bytes are normally zero. However, on some systems, if the section is marked as containing code and the fill value is omitted, the space is filled with no-op instructions.

The third expression is also absolute, and is also optional. If it is present, it is the maximum number of bytes that should be skipped by this alignment directive. If doing the alignment would require skipping more bytes than the specified maximum, then the alignment is not done at all. You can omit the fill value (the second argument) entirely by simply using two commas after the required alignment; this can be useful if you want the alignment to be filled with no-op instructions when appropriate.

The `.balignw` and `.balignl` directives are variants of the `.balign` directive. The `.balignw` directive treats the fill pattern as a two byte word value. The `.balignl` directive treats the fill pattern as a four byte longword value. For example, `.balignw 4,0x368d` will align to a multiple of 4. If it skips two bytes, they will be filled in with the value 0x368d (the exact placement of the bytes depends upon the endianness of the processor). If it skips 1 or 3 bytes, the fill value is undefined.

7.7 `.byte expressions`

`.byte` expects zero or more expressions, separated by commas. Each expression is assembled into the next byte.

7.8 `.comm symbol , length`

`.comm` declares a common symbol named *symbol*. When linking, a common symbol in one object file may be merged with a defined or common symbol of the same name in another object file. If `ld` does not see a definition for the symbol—just one or more common symbols—then it will allocate *length* bytes of uninitialized memory. *length* must be an absolute expression. If `ld` sees multiple common symbols with the same name, and they do not all have the same size, it will allocate space using the largest size.

When using ELF, the `.comm` directive takes an optional third argument. This is the desired alignment of the symbol, specified as a byte boundary (for example, an alignment of 16 means that the least significant 4 bits of the address should be zero). The alignment must be an absolute expression, and it must be a power of two. If `ld` allocates uninitialized memory for the common symbol, it will use the alignment when placing the symbol. If no alignment is specified, `as` will set the alignment to the largest power of two less than or equal to the size of the symbol, up to a maximum of 16.

The syntax for `.comm` differs slightly on the HPPA. The syntax is '*symbol .comm , length*'; *symbol* is optional.

7.9 `.data subsection`

`.data` tells `as` to assemble the following statements onto the end of the data subsection numbered *subsection* (which is an absolute expression). If *subsection* is omitted, it defaults to zero.

7.10 `.def name`

Begin defining debugging information for a symbol *name*; the definition extends until the `.endef` directive is encountered.

This directive is only observed when `as` is configured for COFF format output; when producing `b.out`, '`.def`' is recognized, but ignored.

7.11 `.desc symbol, abs-expression`

This directive sets the descriptor of the symbol (see [Section 5.5 \[Symbol Attributes\], page 28](#)) to the low 16 bits of an absolute expression.

The '`.desc`' directive is not available when `as` is configured for COFF output; it is only for `a.out` or `b.out` object format. For the sake of compatibility, `as` accepts it, but produces no output, when configured for COFF.

7.12 `.dim`

This directive is generated by compilers to include auxiliary debugging information in the symbol table. It is only permitted inside `.def/.endef` pairs.

'`.dim`' is only meaningful when generating COFF format output; when `as` is generating `b.out`, it accepts this directive but ignores it.

7.13 `.double` *flonums*

`.double` expects zero or more flonums, separated by commas. It assembles floating point numbers. The exact kind of floating point numbers emitted depends on how `as` is configured. See [Chapter 8 \[Machine Dependencies\]](#), page 53.

7.14 `.eject`

Force a page break at this point, when generating assembly listings.

7.15 `.else`

`.else` is part of the `as` support for conditional assembly; see [Section 7.33 \[.if\]](#), page 39. It marks the beginning of a section of code to be assembled if the condition for the preceding `.if` was false.

7.16 `.end`

`.end` marks the end of the assembly file. `as` does not process anything in the file past the `.end` directive.

7.17 `.endef`

This directive flags the end of a symbol definition begun with `.def`.

‘`.endef`’ is only meaningful when generating COFF format output; if `as` is configured to generate `b.out`, it accepts this directive but ignores it.

7.18 `.endfunc`

`.endfunc` marks the end of a function specified with `.func`.

7.19 `.endif`

`.endif` is part of the `as` support for conditional assembly; it marks the end of a block of code that is only assembled conditionally. See [Section 7.33 \[.if\]](#), page 39.

7.20 `.equ` *symbol, expression*

This directive sets the value of *symbol* to *expression*. It is synonymous with ‘`.set`’; see [Section 7.59 \[.set\]](#), page 47.

The syntax for `equ` on the HPPA is ‘*symbol* `.equ` *expression*’.

7.21 `.equiv` *symbol*, *expression*

The `.equiv` directive is like `.equ` and `.set`, except that the assembler will signal an error if *symbol* is already defined.

Except for the contents of the error message, this is roughly equivalent to

```
.ifndef SYM
.err
.endif
.equ SYM,VAL
```

7.22 `.err`

If `as` assembles a `.err` directive, it will print an error message and, unless the `-Z` option was used, it will not generate an object file. This can be used to signal error an conditionally compiled code.

7.23 `.exitm`

Exit early from the current macro definition. See [Section 7.46 \[Macro\]](#), page 43.

7.24 `.extern`

`.extern` is accepted in the source program—for compatibility with other assemblers—but it is ignored. `as` treats all undefined symbols as external.

7.25 `.fail` *expression*

Generates an error or a warning. If the value of the *expression* is 500 or more, `as` will print a warning message. If the value is less than 500, `as` will print an error message. The message will include the value of *expression*. This can occasionally be useful inside complex nested macros or conditional assembly.

7.26 `.file` *string*

`.file` tells `as` that we are about to start a new logical file. *string* is the new file name. In general, the filename is recognized whether or not it is surrounded by quotes `"`; but if you wish to specify an empty file name, you must give the quotes—`"`. This statement may go away in future: it is only recognized to be compatible with old `as` programs. In some configurations of `as`, `.file` has already been removed to avoid conflicts with other assemblers. See [Chapter 8 \[Machine Dependencies\]](#), page 53.

7.27 `.fill repeat , size , value`

result, *size* and *value* are absolute expressions. This emits *repeat* copies of *size* bytes. *Repeat* may be zero or more. *Size* may be zero or more, but if it is more than 8, then it is deemed to have the value 8, compatible with other people's assemblers. The contents of each *repeat* bytes is taken from an 8-byte number. The highest order 4 bytes are zero. The lowest order 4 bytes are *value* rendered in the byte-order of an integer on the computer as is assembling for. Each *size* bytes in a repetition is taken from the lowest order *size* bytes of this number. Again, this bizarre behavior is compatible with other people's assemblers.

size and *value* are optional. If the second comma and *value* are absent, *value* is assumed zero. If the first comma and following tokens are absent, *size* is assumed to be 1.

7.28 `.float flonums`

This directive assembles zero or more flonums, separated by commas. It has the same effect as `.single`. The exact kind of floating point numbers emitted depends on how `as` is configured. See [Chapter 8 \[Machine Dependencies\]](#), page 53.

7.29 `.func name[,label]`

`.func` emits debugging information to denote function *name*, and is ignored unless the file is assembled with debugging enabled. Only `--gstabs` is currently supported. *label* is the entry point of the function and if omitted *name* prepended with the 'leading char' is used. 'leading char' is usually `_` or nothing, depending on the target. All functions are currently defined to have void return type. The function must be terminated with `.endfunc`.

7.30 `.global symbol, .globl symbol`

`.global` makes the symbol visible to `ld`. If you define *symbol* in your partial program, its value is made available to other partial programs that are linked with it. Otherwise, *symbol* takes its attributes from a symbol of the same name from another file linked into the same program.

Both spellings (`' .globl'` and `' .global'`) are accepted, for compatibility with other assemblers.

On the HPPA, `.global` is not always enough to make it accessible to other partial programs. You may need the HPPA-only `.EXPORT` directive as well. See [Section 8.8.5 \[HPPA Assembler Directives\]](#), page 75.

7.31 `.hword expressions`

This expects zero or more *expressions*, and emits a 16 bit number for each.

This directive is a synonym for `' .short'`; depending on the target architecture, it may also be a synonym for `' .word'`.

7.32 `.ident`

This directive is used by some assemblers to place tags in object files. `as` simply accepts the directive for source-file compatibility with such assemblers, but does not actually emit anything for it.

7.33 `.if absolute expression`

`.if` marks the beginning of a section of code which is only considered part of the source program being assembled if the argument (which must be an *absolute expression*) is non-zero. The end of the conditional section of code must be marked by `.endif` (see [Section 7.19 \[endif\], page 36](#)); optionally, you may include code for the alternative condition, flagged by `.else` (see [Section 7.15 \[else\], page 36](#)).

The following variants of `.if` are also supported:

`.ifdef symbol`

Assembles the following section of code if the specified *symbol* has been defined.

`.ifc string1,string2`

Assembles the following section of code if the two strings are the same. The strings may be optionally quoted with single quotes. If they are not quoted, the first string stops at the first comma, and the second string stops at the end of the line. Strings which contain whitespace should be quoted. The string comparison is case sensitive.

`.ifeq absolute expression`

Assembles the following section of code if the argument is zero.

`.ifeqs string1,string2`

Another form of `.ifc`. The strings must be quoted using double quotes.

`.ifge absolute expression`

Assembles the following section of code if the argument is greater than or equal to zero.

`.ifgt absolute expression`

Assembles the following section of code if the argument is greater than zero.

`.ifle absolute expression`

Assembles the following section of code if the argument is less than or equal to zero.

`.iflt absolute expression`

Assembles the following section of code if the argument is less than zero.

`.ifnc string1,string2.`

Like `.ifc`, but the sense of the test is reversed: this assembles the following section of code if the two strings are not the same.

`.ifndef symbol`

`.ifnotdef symbol`

Assembles the following section of code if the specified *symbol* has not been defined. Both spelling variants are equivalent.

.ifne *absolute expression*

Assembles the following section of code if the argument is not equal to zero (in other words, this is equivalent to `.if`).

.ifnes *string1, string2*

Like `.ifeqs`, but the sense of the test is reversed: this assembles the following section of code if the two strings are not the same.

7.34 .include "file"

This directive provides a way to include supporting files at specified points in your source program. The code from *file* is assembled as if it followed the point of the `.include`; when the end of the included file is reached, assembly of the original file continues. You can control the search paths used with the `-I` command-line option (see [Chapter 2 \[Command-Line Options\]](#), page 9). Quotation marks are required around *file*.

7.35 .int expressions

Expect zero or more *expressions*, of any section, separated by commas. For each expression, emit a number that, at run time, is the value of that expression. The byte order and bit size of the number depends on what kind of target the assembly is for.

7.36 .irp symbol, values...

Evaluate a sequence of statements assigning different values to *symbol*. The sequence of statements starts at the `.irp` directive, and is terminated by an `.endr` directive. For each *value*, *symbol* is set to *value*, and the sequence of statements is assembled. If no *value* is listed, the sequence of statements is assembled once, with *symbol* set to the null string. To refer to *symbol* within the sequence of statements, use `\symbol`.

For example, assembling

```
.irp    param,1,2,3
move   d\param,sp@-
.endr
```

is equivalent to assembling

```
move   d1,sp@-
move   d2,sp@-
move   d3,sp@-
```

7.37 .irpc symbol, values...

Evaluate a sequence of statements assigning different values to *symbol*. The sequence of statements starts at the `.irpc` directive, and is terminated by an `.endr` directive. For each character in *value*, *symbol* is set to the character, and the sequence of statements is assembled. If no *value* is listed, the sequence of statements is assembled once, with *symbol* set to the null string. To refer to *symbol* within the sequence of statements, use `\symbol`.

For example, assembling

```
.irpc    param,123
move    d\param,sp@-
.endr
```

is equivalent to assembling

```
move    d1,sp@-
move    d2,sp@-
move    d3,sp@-
```

7.38 `.lcomm symbol , length`

Reserve *length* (an absolute expression) bytes for a local common denoted by *symbol*. The section and value of *symbol* are those of the new local common. The addresses are allocated in the bss section, so that at run-time the bytes start off zeroed. *Symbol* is not declared global (see [Section 7.30 \[.global\], page 38](#)), so is normally not visible to `ld`.

Some targets permit a third argument to be used with `.lcomm`. This argument specifies the desired alignment of the symbol in the bss section.

The syntax for `.lcomm` differs slightly on the HPPA. The syntax is '*symbol .lcomm, length*'; *symbol* is optional.

7.39 `.lflags`

`as` accepts this directive, for compatibility with other assemblers, but ignores it.

7.40 `.line line-number`

Change the logical line number. *line-number* must be an absolute expression. The next line has that logical line number. Therefore any other statements on the current line (after a statement separator character) are reported as on logical line number *line-number* - 1. One day `as` will no longer support this directive: it is recognized only for compatibility with existing assembler programs.

Warning: In the AMD29K configuration of `as`, this command is not available; use the synonym `.ln` in that context.

Even though this is a directive associated with the `a.out` or `b.out` object-code formats, `as` still recognizes it when producing COFF output, and treats '`.line`' as though it were the COFF '`.ln`' if it is found outside a `.def/.endef` pair.

Inside a `.def`, '`.line`' is, instead, one of the directives used by compilers to generate auxiliary symbol information for debugging.

7.41 `.linkonce [type]`

Mark the current section so that the linker only includes a single copy of it. This may be used to include the same section in several different object files, but ensure that the linker will only include it once in the final output file. The `.linkonce` pseudo-op must be used for each instance of the section. Duplicate sections are detected based on the section name, so it should be unique.

This directive is only supported by a few object file formats; as of this writing, the only object file format which supports it is the Portable Executable format used on Windows NT.

The *type* argument is optional. If specified, it must be one of the following strings. For example:

```
.linkonce same_size
```

Not all types may be supported on all object file formats.

discard Silently discard duplicate sections. This is the default.

one_only Warn if there are duplicate sections, but still keep only one copy.

same_size
Warn if any of the duplicates have different sizes.

same_contents
Warn if any of the duplicates do not have exactly the same contents.

7.42 `.ln` *line-number*

`.ln` is a synonym for `.line`.

7.43 `.mri` *val*

If *val* is non-zero, this tells `as` to enter MRI mode. If *val* is zero, this tells `as` to exit MRI mode. This change affects code assembled until the next `.mri` directive, or until the end of the file. See [Section 2.7 \[MRI mode\], page 10](#).

7.44 `.list`

Control (in conjunction with the `.nolist` directive) whether or not assembly listings are generated. These two directives maintain an internal counter (which is zero initially). `.list` increments the counter, and `.nolist` decrements it. Assembly listings are generated whenever the counter is greater than zero.

By default, listings are disabled. When you enable them (with the `-a` command line option; see [Chapter 2 \[Command-Line Options\], page 9](#)), the initial value of the listing counter is one.

7.45 `.long` *expressions*

`.long` is the same as `.int`, see [Section 7.35 \[.int\], page 40](#).

7.46 .macro

The commands `.macro` and `.endm` allow you to define macros that generate assembly output. For example, this definition specifies a macro `sum` that puts a sequence of numbers into memory:

```
.macro sum from=0, to=5
.long \from
.if \to-\from
sum "(\from+1)",\to
.endif
.endm
```

With that definition, `SUM 0,5` is equivalent to this assembly input:

```
.long 0
.long 1
.long 2
.long 3
.long 4
.long 5
```

`.macro macname`

`.macro macname macargs ...`

Begin the definition of a macro called *macname*. If your macro definition requires arguments, specify their names after the macro name, separated by commas or spaces. You can supply a default value for any macro argument by following the name with `=deft`. For example, these are all valid `.macro` statements:

`.macro comm`

Begin the definition of a macro called *comm*, which takes no arguments.

`.macro plus1 p, p1`

`.macro plus1 p p1`

Either statement begins the definition of a macro called *plus1*, which takes two arguments; within the macro definition, write `\p` or `\p1` to evaluate the arguments.

`.macro reserve_str p1=0 p2`

Begin the definition of a macro called *reserve_str*, with two arguments. The first argument has a default value, but not the second. After the definition is complete, you can call the macro either as `reserve_str a,b` (with `\p1` evaluating to *a* and `\p2` evaluating to *b*), or as `reserve_str ,b` (with `\p1` evaluating as the default, in this case `0`, and `\p2` evaluating to *b*).

When you call a macro, you can specify the argument values either by position, or by keyword. For example, `sum 9,17` is equivalent to `sum to=17, from=9`.

`.endm`

Mark the end of a macro definition.

`.exitm`

Exit early from the current macro definition.

`\@` `as` maintains a counter of how many macros it has executed in this pseudo-variable; you can copy that number to your output with ‘`\@`’, but *only within a macro definition*.

7.47 `.nolist`

Control (in conjunction with the `.list` directive) whether or not assembly listings are generated. These two directives maintain an internal counter (which is zero initially). `.list` increments the counter, and `.nolist` decrements it. Assembly listings are generated whenever the counter is greater than zero.

7.48 `.octa bignums`

This directive expects zero or more bignums, separated by commas. For each bignum, it emits a 16-byte integer.

The term “octa” comes from contexts in which a “word” is two bytes; hence *octa*-word for 16 bytes.

7.49 `.org new-lc , fill`

Advance the location counter of the current section to *new-lc*. *new-lc* is either an absolute expression or an expression with the same section as the current subsection. That is, you can’t use `.org` to cross sections: if *new-lc* has the wrong section, the `.org` directive is ignored. To be compatible with former assemblers, if the section of *new-lc* is absolute, `as` issues a warning, then pretends the section of *new-lc* is the same as the current subsection.

`.org` may only increase the location counter, or leave it unchanged; you cannot use `.org` to move the location counter backwards.

Because `as` tries to assemble programs in one pass, *new-lc* may not be undefined. If you really detest this restriction we eagerly await a chance to share your improved assembler.

Beware that the origin is relative to the start of the section, not to the start of the subsection. This is compatible with other people’s assemblers.

When the location counter (of the current subsection) is advanced, the intervening bytes are filled with *fill* which should be an absolute expression. If the comma and *fill* are omitted, *fill* defaults to zero.

7.50 `.p2align[w1] abs-expr, abs-expr, abs-expr`

Pad the location counter (in the current subsection) to a particular storage boundary. The first expression (which must be absolute) is the number of low-order zero bits the location counter must have after advancement. For example ‘`.p2align 3`’ advances the location counter until it is a multiple of 8. If the location counter is already a multiple of 8, no change is needed.

The second expression (also absolute) gives the fill value to be stored in the padding bytes. It (and the comma) may be omitted. If it is omitted, the padding bytes are normally

zero. However, on some systems, if the section is marked as containing code and the fill value is omitted, the space is filled with no-op instructions.

The third expression is also absolute, and is also optional. If it is present, it is the maximum number of bytes that should be skipped by this alignment directive. If doing the alignment would require skipping more bytes than the specified maximum, then the alignment is not done at all. You can omit the fill value (the second argument) entirely by simply using two commas after the required alignment; this can be useful if you want the alignment to be filled with no-op instructions when appropriate.

The `.p2alignw` and `.p2alignl` directives are variants of the `.p2align` directive. The `.p2alignw` directive treats the fill pattern as a two byte word value. The `.p2alignl` directive treats the fill pattern as a four byte longword value. For example, `.p2alignw 2,0x368d` will align to a multiple of 4. If it skips two bytes, they will be filled in with the value `0x368d` (the exact placement of the bytes depends upon the endianness of the processor). If it skips 1 or 3 bytes, the fill value is undefined.

7.51 `.print string`

`as` will print *string* on the standard output during assembly. You must put *string* in double quotes.

7.52 `.psize lines , columns`

Use this directive to declare the number of lines—and, optionally, the number of columns—to use for each page, when generating listings.

If you do not use `.psize`, listings use a default line-count of 60. You may omit the comma and *columns* specification; the default width is 200 columns.

`as` generates formfeeds whenever the specified number of lines is exceeded (or whenever you explicitly request one, using `.eject`).

If you specify *lines* as 0, no formfeeds are generated save those explicitly specified with `.eject`.

7.53 `.purgem name`

Undefine the macro *name*, so that later uses of the string will not be expanded. See [Section 7.46 \[Macro\]](#), page 43.

7.54 `.quad bignums`

`.quad` expects zero or more bignums, separated by commas. For each bignum, it emits an 8-byte integer. If the bignum won't fit in 8 bytes, it prints a warning message; and just takes the lowest order 8 bytes of the bignum.

The term “quad” comes from contexts in which a “word” is two bytes; hence *quad*-word for 8 bytes.

7.55 `.rept count`

Repeat the sequence of lines between the `.rept` directive and the next `.endr` directive *count* times.

For example, assembling

```
.rept 3
.long 0
.endr
```

is equivalent to assembling

```
.long 0
.long 0
.long 0
```

7.56 `.sbttl "subheading"`

Use *subheading* as the title (third line, immediately after the title line) when generating assembly listings.

This directive affects subsequent pages, as well as the current page if it appears within ten lines of the top of a page.

7.57 `.scl class`

Set the storage-class value for a symbol. This directive may only be used inside a `.def/.endef` pair. Storage class may flag whether a symbol is static or external, or it may record further symbolic debugging information.

The `‘.scl’` directive is primarily associated with COFF output; when configured to generate `b.out` output format, `as` accepts this directive but ignores it.

7.58 `.section name`

Use the `.section` directive to assemble the following code into a section named *name*.

This directive is only supported for targets that actually support arbitrarily named sections; on `a.out` targets, for example, it is not accepted, even with a standard `a.out` section name.

For COFF targets, the `.section` directive is used in one of the following ways:

```
.section name[, "flags"]
.section name[, subsegment]
```

If the optional argument is quoted, it is taken as flags to use for the section. Each flag is a single character. The following flags are recognized:

<code>b</code>	bss section (uninitialized data)
<code>n</code>	section is not loaded
<code>w</code>	writable section
<code>d</code>	data section

r	read-only section
x	executable section
s	shared section (meaningful for PE targets)

If no flags are specified, the default flags depend upon the section name. If the section name is not recognized, the default will be for the section to be loaded and writable.

If the optional argument to the `.section` directive is not quoted, it is taken as a sub-segment number (see [Section 4.4 \[Sub-Sections\], page 23](#)).

For ELF targets, the `.section` directive is used like this:

```
.section name[, "flags"[, @type]]
```

The optional *flags* argument is a quoted string which may contain any combination of the following characters:

a	section is allocatable
w	section is writable
x	section is executable

The optional *type* argument may contain one of the following constants:

`@progbits`

section contains data

`@nobits`

section does not contain data (i.e., section only occupies space)

If no flags are specified, the default flags depend upon the section name. If the section name is not recognized, the default will be for the section to have none of the above flags: it will not be allocated in memory, nor writable, nor executable. The section will contain data.

For ELF targets, the assembler supports another type of `.section` directive for compatibility with the Solaris assembler:

```
.section "name"[, flags...]
```

Note that the section name is quoted. There may be a sequence of comma separated flags:

<code>#alloc</code>	section is allocatable
<code>#write</code>	section is writable
<code>#execinstr</code>	section is executable

7.59 `.set symbol, expression`

Set the value of *symbol* to *expression*. This changes *symbol*'s value and type to conform to *expression*. If *symbol* was flagged as external, it remains flagged (see [Section 5.5 \[Symbol Attributes\], page 28](#)).

You may `.set` a symbol many times in the same assembly.

If you `.set` a global symbol, the value stored in the object file is the last value stored into it.

The syntax for `set` on the HPPA is '`symbol .set expression`'.

7.60 `.short` *expressions*

`.short` is normally the same as `‘.word’`. See [Section 7.77 \[‘.word’\]](#), page 52.

In some configurations, however, `.short` and `.word` generate numbers of different lengths; see [Chapter 8 \[Machine Dependencies\]](#), page 53.

7.61 `.single` *flonums*

This directive assembles zero or more flonums, separated by commas. It has the same effect as `.float`. The exact kind of floating point numbers emitted depends on how `as` is configured. See [Chapter 8 \[Machine Dependencies\]](#), page 53.

7.62 `.size`

This directive is generated by compilers to include auxiliary debugging information in the symbol table. It is only permitted inside `.def/.endif` pairs.

`‘.size’` is only meaningful when generating COFF format output; when `as` is generating `b.out`, it accepts this directive but ignores it.

7.63 `.sleb128` *expressions*

sleb128 stands for “signed little endian base 128.” This is a compact, variable length representation of numbers used by the DWARF symbolic debugging format. See [Section 7.75 \[Uleb128\]](#), page 51.

7.64 `.skip` *size* , *fill*

This directive emits *size* bytes, each of value *fill*. Both *size* and *fill* are absolute expressions. If the comma and *fill* are omitted, *fill* is assumed to be zero. This is the same as `‘.space’`.

7.65 `.space` *size* , *fill*

This directive emits *size* bytes, each of value *fill*. Both *size* and *fill* are absolute expressions. If the comma and *fill* are omitted, *fill* is assumed to be zero. This is the same as `‘.skip’`.

Warning: `.space` has a completely different meaning for HPPA targets; use `.block` as a substitute. See *HP9000 Series 800 Assembly Language Reference Manual* (HP 92432-90001) for the meaning of the `.space` directive. See [Section 8.8.5 \[HPPA Assembler Directives\]](#), page 75, for a summary.

On the AMD 29K, this directive is ignored; it is accepted for compatibility with other AMD 29K assemblers.

Warning: In most versions of the GNU assembler, the directive `.space` has the effect of `.block`. See [Chapter 8 \[Machine Dependencies\]](#), page 53.

7.66 .stabd, .stabn, .stabs

There are three directives that begin ‘.stab’. All emit symbols (see [Chapter 5 \[Symbols\], page 27](#)), for use by symbolic debuggers. The symbols are not entered in the `as` hash table: they cannot be referenced elsewhere in the source file. Up to five fields are required:

<i>string</i>	This is the symbol’s name. It may contain any character except ‘\000’, so is more general than ordinary symbol names. Some debuggers used to code arbitrarily complex structures into symbol names using this field.
<i>type</i>	An absolute expression. The symbol’s type is set to the low 8 bits of this expression. Any bit pattern is permitted, but <code>ld</code> and debuggers choke on silly bit patterns.
<i>other</i>	An absolute expression. The symbol’s “other” attribute is set to the low 8 bits of this expression.
<i>desc</i>	An absolute expression. The symbol’s descriptor is set to the low 16 bits of this expression.
<i>value</i>	An absolute expression which becomes the symbol’s value.

If a warning is detected while reading a `.stabd`, `.stabn`, or `.stabs` statement, the symbol has probably already been created; you get a half-formed symbol in your object file. This is compatible with earlier assemblers!

`.stabd` *type* , *other* , *desc*

The “name” of the symbol generated is not even an empty string. It is a null pointer, for compatibility. Older assemblers used a null pointer so they didn’t waste space in object files with empty strings.

The symbol’s value is set to the location counter, relocatably. When your program is linked, the value of this symbol is the address of the location counter when the `.stabd` was assembled.

`.stabn` *type* , *other* , *desc* , *value*

The name of the symbol is set to the empty string “”.

`.stabs` *string* , *type* , *other* , *desc* , *value*

All five fields are specified.

7.67 .string "*str*"

Copy the characters in *str* to the object file. You may specify more than one string to copy, separated by commas. Unless otherwise specified for a particular machine, the assembler marks the end of each string with a 0 byte. You can use any of the escape sequences described in [Section 3.6.1.1 \[Strings\], page 17](#).

7.68 `.struct expression`

Switch to the absolute section, and set the section offset to *expression*, which must be an absolute expression. You might use this as follows:

```
.struct 0
field1:
    .struct field1 + 4
field2:
    .struct field2 + 4
field3:
```

This would define the symbol `field1` to have the value 0, the symbol `field2` to have the value 4, and the symbol `field3` to have the value 8. Assembly would be left in the absolute section, and you would need to use a `.section` directive of some sort to change to some other section before further assembly.

7.69 `.symver`

Use the `.symver` directive to bind symbols to specific version nodes within a source file. This is only supported on ELF platforms, and is typically used when assembling files to be linked into a shared library. There are cases where it may make sense to use this in objects to be bound into an application itself so as to override a versioned symbol from a shared library.

For ELF targets, the `.symver` directive is used like this:

```
.symver name, name2@nodename
```

In this case, the symbol *name* must exist and be defined within the file being assembled. The `.versym` directive effectively creates a symbol alias with the name *name2@nodename*, and in fact the main reason that we just don't try and create a regular alias is that the `@` character isn't permitted in symbol names. The *name2* part of the name is the actual name of the symbol by which it will be externally referenced. The name *name* itself is merely a name of convenience that is used so that it is possible to have definitions for multiple versions of a function within a single source file, and so that the compiler can unambiguously know which version of a function is being mentioned. The *nodename* portion of the alias should be the name of a node specified in the version script supplied to the linker when building a shared library. If you are attempting to override a versioned symbol from a shared library, then *nodename* should correspond to the nodename of the symbol you are trying to override.

7.70 `.tag structname`

This directive is generated by compilers to include auxiliary debugging information in the symbol table. It is only permitted inside `.def/.endef` pairs. Tags are used to link structure definitions in the symbol table with instances of those structures.

'`.tag`' is only used when generating COFF format output; when `as` is generating `b.out`, it accepts this directive but ignores it.

7.71 `.text` *subsection*

Tells `as` to assemble the following statements onto the end of the text subsection numbered *subsection*, which is an absolute expression. If *subsection* is omitted, subsection number zero is used.

7.72 `.title` "*heading*"

Use *heading* as the title (second line, immediately after the source file name and pagenumber) when generating assembly listings.

This directive affects subsequent pages, as well as the current page if it appears within ten lines of the top of a page.

7.73 `.type` *int*

This directive, permitted only within `.def/.endif` pairs, records the integer *int* as the type attribute of a symbol table entry.

'`.type`' is associated only with COFF format output; when `as` is configured for `b.out` output, it accepts this directive but ignores it.

7.74 `.val` *addr*

This directive, permitted only within `.def/.endif` pairs, records the address *addr* as the value attribute of a symbol table entry.

'`.val`' is used only for COFF output; when `as` is configured for `b.out`, it accepts this directive but ignores it.

7.75 `.uleb128` *expressions*

uleb128 stands for "unsigned little endian base 128." This is a compact, variable length representation of numbers used by the DWARF symbolic debugging format. See [Section 7.63 \[Sleb128\]](#), page 48.

7.76 `.internal`, `.hidden`, `.protected`

These directives can be used to set the visibility of a specified symbol. By default a symbol's visibility is set by its binding (local, global or weak), but these directives can be used to override that.

A visibility of `protected` means that any references to the symbol from within the component that defines the symbol must be resolved to the definition in that component, even if a definition in another component would normally preempt this.

A visibility of `hidden` means that the symbol is not visible to other components. Such a symbol is always considered to be protected as well.

A visibility of `internal` is the same as a visibility of `hidden`, except that some extra, processor specific processing must also be performed upon the symbol.

For ELF targets, the directives are used like this:

```
.internal name
.hidden name
.protected name
```

7.77 `.word` *expressions*

This directive expects zero or more *expressions*, of any section, separated by commas.

The size of the number emitted, and its byte order, depend on what target computer the assembly is for.

Warning: Special Treatment to support Compilers

Machines with a 32-bit address space, but that do less than 32-bit addressing, require the following special treatment. If the machine of interest to you does 32-bit addressing (or doesn't require it; see [Chapter 8 \[Machine Dependencies\], page 53](#)), you can ignore this issue.

In order to assemble compiler output into something that works, `as` occasionally does strange things to `.word` directives. Directives of the form `.word sym1-sym2` are often emitted by compilers as part of jump tables. Therefore, when `as` assembles a directive of the form `.word sym1-sym2`, and the difference between `sym1` and `sym2` does not fit in 16 bits, `as` creates a *secondary jump table*, immediately before the next label. This secondary jump table is preceded by a short-jump to the first byte after the secondary table. This short-jump prevents the flow of control from accidentally falling into the new table. Inside the table is a long-jump to `sym2`. The original `.word` contains `sym1` minus the address of the long-jump to `sym2`.

If there were several occurrences of `.word sym1-sym2` before the secondary jump table, all of them are adjusted. If there was a `.word sym3-sym4`, that also did not fit in sixteen bits, a long-jump to `sym4` is included in the secondary jump table, and the `.word` directives are adjusted to contain `sym3` minus the address of the long-jump to `sym4`; and so on, for as many entries in the original jump table as necessary.

7.78 *Deprecated Directives*

One day these directives won't work. They are included for compatibility with older assemblers.

```
.abort
.line
```

8 Machine Dependent Features

The machine instruction sets are (almost by definition) different on each machine where `as` runs. Floating point representations vary as well, and `as` often supports a few additional directives or command-line options for compatibility with other assemblers on a particular platform. Finally, some versions of `as` support special pseudo-instructions for branch optimization.

This chapter discusses most of these differences, though it does not include details on any machine's instruction set. For details on that subject, see the hardware manufacturer's manual.

8.1 ARC Dependent Features

8.1.1 Options

The ARC chip family includes several successive levels (or other variants) of chip, using the same core instruction set, but including a few additional instructions at each level.

By default, `as` assumes the core instruction set (ARC base). The `.cpu` pseudo-op is intended to be used to select the variant.

`-mbig-endian`
`-mlittle-endian`

Any ARC configuration of `as` can select big-endian or little-endian output at run time (unlike most other GNU development tools, which must be configured for one or the other). Use `'-mbig-endian'` to select big-endian output, and `'-mlittle-endian'` for little-endian.

8.1.2 Floating Point

The ARC cpu family currently does not have hardware floating point support. Software floating point support is provided by `GCC` and uses IEEE floating-point numbers.

8.1.3 ARC Machine Directives

The ARC version of `as` supports the following additional machine directives:

`.cpu` This must be followed by the desired cpu. The ARC is intended to be customizable, `.cpu` is used to select the desired variant [though currently there are none].

8.2 AMD 29K Dependent Features

8.2.1 Options

`as` has no additional command-line options for the AMD 29K family.

8.2.2 Syntax

8.2.2.1 Macros

The macro syntax used on the AMD 29K is like that described in the AMD 29K Family Macro Assembler Specification. Normal `as` macros should still work.

8.2.2.2 Special Characters

‘;’ is the line comment character.

The character ‘?’ is permitted in identifiers (but may not begin an identifier).

8.2.2.3 Register Names

General-purpose registers are represented by predefined symbols of the form ‘`GR nnn` ’ (for global registers) or ‘`LR nnn` ’ (for local registers), where nnn represents a number between 0 and 127, written with no leading zeros. The leading letters may be in either upper or lower case; for example, ‘`gr13`’ and ‘`LR7`’ are both valid register names.

You may also refer to general-purpose registers by specifying the register number as the result of an expression (prefixed with ‘`%`’ to flag the expression as a register number):

`%%expression`

—where *expression* must be an absolute expression evaluating to a number between 0 and 255. The range [0, 127] refers to global registers, and the range [128, 255] to local registers.

In addition, `as` understands the following protected special-purpose register names for the AMD 29K family:

<code>vab</code>	<code>chd</code>	<code>pc0</code>
<code>ops</code>	<code>chc</code>	<code>pc1</code>
<code>cps</code>	<code>rbp</code>	<code>pc2</code>
<code>cfg</code>	<code>tmc</code>	<code>mmu</code>
<code>cha</code>	<code>tmr</code>	<code>lru</code>

These unprotected special-purpose register names are also recognized:

<code>ipc</code>	<code>alu</code>	<code>fpe</code>
<code>ipa</code>	<code>bp</code>	<code>inte</code>
<code>ipb</code>	<code>fc</code>	<code>fps</code>
<code>q</code>	<code>cr</code>	<code>exop</code>

8.2.3 Floating Point

The AMD 29K family uses IEEE floating-point numbers.

8.2.4 AMD 29K Machine Directives

.block *size* , *fill*

This directive emits *size* bytes, each of value *fill*. Both *size* and *fill* are absolute expressions. If the comma and *fill* are omitted, *fill* is assumed to be zero.

In other versions of the GNU assembler, this directive is called ‘.space’.

.cputype This directive is ignored; it is accepted for compatibility with other AMD 29K assemblers.

.file This directive is ignored; it is accepted for compatibility with other AMD 29K assemblers.

Warning: in other versions of the GNU assembler, **.file** is used for the directive called **.app-file** in the AMD 29K support.

.line This directive is ignored; it is accepted for compatibility with other AMD 29K assemblers.

.sect This directive is ignored; it is accepted for compatibility with other AMD 29K assemblers.

.use *section name*

Establishes the section and subsection for the following code; *section name* may be one of **.text**, **.data**, **.data1**, or **.lit**. With one of the first three *section name* options, ‘.use’ is equivalent to the machine directive *section name*; the remaining case, ‘.use **.lit**’, is the same as ‘.data 200’.

8.2.5 Opcodes

as implements all the standard AMD 29K opcodes. No additional pseudo-instructions are needed on this family.

For information on the 29K machine instruction set, see *Am29000 User’s Manual*, Advanced Micro Devices, Inc.

8.3 ARM Dependent Features

8.3.1 Options

`-marm`

`[2|250|3|6|60|600|610|620|7|7m|7d|7dm|7di|7dmi|70|700|700i|710|710c|7100|7500|7500fe|7td`

This option specifies the target processor. The assembler will issue an error message if an attempt is made to assemble an instruction which will not execute on the target processor.

`-marmv [2|2a|3|3m|4|4t|5|5t]`

This option specifies the target architecture. The assembler will issue an error message if an attempt is made to assemble an instruction which will not execute on the target architecture.

`-mthumb` This option specifies that only Thumb instructions should be assembled.

`-mall` This option specifies that any Arm or Thumb instruction should be assembled.

`-mfpa [10|11]`

This option specifies the floating point architecture in use on the target processor.

`-mfpe-old`

Do not allow the assemble of floating point multiple instructions.

`-mno-fpu` Do not allow the assembly of any floating point instructions.

`-mthumb-interwork`

This option specifies that the output generated by the assembler should be marked as supporting interworking.

`-mapcs [26|32]`

This option specifies that the output generated by the assembler should be marked as supporting the indicated version of the Arm Procedure. Calling Standard.

`-mapcs-float`

This indicates the the floating point variant of the APCS should be used. In this variant floating point arguments are passed in FP registers rather than integer registers.

`-mapcs-reentrant`

This indicates that the reentrant variant of the APCS should be used. This variant supports position independent code.

`-EB` This option specifies that the output generated by the assembler should be marked as being encoded for a big-endian processor.

`-EL` This option specifies that the output generated by the assembler should be marked as being encoded for a little-endian processor.

- k This option enables the generation of PIC (position independent code).
- moabi This indicates that the code should be assembled using the old ARM ELF conventions, based on a beta release of the ARM-ELF specifications, rather than the default conventions which are based on the final release of the ARM-ELF specifications.

8.3.2 Syntax

8.3.2.1 Special Characters

The presence of a ‘@’ on a line indicates the start of a comment that extends to the end of the current line. If a ‘#’ appears as the first character of a line, the whole line is treated as a comment.

On ARM systems running the GNU/Linux operating system, ‘;’ can be used instead of a newline to separate statements.

Either ‘#’ or ‘\$’ can be used to indicate immediate operands.

TODO Explain about /data modifier on symbols.

8.3.2.2 Register Names

TODO Explain about ARM register naming, and the predefined names.

8.3.3 Floating Point

The ARM family uses IEEE floating-point numbers.

8.3.4 ARM Machine Directives

name .req *register name*

This creates an alias for *register name* called *name*. For example:

```
foo .req r0
```

.code [16|32]

This directive selects the instruction set being generated. The value 16 selects Thumb, with the value 32 selecting ARM.

.thumb This performs the same action as *.code 16*.

.arm This performs the same action as *.code 32*.

.force_thumb

This directive forces the selection of Thumb instructions, even if the target processor does not support those instructions

.thumb_func

This directive specifies that the following symbol is the name of a Thumb encoded function. This information is necessary in order to allow the assembler and linker to generate correct code for interworking between Arm and Thumb instructions and should be used even if interworking is not going to be performed.

.thumb_set

This performs the equivalent of a `.set` directive in that it creates a symbol which is an alias for another symbol (possibly not yet defined). This directive also has the added property in that it marks the aliased symbol as being a thumb function entry point, in the same way that the `.thumb_func` directive does.

.ltorg

This directive causes the current contents of the literal pool to be dumped into the current section (which is assumed to be the `.text` section) at the current location (aligned to a word boundary).

.pool

This is a synonym for `.ltorg`.

8.3.5 Opcodes

`as` implements all the standard ARM opcodes. It also implements several pseudo opcodes, including several synthetic load instructions.

NOP

`nop`

This pseudo op will always evaluate to a legal ARM instruction that does nothing. Currently it will evaluate to `MOV r0, r0`.

LDR

`ldr <register> , = <expression>`

If expression evaluates to a numeric constant then a `MOV` or `MVN` instruction will be used in place of the `LDR` instruction, if the constant can be generated by either of these instructions. Otherwise the constant will be placed into the nearest literal pool (if it not already there) and a PC relative `LDR` instruction will be generated.

ADR

`adr <register> <label>`

This instruction will load the address of *label* into the indicated register. The instruction will evaluate to a PC relative `ADD` or `SUB` instruction depending upon where the label is located. If the label is out of range, or if it is not defined in the same file (and section) as the `ADR` instruction, then an error will be generated. This instruction will not make use of the literal pool.

ADRL

`adrl <register> <label>`

This instruction will load the address of *label* into the indicated register. The instruction will evaluate to one or two a PC relative `ADD` or `SUB` instructions depending upon where the label is located. If a second instruction is not needed a `NOP` instruction will be generated in its place, so that this instruction is always 8 bytes long.

If the label is out of range, or if it is not defined in the same file (and section) as the `ADRL` instruction, then an error will be generated. This instruction will not make use of the literal pool.

For information on the ARM or Thumb instruction sets, see *ARM Software Development Toolkit Reference Manual*, Advanced RISC Machines Ltd.

8.4 D10V Dependent Features

8.4.1 D10V Options

The Mitsubishi D10V version of `as` has a few machine dependent options.

`-O` The D10V can often execute two sub-instructions in parallel. When this option is used, `as` will attempt to optimize its output by detecting when instructions can be executed in parallel.

`--nowarnswap`
To optimize execution performance, `as` will sometimes swap the order of instructions. Normally this generates a warning. When this option is used, no warning will be generated when instructions are swapped.

8.4.2 Syntax

The D10V syntax is based on the syntax in Mitsubishi's D10V architecture manual. The differences are detailed below.

8.4.2.1 Size Modifiers

The D10V version of `as` uses the instruction names in the D10V Architecture Manual. However, the names in the manual are sometimes ambiguous. There are instruction names that can assemble to a short or long form opcode. How does the assembler pick the correct form? `as` will always pick the smallest form if it can. When dealing with a symbol that is not defined yet when a line is being assembled, it will always use the long form. If you need to force the assembler to use either the short or long form of the instruction, you can append either `.s` (short) or `.l` (long) to it. For example, if you are writing an assembly program and you want to do a branch to a symbol that is defined later in your program, you can write `bra.s foo`. `Objdump` and `GDB` will always append `.s` or `.l` to instructions which have both short and long forms.

8.4.2.2 Sub-Instructions

The D10V assembler takes as input a series of instructions, either one-per-line, or in the special two-per-line format described in the next section. Some of these instructions will be short-form or sub-instructions. These sub-instructions can be packed into a single instruction. The assembler will do this automatically. It will also detect when it should not pack instructions. For example, when a label is defined, the next instruction will never be packaged with the previous one. Whenever a branch and link instruction is called, it will not be packaged with the next instruction so the return address will be valid. Nops are automatically inserted when necessary.

If you do not want the assembler automatically making these decisions, you can control the packaging and execution type (parallel or sequential) with the special execution symbols described in the next section.

8.4.2.3 Special Characters

‘;’ and ‘#’ are the line comment characters. Sub-instructions may be executed in order, in reverse-order, or in parallel. Instructions listed in the standard one-per-line format will be executed sequentially. To specify the executing order, use the following symbols:

‘->’ Sequential with instruction on the left first.
 ‘<-’ Sequential with instruction on the right first.
 ‘||’ Parallel

The D10V syntax allows either one instruction per line, one instruction per line with the execution symbol, or two instructions per line. For example

```
abs a1 -> abs r0
```

Execute these sequentially. The instruction on the right is in the right container and is executed second.

```
abs r0 <- abs a1
```

Execute these reverse-sequentially. The instruction on the right is in the right container, and is executed first.

```
ld2w r2,@r8+ || mac a0,r0,r7
```

Execute these in parallel.

```
ld2w r2,@r8+ ||
mac a0,r0,r7
```

Two-line format. Execute these in parallel.

```
ld2w r2,@r8+
mac a0,r0,r7
```

Two-line format. Execute these sequentially. Assembler will put them in the proper containers.

```
ld2w r2,@r8+ ->
mac a0,r0,r7
```

Two-line format. Execute these sequentially. Same as above but second instruction will always go into right container.

Since ‘\$’ has no special meaning, you may use it in symbol names.

8.4.2.4 Register Names

You can use the predefined symbols ‘r0’ through ‘r15’ to refer to the D10V registers. You can also use ‘sp’ as an alias for ‘r15’. The accumulators are ‘a0’ and ‘a1’. There are special register-pair names that may optionally be used in opcodes that require even-numbered registers. Register names are not case sensitive.

Register Pairs

```
r0-r1
r2-r3
r4-r5
```


r6-r7

r8-r9

r10-r11

r12-r13

r14-r15

The D10V also has predefined symbols for these control registers and status bits:

psw	Processor Status Word
bpsw	Backup Processor Status Word
pc	Program Counter
bpc	Backup Program Counter
rpt_c	Repeat Count
rpt_s	Repeat Start address
rpt_e	Repeat End address
mod_s	Modulo Start address
mod_e	Modulo End address
iba	Instruction Break Address
f0	Flag 0
f1	Flag 1
c	Carry flag

8.4.2.5 Addressing Modes

as understands the following addressing modes for the D10V. *Rn* in the following refers to any of the numbered registers, but *not* the control registers.

<i>Rn</i>	Register direct
@ <i>Rn</i>	Register indirect
@ <i>Rn</i> +	Register indirect with post-increment
@ <i>Rn</i> -	Register indirect with post-decrement
@-SP	Register indirect with pre-decrement
@(<i>disp</i> , <i>Rn</i>)	Register indirect with displacement
<i>addr</i>	PC relative address (for branch or rep).
# <i>imm</i>	Immediate data (the '#' is optional and ignored)

8.4.2.6 @WORD Modifier

Any symbol followed by `@word` will be replaced by the symbol's value shifted right by 2. This is used in situations such as loading a register with the address of a function (or any other code fragment). For example, if you want to load a register with the location of the function `main` then jump to that function, you could do it as follows:

```
ldi    r2, main@word
jmp    r2
```

8.4.3 Floating Point

The D10V has no hardware floating point, but the `.float` and `.double` directives generates IEEE floating-point numbers for compatibility with other development tools.

8.4.4 Opcodes

For detailed information on the D10V machine instruction set, see *D10V Architecture: A VLIW Microprocessor for Multimedia Applications* (Mitsubishi Electric Corp.). `as` implements all the standard D10V opcodes. The only changes are those described in the section on size modifiers

8.5 D30V Dependent Features

8.5.1 D30V Options

The Mitsubishi D30V version of `as` has a few machine dependent options.

- '-O' The D30V can often execute two sub-instructions in parallel. When this option is used, `as` will attempt to optimize its output by detecting when instructions can be executed in parallel.
- '-n' When this option is used, `as` will issue a warning every time it adds a nop instruction.
- '-N' When this option is used, `as` will issue a warning if it needs to insert a nop after a 32-bit multiply before a load or 16-bit multiply instruction.

8.5.2 Syntax

The D30V syntax is based on the syntax in Mitsubishi's D30V architecture manual. The differences are detailed below.

8.5.2.1 Size Modifiers

The D30V version of `as` uses the instruction names in the D30V Architecture Manual. However, the names in the manual are sometimes ambiguous. There are instruction names that can assemble to a short or long form opcode. How does the assembler pick the correct form? `as` will always pick the smallest form if it can. When dealing with a symbol that is not defined yet when a line is being assembled, it will always use the long form. If you need to force the assembler to use either the short or long form of the instruction, you can append either `.s` (short) or `.l` (long) to it. For example, if you are writing an assembly program and you want to do a branch to a symbol that is defined later in your program, you can write `bra.s foo`. `Objdump` and `GDB` will always append `.s` or `.l` to instructions which have both short and long forms.

8.5.2.2 Sub-Instructions

The D30V assembler takes as input a series of instructions, either one-per-line, or in the special two-per-line format described in the next section. Some of these instructions will be short-form or sub-instructions. These sub-instructions can be packed into a single instruction. The assembler will do this automatically. It will also detect when it should not pack instructions. For example, when a label is defined, the next instruction will never be packaged with the previous one. Whenever a branch and link instruction is called, it will not be packaged with the next instruction so the return address will be valid. Nops are automatically inserted when necessary.

If you do not want the assembler automatically making these decisions, you can control the packaging and execution type (parallel or sequential) with the special execution symbols described in the next section.

8.5.2.3 Special Characters

‘;’ and ‘#’ are the line comment characters. Sub-instructions may be executed in order, in reverse-order, or in parallel. Instructions listed in the standard one-per-line format will be executed sequentially unless you use the ‘-0’ option.

To specify the executing order, use the following symbols:

‘->’ Sequential with instruction on the left first.
 ‘<-’ Sequential with instruction on the right first.
 ‘||’ Parallel

The D30V syntax allows either one instruction per line, one instruction per line with the execution symbol, or two instructions per line. For example

```
abs r2,r3 -> abs r4,r5
```

Execute these sequentially. The instruction on the right is in the right container and is executed second.

```
abs r2,r3 <- abs r4,r5
```

Execute these reverse-sequentially. The instruction on the right is in the right container, and is executed first.

```
abs r2,r3 || abs r4,r5
```

Execute these in parallel.

```
ldw r2,@(r3,r4) ||
```

```
mulx r6,r8,r9
```

Two-line format. Execute these in parallel.

```
mulx a0,r8,r9
```

```
stw r2,@(r3,r4)
```

Two-line format. Execute these sequentially unless ‘-0’ option is used. If the ‘-0’ option is used, the assembler will determine if the instructions could be done in parallel (the above two instructions can be done in parallel), and if so, emit them as parallel instructions. The assembler will put them in the proper containers. In the above example, the assembler will put the ‘stw’ instruction in left container and the ‘mulx’ instruction in the right container.

```
stw r2,@(r3,r4) ->
```

```
mulx a0,r8,r9
```

Two-line format. Execute the ‘stw’ instruction followed by the ‘mulx’ instruction sequentially. The first instruction goes in the left container and the second instruction goes into right container. The assembler will give an error if the machine ordering constraints are violated.

```
stw r2,@(r3,r4) <-
```

```
mulx a0,r8,r9
```

Same as previous example, except that the ‘mulx’ instruction is executed before the ‘stw’ instruction.

Since ‘\$’ has no special meaning, you may use it in symbol names.

8.5.2.4 Guarded Execution

`as` supports the full range of guarded execution directives for each instruction. Just append the directive after the instruction proper. The directives are:

<code>/tx</code>	Execute the instruction if flag <code>f0</code> is true.
<code>/fx</code>	Execute the instruction if flag <code>f0</code> is false.
<code>/xt</code>	Execute the instruction if flag <code>f1</code> is true.
<code>/xf</code>	Execute the instruction if flag <code>f1</code> is false.
<code>/tt</code>	Execute the instruction if both flags <code>f0</code> and <code>f1</code> are true.
<code>/tf</code>	Execute the instruction if flag <code>f0</code> is true and flag <code>f1</code> is false.

8.5.2.5 Register Names

You can use the predefined symbols `r0` through `r63` to refer to the D30V registers. You can also use `sp` as an alias for `r63` and `link` as an alias for `r62`. The accumulators are `a0` and `a1`.

The D30V also has predefined symbols for these control registers and status bits:

<code>psw</code>	Processor Status Word
<code>bpsw</code>	Backup Processor Status Word
<code>pc</code>	Program Counter
<code>bpc</code>	Backup Program Counter
<code>rpt_c</code>	Repeat Count
<code>rpt_s</code>	Repeat Start address
<code>rpt_e</code>	Repeat End address
<code>mod_s</code>	Modulo Start address
<code>mod_e</code>	Modulo End address
<code>iba</code>	Instruction Break Address
<code>f0</code>	Flag 0
<code>f1</code>	Flag 1
<code>f2</code>	Flag 2
<code>f3</code>	Flag 3
<code>f4</code>	Flag 4
<code>f5</code>	Flag 5
<code>f6</code>	Flag 6
<code>f7</code>	Flag 7
<code>s</code>	Same as flag 4 (saturation flag)

<code>v</code>	Same as flag 5 (overflow flag)
<code>va</code>	Same as flag 6 (sticky overflow flag)
<code>c</code>	Same as flag 7 (carry/borrow flag)
<code>b</code>	Same as flag 7 (carry/borrow flag)

8.5.2.6 Addressing Modes

`as` understands the following addressing modes for the D30V. `Rn` in the following refers to any of the numbered registers, but *not* the control registers.

<code>Rn</code>	Register direct
<code>@Rn</code>	Register indirect
<code>@Rn+</code>	Register indirect with post-increment
<code>@Rn-</code>	Register indirect with post-decrement
<code>@-SP</code>	Register indirect with pre-decrement
<code>@(disp, Rn)</code>	Register indirect with displacement
<code>addr</code>	PC relative address (for branch or rep).
<code>#imm</code>	Immediate data (the '#' is optional and ignored)

8.5.3 Floating Point

The D30V has no hardware floating point, but the `.float` and `.double` directives generate IEEE floating-point numbers for compatibility with other development tools.

8.5.4 Opcodes

For detailed information on the D30V machine instruction set, see *D30V Architecture: A VLIW Microprocessor for Multimedia Applications* (Mitsubishi Electric Corp.). `as` implements all the standard D30V opcodes. The only changes are those described in the section on size modifiers

8.6 H8/300 Dependent Features

8.6.1 Options

as has no additional command-line options for the Hitachi H8/300 family.

8.6.2 Syntax

8.6.2.1 Special Characters

';' is the line comment character.

'\$' can be used instead of a newline to separate statements. Therefore *you may not use '\$' in symbol names* on the H8/300.

8.6.2.2 Register Names

You can use predefined symbols of the form '**rn**' and '**rn1**' to refer to the H8/300 registers as sixteen 8-bit general-purpose registers. *n* is a digit from '0' to '7'; for instance, both '**r0h**' and '**r7l**' are valid register names.

You can also use the eight predefined symbols '**rn**' to refer to the H8/300 registers as 16-bit registers (you must use this form for addressing).

On the H8/300H, you can also use the eight predefined symbols '**ern**' ('**er0**' . . . '**er7**') to refer to the 32-bit general purpose registers.

The two control registers are called **pc** (program counter; a 16-bit register, except on the H8/300H where it is 24 bits) and **ccr** (condition code register; an 8-bit register). **r7** is used as the stack pointer, and can also be called **sp**.

8.6.2.3 Addressing Modes

as understands the following addressing modes for the H8/300:

rn	Register direct
@rn	Register indirect
@(d, rn)	
@(d:16, rn)	
@(d:24, rn)	Register indirect: 16-bit or 24-bit displacement <i>d</i> from register <i>n</i> . (24-bit displacements are only meaningful on the H8/300H.)
@rn+	Register indirect with post-increment
@-rn	Register indirect with pre-decrement
@aa	
@aa:8	
@aa:16	
@aa:24	Absolute address aa . (The address size ':24' only makes sense on the H8/300H.)

`#xx`

`#xx:8`

`#xx:16`

`#xx:32` Immediate data `xx`. You may specify the `':8'`, `':16'`, or `':32'` for clarity, if you wish; but `as` neither requires this nor uses it—the data size required is taken from context.

`@@aa`

`@@aa:8` Memory indirect. You may specify the `':8'` for clarity, if you wish; but `as` neither requires this nor uses it.

8.6.3 Floating Point

The H8/300 family has no hardware floating point, but the `.float` directive generates IEEE floating-point numbers for compatibility with other development tools.

8.6.4 H8/300 Machine Directives

`as` has only one machine-dependent directive for the H8/300:

`.h8300h` Recognize and emit additional instructions for the H8/300H variant, and also make `.int` emit 32-bit numbers rather than the usual (16-bit) for the H8/300 family.

On the H8/300 family (including the H8/300H) ‘`.word`’ directives generate 16-bit numbers.

8.6.5 Opcodes

For detailed information on the H8/300 machine instruction set, see *H8/300 Series Programming Manual* (Hitachi ADE-602-025). For information specific to the H8/300H, see *H8/300H Series Programming Manual* (Hitachi).

`as` implements all the standard H8/300 opcodes. No additional pseudo-instructions are needed on this family.

Four H8/300 instructions (`add`, `cmp`, `mov`, `sub`) are defined with variants using the suffixes ‘`.b`’, ‘`.w`’, and ‘`.l`’ to specify the size of a memory operand. `as` supports these suffixes, but does not require them; since one of the operands is always a register, `as` can deduce the correct size.

For example, since `r0` refers to a 16-bit register,

```
mov    r0,@foo
```

is equivalent to

```
mov.w r0,@foo
```

If you use the size suffixes, `as` issues a warning when the suffix and the register size do not match.

8.7 H8/500 Dependent Features

8.7.1 Options

as has no additional command-line options for the Hitachi H8/500 family.

8.7.2 Syntax

8.7.2.1 Special Characters

'!' is the line comment character.

';' can be used instead of a newline to separate statements.

Since '\$' has no special meaning, you may use it in symbol names.

8.7.2.2 Register Names

You can use the predefined symbols 'r0', 'r1', 'r2', 'r3', 'r4', 'r5', 'r6', and 'r7' to refer to the H8/500 registers.

The H8/500 also has these control registers:

cp	code pointer
dp	data pointer
bp	base pointer
tp	stack top pointer
ep	extra pointer
sr	status register
ccr	condition code register

All registers are 16 bits long. To represent 32 bit numbers, use two adjacent registers; for distant memory addresses, use one of the segment pointers (cp for the program counter; dp for r0-r3; ep for r4 and r5; and tp for r6 and r7).

8.7.2.3 Addressing Modes

as understands the following addressing modes for the H8/500:

Rn	Register direct
@Rn	Register indirect
@(d:8, Rn)	Register indirect with 8 bit signed displacement
@(d:16, Rn)	Register indirect with 16 bit signed displacement
@-Rn	Register indirect with pre-decrement

<code>@Rn+</code>	Register indirect with post-increment
<code>@aa:8</code>	8 bit absolute address
<code>@aa:16</code>	16 bit absolute address
<code>#xx:8</code>	8 bit immediate
<code>#xx:16</code>	16 bit immediate

8.7.3 Floating Point

The H8/500 family has no hardware floating point, but the `.float` directive generates IEEE floating-point numbers for compatibility with other development tools.

8.7.4 H8/500 Machine Directives

`as` has no machine-dependent directives for the H8/500. However, on this platform the `.int` and `.word` directives generate 16-bit numbers.

8.7.5 Opcodes

For detailed information on the H8/500 machine instruction set, see *H8/500 Series Programming Manual* (Hitachi M21T001).

`as` implements all the standard H8/500 opcodes. No additional pseudo-instructions are needed on this family.

8.8 HPPA Dependent Features

8.8.1 Notes

As a back end for GNU CC `as` has been thoroughly tested and should work extremely well. We have tested it only minimally on hand written assembly code and no one has tested it much on the assembly output from the HP compilers.

The format of the debugging sections has changed since the original `as` port (version 1.3X) was released; therefore, you must rebuild all HPPA objects and libraries with the new assembler so that you can debug the final executable.

The HPPA `as` port generates a small subset of the relocations available in the SOM and ELF object file formats. Additional relocation support will be added as it becomes necessary.

8.8.2 Options

`as` has no machine-dependent command-line options for the HPPA.

8.8.3 Syntax

The assembler syntax closely follows the HPPA instruction set reference manual; assembler directives and general syntax closely follow the HPPA assembly language reference manual, with a few noteworthy differences.

First, a colon may immediately follow a label definition. This is simply for compatibility with how most assembly language programmers write code.

Some obscure expression parsing problems may affect hand written code which uses the `spop` instructions, or code which makes significant use of the `!` line separator.

`as` is much less forgiving about missing arguments and other similar oversights than the HP assembler. `as` notifies you of missing arguments as syntax errors; this is regarded as a feature, not a bug.

Finally, `as` allows you to use an external symbol without explicitly importing the symbol. *Warning:* in the future this will be an error for HPPA targets.

Special characters for HPPA targets include:

‘`;`’ is the line comment character.

‘`!`’ can be used instead of a newline to separate statements.

Since ‘`$`’ has no special meaning, you may use it in symbol names.

8.8.4 Floating Point

The HPPA family uses IEEE floating-point numbers.

8.8.5 HPPA Assembler Directives

`as` for the HPPA supports many additional directives for compatibility with the native assembler. This section describes them only briefly. For detailed information on HPPA-specific assembler directives, see *HP9000 Series 800 Assembly Language Reference Manual* (HP 92432-90001).

`as` does *not* support the following assembler directives described in the HP manual:

```
.endm          .liston
.enter         .locct
.leave        .macro
.listoff
```

Beyond those implemented for compatibility, `as` supports one additional assembler directive for the HPPA: `.param`. It conveys register argument locations for static functions. Its syntax closely follows the `.export` directive.

These are the additional directives in `as` for the HPPA:

```
.block n
.blockz n Reserve n bytes of storage, and initialize them to zero.
.call    Mark the beginning of a procedure call. Only the special case with no arguments
         is allowed.
.callinfo [ param=value, ... ] [ flag, ... ]
         Specify a number of parameters and flags that define the environment for a
         procedure.
         param may be any of ‘frame’ (frame size), ‘entry_gr’ (end of general register
         range), ‘entry_fr’ (end of float register range), ‘entry_sr’ (end of space
         register range).
         The values for flag are ‘calls’ or ‘caller’ (proc has subroutines), ‘no_calls’
         (proc does not call subroutines), ‘save_rp’ (preserve return pointer), ‘save_sp’
         (proc preserves stack pointer), ‘no_unwind’ (do not unwind this proc),
         ‘hpx_int’ (proc is interrupt routine).
.code    Assemble into the standard section called ‘$TEXT$’, subsection ‘$CODE$’.
.copyright "string"
         In the SOM object format, insert string into the object code, marked as a
         copyright string.
.copyright "string"
         In the ELF object format, insert string into the object code, marked as a version
         string.
.enter   Not yet supported; the assembler rejects programs containing this directive.
.entry   Mark the beginning of a procedure.
.exit    Mark the end of a procedure.
.export name [ ,typ ] [ ,param=r ]
         Make a procedure name available to callers. typ, if present, must be one
         of ‘absolute’, ‘code’ (ELF only, not SOM), ‘data’, ‘entry’, ‘data’, ‘entry’,
         ‘millicode’, ‘plabel’, ‘pri_prog’, or ‘sec_prog’.
```

param, if present, provides either relocation information for the procedure arguments and result, or a privilege level. *param* may be ‘*argwn*’ (where *n* ranges from 0 to 3, and indicates one of four one-word arguments); ‘*rtnval*’ (the procedure’s result); or ‘*priv_lev*’ (privilege level). For arguments or the result, *r* specifies how to relocate, and must be one of ‘*no*’ (not relocatable), ‘*gr*’ (argument is in general register), ‘*fr*’ (in floating point register), or ‘*fu*’ (upper half of float register). For ‘*priv_lev*’, *r* is an integer.

- `.half n` Define a two-byte integer constant *n*; synonym for the portable `as` directive `.short`.
- `.import name [,typ]`
Converse of `.export`; make a procedure available to call. The arguments use the same conventions as the first two arguments for `.export`.
- `.label name`
Define *name* as a label for the current assembly location.
- `.leave` Not yet supported; the assembler rejects programs containing this directive.
- `.origin lc`
Advance location counter to *lc*. Synonym for the [No value for ‘‘*as*’’] portable directive `.org`.
- `.param name [,typ] [,param=r]`
Similar to `.export`, but used for static procedures.
- `.proc` Use preceding the first statement of a procedure.
- `.procend` Use following the last statement of a procedure.
- `label .reg expr`
Synonym for `.equ`; define *label* with the absolute expression *expr* as its value.
- `.space secname [,params]`
Switch to section *secname*, creating a new section by that name if necessary. You may only use *params* when creating a new section, not when switching to an existing one. *secname* may identify a section by number rather than by name.

If specified, the list *params* declares attributes of the section, identified by keywords. The keywords recognized are ‘*spnum=exp*’ (identify this section by the number *exp*, an absolute expression), ‘*sort=exp*’ (order sections according to this sort key when linking; *exp* is an absolute expression), ‘*unloadable*’ (section contains no loadable data), ‘*notdefined*’ (this section defined elsewhere), and ‘*private*’ (data in this section not available to other programs).
- `.spnum secnam`
Allocate four bytes of storage, and initialize them with the section number of the section named *secnam*. (You can define the section number with the HPPA `.space` directive.)
- `.string "str"`
Copy the characters in the string *str* to the object file. See [Section 3.6.1.1 \[Strings\], page 17](#), for information on escape sequences you can use in `as` strings.

Warning! The HPPA version of `.string` differs from the usual as definition: it does *not* write a zero byte after copying `str`.

`.stringz "str"`

Like `.string`, but appends a zero byte after copying `str` to object file.

`.subspa name [,params]`

`.nsubspa name [,params]`

Similar to `.space`, but selects a subsection `name` within the current section. You may only specify `params` when you create a subsection (in the first instance of `.subspa` for this `name`).

If specified, the list `params` declares attributes of the subsection, identified by keywords. The keywords recognized are ‘`quad=expr`’ (“quadrant” for this subsection), ‘`align=expr`’ (alignment for beginning of this subsection; a power of two), ‘`access=expr`’ (value for “access rights” field), ‘`sort=expr`’ (sorting order for this subspace in link), ‘`code_only`’ (subsection contains only code), ‘`unloadable`’ (subsection cannot be loaded into memory), ‘`common`’ (subsection is common block), ‘`dup_comm`’ (initialized data may have duplicate names), or ‘`zero`’ (subsection is all zeros, do not write in object file).

`.nsubspa` always creates a new subspace with the given name, even if one with the same name already exists.

`.version "str"`

Write `str` as version identifier in object code.

8.8.6 Opcodes

For detailed information on the HPPA machine instruction set, see *PA-RISC Architecture and Instruction Set Reference Manual* (HP 09740-90039).

8.9 ESA/390 Dependent Features

8.9.1 Notes

The ESA/390 `as` port is currently intended to be a back-end for the GNU `CC` compiler. It is not HLASM compatible, although it does support a subset of some of the HLASM directives. The only supported binary file format is ELF; none of the usual MVS/VM/OE/USS object file formats, such as ESD or XSD, are supported.

When used with the GNU `CC` compiler, the ESA/390 `as` will produce correct, fully relocated, functional binaries, and has been used to compile and execute large projects. However, many aspects should still be considered experimental; these include shared library support, dynamically loadable objects, and any relocation other than the 31-bit relocation.

8.9.2 Options

`as` has no machine-dependent command-line options for the ESA/390.

8.9.3 Syntax

The opcode/operand syntax follows the ESA/390 Principles of Operation manual; assembler directives and general syntax are loosely based on the prevailing AT&T/SVR4/ELF/Solaris style notation. HLASM-style directives are *not* supported for the most part, with the exception of those described herein.

A leading dot in front of directives is optional, and the case of directives is ignored; thus for example, `.using` and `USING` have the same effect.

A colon may immediately follow a label definition. This is simply for compatibility with how most assembly language programmers write code.

`#` is the line comment character.

`;` can be used instead of a newline to separate statements.

Since `$` has no special meaning, you may use it in symbol names.

Registers can be given the symbolic names `r0..r15`, `fp0`, `fp2`, `fp4`, `fp6`. By using these symbolic names, `as` can detect simple syntax errors. The name `rarg` or `r.arg` is a synonym for `r11`, `rtca` or `r.tca` for `r12`, `sp`, `r.sp`, `dsa` `r.dsa` for `r13`, `lr` or `r.lr` for `r14`, `rbase` or `r.base` for `r3` and `rpgt` or `r.pgt` for `r4`.

`*` is the current location counter. Unlike `.` it is always relative to the last `USING` directive. Note that this means that expressions cannot use multiplication, as any occurrence of `*` will be interpreted as a location counter.

All labels are relative to the last `USING`. Thus, branches to a label always imply the use of `base+displacement`.

Many of the usual forms of address constants / address literals are supported. Thus,

```
.using *,r3
L r15,=A(some_routine)
LM r6,r7,=V(some_longlong_extern)
A r1,=F'12'
AH r0,=H'42'
```



```
ME r6,=E'3.1416'
MD r6,=D'3.14159265358979'
O r6,=XL4'cacad0d0'
.ltorg
```

should all behave as expected: that is, an entry in the literal pool will be created (or reused if it already exists), and the instruction operands will be the displacement into the literal pool using the current base register (as last declared with the `.using` directive).

8.9.4 Floating Point

The assembler generates only IEEE floating-point numbers. The older floating point formats are not supported.

8.9.5 ESA/390 Assembler Directives

as for the ESA/390 supports all of the standard ELF/SVR4 assembler directives that are documented in the main part of this documentation. Several additional directives are supported in order to implement the ESA/390 addressing model. The most important of these are `.using` and `.ltorg`

These are the additional directives in `as` for the ESA/390:

- `.dc` A small subset of the usual DC directive is supported.
- `.drop regno` Stop using *regno* as the base register. The *regno* must have been previously declared with a `.using` directive in the same section as the current section.
- `.ebcdic string` Emit the EBCDIC equivalent of the indicated string. The emitted string will be null terminated. Note that the directives `.string` etc. emit ascii strings by default.
- `EQU` The standard HLASM-style EQU directive is not supported; however, the standard `as` directive `.equ` can be used to the same effect.
- `.ltorg` Dump the literal pool accumulated so far; begin a new literal pool. The literal pool will be written in the current section; in order to generate correct assembly, a `.using` must have been previously specified in the same section.
- `.using expr, regno` Use *regno* as the base register for all subsequent RX, RS, and SS form instructions. The *expr* will be evaluated to obtain the base address; usually, *expr* will merely be `'*`.

This assembler allows two `.using` directives to be simultaneously outstanding, one in the `.text` section, and one in another section (typically, the `.data` section). This feature allows dynamically loaded objects to be implemented in a relatively straightforward way. A `.using` directive must always be specified in the `.text` section; this will specify the base register that will be used for branches in the `.text` section. A second `.using` may be specified in another section; this will specify the base register that is used for non-label address

literals. When a second `.using` is specified, then the subsequent `.ltorg` must be put in the same section; otherwise an error will result.

Thus, for example, the following code uses `r3` to address branch targets and `r4` to address the literal pool, which has been written to the `.data` section. The is, the constants `=A(some_routine)`, `=H'42'` and `=E'3.1416'` will all appear in the `.data` section.

```
.data
.using LITPOOL,r4
.text
BASR r3,0
.using *,r3
      B          START
.long LITPOOL
START:
L r4,4(,r3)
L r15,=A(some_routine)
LTR r15,r15
BNE LABEL
AH r0,=H'42'
LABEL:
ME r6,=E'3.1416'
.data
LITPOOL:
.ltorg
```

Note that this dual-`.using` directive semantics extends and is not compatible with HLASM semantics. Note that this assembler directive does not support the full range of HLASM semantics.

8.9.6 Opcodes

For detailed information on the ESA/390 machine instruction set, see *ESA/390 Principles of Operation* (IBM Publication Number DZ9AR004).

8.10 80386 Dependent Features

8.10.1 Options

The 80386 has no machine dependent options.

8.10.2 AT&T Syntax versus Intel Syntax

In order to maintain compatibility with the output of `gcc`, `as` supports AT&T System V/386 assembler syntax. This is quite different from Intel syntax. We mention these differences because almost all 80386 documents use Intel syntax. Notable differences between the two syntaxes are:

- AT&T immediate operands are preceded by '\$'; Intel immediate operands are undelimited (Intel 'push 4' is AT&T 'pushl \$4'). AT&T register operands are preceded by '%'; Intel register operands are undelimited. AT&T absolute (as opposed to PC relative) jump/call operands are prefixed by '*'; they are undelimited in Intel syntax.
- AT&T and Intel syntax use the opposite order for source and destination operands. Intel 'add eax, 4' is 'addl \$4, %eax'. The 'source, dest' convention is maintained for compatibility with previous Unix assemblers. Note that instructions with more than one source operand, such as the 'enter' instruction, do *not* have reversed order. [Section 8.10.11 \[i386-Bugs\], page 86](#).
- In AT&T syntax the size of memory operands is determined from the last character of the instruction mnemonic. Mnemonic suffixes of 'b', 'w', and 'l' specify byte (8-bit), word (16-bit), and long (32-bit) memory references. Intel syntax accomplishes this by prefixing memory operands (*not* the instruction mnemonics) with 'byte ptr', 'word ptr', and 'dword ptr'. Thus, Intel 'mov al, byte ptr foo' is 'movb foo, %al' in AT&T syntax.
- Immediate form long jumps and calls are 'lcall/ljmp \$section, \$offset' in AT&T syntax; the Intel syntax is 'call/jmp far section:offset'. Also, the far return instruction is 'lret \$stack-adjust' in AT&T syntax; Intel syntax is 'ret far stack-adjust'.
- The AT&T assembler does not provide support for multiple section programs. Unix style systems expect all programs to be single sections.

8.10.3 Instruction Naming

Instruction mnemonics are suffixed with one character modifiers which specify the size of operands. The letters 'b', 'w', and 'l' specify byte, word, and long operands. If no suffix is specified by an instruction then `as` tries to fill in the missing suffix based on the destination register operand (the last one by convention). Thus, 'mov %ax, %bx' is equivalent to 'movw %ax, %bx'; also, 'mov \$1, %bx' is equivalent to 'movw \$1, %bx'. Note that this is incompatible with the AT&T Unix assembler which assumes that a missing mnemonic suffix implies long operand size. (This incompatibility does not affect compiler output since compilers always explicitly specify the mnemonic suffix.)

Almost all instructions have the same names in AT&T and Intel format. There are a few exceptions. The sign extend and zero extend instructions need two sizes to specify them. They need a size to sign/zero extend *from* and a size to zero extend *to*. This is accomplished

by using two instruction mnemonic suffixes in AT&T syntax. Base names for sign extend and zero extend are ‘movs...’ and ‘movz...’ in AT&T syntax (‘movsx’ and ‘movzx’ in Intel syntax). The instruction mnemonic suffixes are tacked on to this base name, the *from* suffix before the *to* suffix. Thus, ‘movsbl %al, %edx’ is AT&T syntax for “move sign extend *from* %al *to* %edx.” Possible suffixes, thus, are ‘bl’ (from byte to long), ‘bw’ (from byte to word), and ‘wl’ (from word to long).

The Intel-syntax conversion instructions

- ‘cbw’ — sign-extend byte in ‘%al’ to word in ‘%ax’,
- ‘cwde’ — sign-extend word in ‘%ax’ to long in ‘%eax’,
- ‘cwd’ — sign-extend word in ‘%ax’ to long in ‘%dx:%ax’,
- ‘cdq’ — sign-extend dword in ‘%eax’ to quad in ‘%edx:%eax’,

are called ‘cbtw’, ‘cwtl’, ‘cwtl’, and ‘cltd’ in AT&T naming. `as` accepts either naming for these instructions.

Far call/jump instructions are ‘lcall’ and ‘ljmp’ in AT&T syntax, but are ‘call far’ and ‘jump far’ in Intel convention.

8.10.4 Register Naming

Register operands are always prefixed with ‘%’. The 80386 registers consist of

- the 8 32-bit registers ‘%eax’ (the accumulator), ‘%ebx’, ‘%ecx’, ‘%edx’, ‘%edi’, ‘%esi’, ‘%ebp’ (the frame pointer), and ‘%esp’ (the stack pointer).
- the 8 16-bit low-ends of these: ‘%ax’, ‘%bx’, ‘%cx’, ‘%dx’, ‘%di’, ‘%si’, ‘%bp’, and ‘%sp’.
- the 8 8-bit registers: ‘%ah’, ‘%al’, ‘%bh’, ‘%bl’, ‘%ch’, ‘%cl’, ‘%dh’, and ‘%dl’ (These are the high-bytes and low-bytes of ‘%ax’, ‘%bx’, ‘%cx’, and ‘%dx’)
- the 6 section registers ‘%cs’ (code section), ‘%ds’ (data section), ‘%ss’ (stack section), ‘%es’, ‘%fs’, and ‘%gs’.
- the 3 processor control registers ‘%cr0’, ‘%cr2’, and ‘%cr3’.
- the 6 debug registers ‘%db0’, ‘%db1’, ‘%db2’, ‘%db3’, ‘%db6’, and ‘%db7’.
- the 2 test registers ‘%tr6’ and ‘%tr7’.
- the 8 floating point register stack ‘%st’ or equivalently ‘%st(0)’, ‘%st(1)’, ‘%st(2)’, ‘%st(3)’, ‘%st(4)’, ‘%st(5)’, ‘%st(6)’, and ‘%st(7)’.

8.10.5 Instruction Prefixes

Instruction prefixes are used to modify the following instruction. They are used to repeat string instructions, to provide section overrides, to perform bus lock operations, and to change operand and address sizes. (Most instructions that normally operate on 32-bit operands will use 16-bit operands if the instruction has an “operand size” prefix.) Instruction prefixes are best written on the same line as the instruction they act upon. For example, the ‘scas’ (scan string) instruction is repeated with:

```
repne scas %es:(%edi),%al
```

You may also place prefixes on the lines immediately preceding the instruction, but this circumvents checks that `as` does with prefixes, and will not work with all prefixes.

Here is a list of instruction prefixes:

- Section override prefixes ‘cs’, ‘ds’, ‘ss’, ‘es’, ‘fs’, ‘gs’. These are automatically added by specifying using the *section:memory-operand* form for memory references.
- Operand/Address size prefixes ‘data16’ and ‘addr16’ change 32-bit operands/addresses into 16-bit operands/addresses, while ‘data32’ and ‘addr32’ change 16-bit ones (in a .code16 section) into 32-bit operands/addresses. These prefixes *must* appear on the same line of code as the instruction they modify. For example, in a 16-bit .code16 section, you might write:

```
addr32 jmpl *(%ebx)
```

- The bus lock prefix ‘lock’ inhibits interrupts during execution of the instruction it precedes. (This is only valid with certain instructions; see a 80386 manual for details).
- The wait for coprocessor prefix ‘wait’ waits for the coprocessor to complete the current instruction. This should never be needed for the 80386/80387 combination.
- The ‘rep’, ‘repe’, and ‘repne’ prefixes are added to string instructions to make them repeat ‘ecx’ times (‘cx’ times if the current address size is 16-bits).

8.10.6 Memory References

An Intel syntax indirect memory reference of the form

```
section:[base + index*scale + disp]
```

is translated into the AT&T syntax

```
section:disp(base, index, scale)
```

where *base* and *index* are the optional 32-bit base and index registers, *disp* is the optional displacement, and *scale*, taking the values 1, 2, 4, and 8, multiplies *index* to calculate the address of the operand. If no *scale* is specified, *scale* is taken to be 1. *section* specifies the optional section register for the memory operand, and may override the default section register (see a 80386 manual for section register defaults). Note that section overrides in AT&T syntax *must* be preceded by a ‘%’. If you specify a section override which coincides with the default section register, **as** does *not* output any section register override prefixes to assemble the given instruction. Thus, section overrides can be specified to emphasize which section register is used for a given memory operand.

Here are some examples of Intel and AT&T style memory references:

AT&T: ‘-4(%ebp)’, Intel: ‘[ebp - 4]’

base is ‘%ebp’; *disp* is ‘-4’. *section* is missing, and the default section is used (‘%ss’ for addressing with ‘%ebp’ as the base register). *index*, *scale* are both missing.

AT&T: ‘foo(,%eax,4)’, Intel: ‘[foo + eax*4]’

index is ‘%eax’ (scaled by a *scale* 4); *disp* is ‘foo’. All other fields are missing. The section register here defaults to ‘%ds’.

AT&T: ‘foo(,1)’; Intel ‘[foo]’

This uses the value pointed to by ‘foo’ as a memory operand. Note that *base* and *index* are both missing, but there is only *one* ‘,’. This is a syntactic exception.

AT&T: `%gs:foo`; Intel `gs:foo`

This selects the contents of the variable `foo` with section register *section* being `%gs`.

Absolute (as opposed to PC relative) call and jump operands must be prefixed with `*`. If no `*` is specified, `as` always chooses PC relative addressing for jump/call labels.

Any instruction that has a memory operand, but no register operand, *must* specify its size (byte, word, or long) with an instruction mnemonic suffix (`b`, `w`, or `l`, respectively).

8.10.7 Handling of Jump Instructions

Jump instructions are always optimized to use the smallest possible displacements. This is accomplished by using byte (8-bit) displacement jumps whenever the target is sufficiently close. If a byte displacement is insufficient a long (32-bit) displacement is used. We do not support word (16-bit) displacement jumps in 32-bit mode (i.e. prefixing the jump instruction with the `data16` instruction prefix), since the 80386 insists upon masking `%eip` to 16 bits after the word displacement is added.

Note that the `jcxz`, `jecxz`, `loop`, `loopz`, `loope`, `loopnz` and `loopne` instructions only come in byte displacements, so that if you use these instructions (`gcc` does not use them) you may get an error message (and incorrect code). The AT&T 80386 assembler tries to get around this problem by expanding `jcxz foo` to

```

        jcxz cx_zero
        jmp  cx_nonzero
cx_zero: jmp  foo
cx_nonzero:
```

8.10.8 Floating Point

All 80387 floating point types except packed BCD are supported. (BCD support may be added without much difficulty). These data types are 16-, 32-, and 64- bit integers, and single (32-bit), double (64-bit), and extended (80-bit) precision floating point. Each supported type has an instruction mnemonic suffix and a constructor associated with it. Instruction mnemonic suffixes specify the operand's data type. Constructors build these data types into memory.

- Floating point constructors are `.float` or `.single`, `.double`, and `.tfloat` for 32-, 64-, and 80-bit formats. These correspond to instruction mnemonic suffixes `s`, `l`, and `t`. `t` stands for 80-bit (ten byte) real. The 80387 only supports this format via the `fldt` (load 80-bit real to stack top) and `fstpt` (store 80-bit real and pop stack) instructions.
- Integer constructors are `.word`, `.long` or `.int`, and `.quad` for the 16-, 32-, and 64-bit integer formats. The corresponding instruction mnemonic suffixes are `s` (single), `l` (long), and `q` (quad). As with the 80-bit real format, the 64-bit `q` format is only present in the `fldq` (load quad integer to stack top) and `fistpq` (store quad integer and pop stack) instructions.

Register to register operations should not use instruction mnemonic suffixes. `fstl %st, %st(1)` will give a warning, and be assembled as if you wrote `fst %st, %st(1)`, since all

register to register operations use 80-bit floating point operands. (Contrast this with `'fstl %st, mem'`, which converts `'%st'` from 80-bit to 64-bit floating point format, then stores the result in the 4 byte location `'mem'`)

8.10.9 Intel's MMX and AMD's 3DNow! SIMD Operations

`as` supports Intel's MMX instruction set (SIMD instructions for integer data), available on Intel's Pentium MMX processors and Pentium II processors, AMD's K6 and K6-2 processors, Cyrix' M2 processor, and probably others. It also supports AMD's 3DNow! instruction set (SIMD instructions for 32-bit floating point data) available on AMD's K6-2 processor and possibly others in the future.

Currently, `as` does not support Intel's floating point SIMD, Katmai (KNI).

The eight 64-bit MMX operands, also used by 3DNow!, are called `'%mm0'`, `'%mm1'`, ... `'%mm7'`. They contain eight 8-bit integers, four 16-bit integers, two 32-bit integers, one 64-bit integer, or two 32-bit floating point values. The MMX registers cannot be used at the same time as the floating point stack.

See Intel and AMD documentation, keeping in mind that the operand order in instructions is reversed from the Intel syntax.

8.10.10 Writing 16-bit Code

While `as` normally writes only "pure" 32-bit i386 code, it also supports writing code to run in real mode or in 16-bit protected mode code segments. To do this, put a `' .code16'` or `' .code16gcc'` directive before the assembly language instructions to be run in 16-bit mode. You can switch `as` back to writing normal 32-bit code with the `' .code32'` directive.

`' .code16gcc'` provides experimental support for generating 16-bit code from `gcc`, and differs from `' .code16'` in that `'call'`, `'ret'`, `'enter'`, `'leave'`, `'push'`, `'pop'`, `'pusha'`, `'popa'`, `'pushf'`, and `'popf'` instructions default to 32-bit size. This is so that the stack pointer is manipulated in the same way over function calls, allowing access to function parameters at the same stack offsets as in 32-bit mode. `' .code16gcc'` also automatically adds address size prefixes where necessary to use the 32-bit addressing modes that `gcc` generates.

The code which `as` generates in 16-bit mode will not necessarily run on a 16-bit pre-80386 processor. To write code that runs on such a processor, you must refrain from using *any* 32-bit constructs which require `as` to output address or operand size prefixes.

Note that writing 16-bit code instructions by explicitly specifying a prefix or an instruction mnemonic suffix within a 32-bit code section generates different machine instructions than those generated for a 16-bit code segment. In a 32-bit code section, the following code generates the machine opcode bytes `'66 6a 04'`, which pushes the value `'4'` onto the stack, decrementing `'%esp'` by 2.

```
pushw $4
```

The same code in a 16-bit code section would generate the machine opcode bytes `'6a 04'` (ie. without the operand size prefix), which is correct since the processor default operand size is assumed to be 16 bits in a 16-bit code section.

8.10.11 AT&T Syntax bugs

The UnixWare assembler, and probably other AT&T derived ix86 Unix assemblers, generate floating point instructions with reversed source and destination registers in certain cases. Unfortunately, gcc and possibly many other programs use this reversed syntax, so we're stuck with it.

For example

```
fsub %st,%st(3)
```

results in '`%st(3)`' being updated to '`%st - %st(3)`' rather than the expected '`%st(3) - %st`'. This happens with all the non-commutative arithmetic floating point operations with two register operands where the source register is '`%st`' and the destination register is '`%st(i)`'.

8.10.12 Notes

There is some trickery concerning the '`mul`' and '`imul`' instructions that deserves mention. The 16-, 32-, and 64-bit expanding multiplies (base opcode '`0xf6`'; extension 4 for '`mul`' and 5 for '`imul`') can be output only in the one operand form. Thus, '`imul %ebx, %eax`' does *not* select the expanding multiply; the expanding multiply would clobber the '`%edx`' register, and this would confuse gcc output. Use '`imul %ebx`' to get the 64-bit product in '`%edx:%eax`'.

We have added a two operand form of '`imul`' when the first operand is an immediate mode expression and the second operand is a register. This is just a shorthand, so that, multiplying '`%eax`' by 69, for example, can be done with '`imul $69, %eax`' rather than '`imul $69, %eax, %eax`'.

8.11 Intel 80960 Dependent Features

8.11.1 i960 Command-line Options

`-ACA | -ACA_A | -ACB | -ACC | -AKA | -AKB | -AKC | -AMC`

Select the 80960 architecture. Instructions or features not supported by the selected architecture cause fatal errors.

'-ACA' is equivalent to '-ACA_A'; '-AKC' is equivalent to '-AMC'. Synonyms are provided for compatibility with other tools.

If you do not specify any of these options, `as` generates code for any instruction or feature that is supported by *some* version of the 960 (even if this means mixing architectures!). In principle, `as` attempts to deduce the minimal sufficient processor type if none is specified; depending on the object code format, the processor type may be recorded in the object file. If it is critical that the `as` output match a specific architecture, specify that architecture explicitly.

`-b`

Add code to collect information about conditional branches taken, for later optimization using branch prediction bits. (The conditional branch instructions have branch prediction bits in the CA, CB, and CC architectures.) If `BR` represents a conditional branch instruction, the following represents the code generated by the assembler when '-b' is specified:

```

                call    increment routine
                .word   0          # pre-counter
Label: BR
                call    increment routine
                .word   0          # post-counter

```

The counter following a branch records the number of times that branch was *not* taken; the difference between the two counters is the number of times the branch *was* taken.

A table of every such `Label` is also generated, so that the external postprocessor `gbr960` (supplied by Intel) can locate all the counters. This table is always labelled '`__BRANCH_TABLE__`'; this is a local symbol to permit collecting statistics for many separate object files. The table is word aligned, and begins with a two-word header. The first word, initialized to 0, is used in maintaining linked lists of branch tables. The second word is a count of the number of entries in the table, which follow immediately: each is a word, pointing to one of the labels illustrated above.

*NEXT	COUNT: <i>N</i>	*BRLAB 1	...	*BRLAB <i>N</i>
-------	-----------------	----------	-----	-----------------

__BRANCH_TABLE__ *layout*

The first word of the header is used to locate multiple branch tables, since each object file may contain one. Normally the links are maintained with a call to an initialization routine, placed at the beginning of each function in the file. The GNU C compiler generates these calls automatically when you give it a ‘-b’ option. For further details, see the documentation of ‘gbr960’.

-no-relax

Normally, Compare-and-Branch instructions with targets that require displacements greater than 13 bits (or that have external targets) are replaced with the corresponding compare (or ‘chkbit’) and branch instructions. You can use the ‘-no-relax’ option to specify that **as** should generate errors instead, if the target displacement is larger than 13 bits.

This option does not affect the Compare-and-Jump instructions; the code emitted for them is *always* adjusted when necessary (depending on displacement size), regardless of whether you use ‘-no-relax’.

8.11.2 Floating Point

as generates IEEE floating-point numbers for the directives ‘.float’, ‘.double’, ‘.extended’, and ‘.single’.

8.11.3 i960 Machine Directives

.bss *symbol, length, align*

Reserve *length* bytes in the bss section for a local *symbol*, aligned to the power of two specified by *align*. *length* and *align* must be positive absolute expressions. This directive differs from ‘.lcomm’ only in that it permits you to specify an alignment. See [Section 7.38 \[.lcomm\]](#), page 41.

.extended *flonums*

.extended expects zero or more flonums, separated by commas; for each flonum, ‘.extended’ emits an IEEE extended-format (80-bit) floating-point number.

.leafproc *call-lab, bal-lab*

You can use the ‘.leafproc’ directive in conjunction with the optimized **callj** instruction to enable faster calls of leaf procedures. If a procedure is known to call no other procedures, you may define an entry point that skips procedure prolog code (and that does not depend on system-supplied saved context), and declare it as the *bal-lab* using ‘.leafproc’. If the procedure also has an entry point that goes through the normal prolog, you can specify that entry point as *call-lab*.

A ‘.leafproc’ declaration is meant for use in conjunction with the optimized call instruction ‘callj’; the directive records the data needed later to choose between converting the ‘callj’ into a **bal** or a **call**.

call-lab is optional; if only one argument is present, or if the two arguments are identical, the single argument is assumed to be the `bal` entry point.

`.sysproc` *name*, *index*

The `.sysproc` directive defines a name for a system procedure. After you define it using `.sysproc`, you can use *name* to refer to the system procedure identified by *index* when calling procedures with the optimized call instruction `callj`.

Both arguments are required; *index* must be between 0 and 31 (inclusive).

8.11.4 i960 Opcodes

All Intel 960 machine instructions are supported; see [Section 8.11.1 \[i960 Command-line Options\]](#), page 87 for a discussion of selecting the instruction subset for a particular 960 architecture.

Some opcodes are processed beyond simply emitting a single corresponding instruction: `callj`, and Compare-and-Branch or Compare-and-Jump instructions with target displacements larger than 13 bits.

8.11.4.1 callj

You can write `callj` to have the assembler or the linker determine the most appropriate form of subroutine call: `call`, `bal`, or `calls`. If the assembly source contains enough information—a `.leafproc` or `.sysproc` directive defining the operand—then `as` translates the `callj`; if not, it simply emits the `callj`, leaving it for the linker to resolve.

8.11.4.2 Compare-and-Branch

The 960 architectures provide combined Compare-and-Branch instructions that permit you to store the branch target in the lower 13 bits of the instruction word itself. However, if you specify a branch target far enough away that its address won't fit in 13 bits, the assembler can either issue an error, or convert your Compare-and-Branch instruction into separate instructions to do the compare and the branch.

Whether `as` gives an error or expands the instruction depends on two choices you can make: whether you use the `-no-relax` option, and whether you use a “Compare and Branch” instruction or a “Compare and Jump” instruction. The “Jump” instructions are *always* expanded if necessary; the “Branch” instructions are expanded when necessary *unless* you specify `-no-relax`—in which case `as` gives an error instead.

These are the Compare-and-Branch instructions, their “Jump” variants, and the instruction pairs they may expand into:

<i>Compare and Branch</i>	<i>Jump</i>	<i>Expanded to</i>
<code>bbc</code>		<code>chkbit; bno</code>
<code>bbs</code>		<code>chkbit; bo</code>
<code>cmpibe</code>	<code>cmpije</code>	<code>cmpi; be</code>
<code>cmpibg</code>	<code>cmpijg</code>	<code>cmpi; bg</code>

cmpibge	cmpijge	cmpi; bge
cmpibl	cmpijl	cmpi; bl
cmpible	cmpijle	cmpi; ble
cmpibno	cmpijno	cmpi; bno
cmpibne	cmpijne	cmpi; bne
cmpibo	cmpijo	cmpi; bo
cmpobe	cmpoje	cmpo; be
cmpobg	cmpojg	cmpo; bg
cmpobge	cmpojge	cmpo; bge
cmpobl	cmpojl	cmpo; bl
cmpoble	cmpojle	cmpo; ble
cmpobne	cmpojne	cmpo; bne

8.12 M680x0 Dependent Features

8.12.1 M680x0 Options

The Motorola 680x0 version of `as` has a few machine dependent options.

You can use the `-1` option to shorten the size of references to undefined symbols. If you do not use the `-1` option, references to undefined symbols are wide enough for a full `long` (32 bits). (Since `as` cannot know where these symbols end up, `as` can only allocate space for the linker to fill in later. Since `as` does not know how far away these symbols are, it allocates as much space as it can.) If you use this option, the references are only one word wide (16 bits). This may be useful if you want the object file to be as small as possible, and you know that the relevant symbols are always less than 17 bits away.

For some configurations, especially those where the compiler normally does not prepend an underscore to the names of user variables, the assembler requires a `%` before any use of a register name. This is intended to let the assembler distinguish between C variables and functions named `a0` through `a7`, and so on. The `%` is always accepted, but is not required for certain configurations, notably `sun3`. The `--register-prefix-optional` option may be used to permit omitting the `%` even for configurations for which it is normally required. If this is done, it will generally be impossible to refer to C variables and functions with the same names as register names.

Normally the character `|` is treated as a comment character, which means that it can not be used in expressions. The `--bitwise-or` option turns `|` into a normal character. In this mode, you must either use C style comments, or start comments with a `#` character at the beginning of a line.

If you use an addressing mode with a base register without specifying the size, `as` will normally use the full 32 bit value. For example, the addressing mode `%a0@(%d0)` is equivalent to `%a0@(%d0:1)`. You may use the `--base-size-default-16` option to tell `as` to default to using the 16 bit value. In this case, `%a0@(%d0)` is equivalent to `%a0@(%d0:w)`. You may use the `--base-size-default-32` option to restore the default behaviour.

If you use an addressing mode with a displacement, and the value of the displacement is not known, `as` will normally assume that the value is 32 bits. For example, if the symbol `disp` has not been defined, `as` will assemble the addressing mode `%a0@(disp,%d0)` as though `disp` is a 32 bit value. You may use the `--disp-size-default-16` option to tell `as` to instead assume that the displacement is 16 bits. In this case, `as` will assemble `%a0@(disp,%d0)` as though `disp` is a 16 bit value. You may use the `--disp-size-default-32` option to restore the default behaviour.

`as` can assemble code for several different members of the Motorola 680x0 family. The default depends upon how `as` was configured when it was built; normally, the default is to assemble code for the 68020 microprocessor. The following options may be used to change the default. These options control which instructions and addressing modes are permitted. The members of the 680x0 family are very similar. For detailed information about the differences, see the Motorola manuals.

`'-m68000'`
`'-m68ec000'`
`'-m68hc000'`
`'-m68hc001'`
`'-m68008'`
`'-m68302'`
`'-m68306'`
`'-m68307'`
`'-m68322'`
`'-m68356'` Assemble for the 68000. `'-m68008'`, `'-m68302'`, and so on are synonyms for `'-m68000'`, since the chips are the same from the point of view of the assembler.

`'-m68010'` Assemble for the 68010.

`'-m68020'`
`'-m68ec020'`
 Assemble for the 68020. This is normally the default.

`'-m68030'`
`'-m68ec030'`
 Assemble for the 68030.

`'-m68040'`
`'-m68ec040'`
 Assemble for the 68040.

`'-m68060'`
`'-m68ec060'`
 Assemble for the 68060.

`'-mcpu32'`
`'-m68330'`
`'-m68331'`
`'-m68332'`
`'-m68333'`
`'-m68334'`
`'-m68336'`
`'-m68340'`
`'-m68341'`
`'-m68349'`
`'-m68360'` Assemble for the CPU32 family of chips.

`'-m5200'` Assemble for the ColdFire family of chips.

`'-m68881'`
`'-m68882'` Assemble 68881 floating point instructions. This is the default for the 68020, 68030, and the CPU32. The 68040 and 68060 always support floating point instructions.

`'-mno-68881'`
 Do not assemble 68881 floating point instructions. This is the default for 68000 and the 68010. The 68040 and 68060 always support floating point instructions, even if this option is used.

`-m68851` Assemble 68851 MMU instructions. This is the default for the 68020, 68030, and 68060. The 68040 accepts a somewhat different set of MMU instructions; `-m68851` and `-m68040` should not be used together.

`-mno-68851`
Do not assemble 68851 MMU instructions. This is the default for the 68000, 68010, and the CPU32. The 68040 accepts a somewhat different set of MMU instructions.

8.12.2 Syntax

This syntax for the Motorola 680x0 was developed at MIT.

The 680x0 version of `as` uses instructions names and syntax compatible with the Sun assembler. Intervening periods are ignored; for example, `mov1` is equivalent to `mov.1`.

In the following table *apc* stands for any of the address registers (`%a0` through `%a7`), the program counter (`%pc`), the zero-address relative to the program counter (`%zpc`), a suppressed address register (`%za0` through `%za7`), or it may be omitted entirely. The use of *size* means one of `w` or `l`, and it may be omitted, along with the leading colon, unless a scale is also specified. The use of *scale* means one of `1`, `2`, `4`, or `8`, and it may always be omitted along with the leading colon.

The following addressing modes are understood:

Immediate

`#number`

Data Register

`%d0` through `%d7`

Address Register

`%a0` through `%a7`

`%a7` is also known as `%sp`, i.e. the Stack Pointer. `%a6` is also known as `%fp`, the Frame Pointer.

Address Register Indirect

`%a0@` through `%a7@`

Address Register Postincrement

`%a0@+` through `%a7@+`

Address Register Predecrement

`%a0@-` through `%a7@-`

Indirect Plus Offset

`apc@(number)`

Index `apc@(number, register : size : scale)`

The *number* may be omitted.

Postindex `apc@(number)@(onumber, register : size : scale)`

The *onumber* or the *register*, but not both, may be omitted.

Preindex ‘*apc*@(*number*, *register*:*size*:*scale*)@(*onumber*)’

The *number* may be omitted. Omitting the *register* produces the Postindex addressing mode.

Absolute ‘*symbol*’, or ‘*digits*’, optionally followed by ‘:b’, ‘:w’, or ‘:l’.

8.12.3 Motorola Syntax

The standard Motorola syntax for this chip differs from the syntax already discussed (see [Section 8.12.2 \[Syntax\], page 93](#)). `as` can accept Motorola syntax for operands, even if MIT syntax is used for other operands in the same instruction. The two kinds of syntax are fully compatible.

In the following table *apc* stands for any of the address registers (‘%a0’ through ‘%a7’), the program counter (‘%pc’), the zero-address relative to the program counter (‘%zpc’), or a suppressed address register (‘%za0’ through ‘%za7’). The use of *size* means one of ‘w’ or ‘l’, and it may always be omitted along with the leading dot. The use of *scale* means one of ‘1’, ‘2’, ‘4’, or ‘8’, and it may always be omitted along with the leading asterisk.

The following additional addressing modes are understood:

Address Register Indirect

‘(%a0)’ through ‘(%a7)’

‘%a7’ is also known as ‘%sp’, i.e. the Stack Pointer. %a6 is also known as ‘%fp’, the Frame Pointer.

Address Register Postincrement

‘(%a0)+’ through ‘(%a7)+’

Address Register Predecrement

‘-(%a0)’ through ‘-(%a7)’

Indirect Plus Offset

‘*number*(%a0)’ through ‘*number*(%a7)’, or ‘*number*(%pc)’.

The *number* may also appear within the parentheses, as in ‘(*number*, %a0)’. When used with the *pc*, the *number* may be omitted (with an address register, omitting the *number* produces Address Register Indirect mode).

Index ‘*number*(*apc*, *register*.*size***scale*)’

The *number* may be omitted, or it may appear within the parentheses. The *apc* may be omitted. The *register* and the *apc* may appear in either order. If both *apc* and *register* are address registers, and the *size* and *scale* are omitted, then the first register is taken as the base register, and the second as the index register.

Postindex ‘([*number*, *apc*], *register*.*size***scale*, *onumber*)’

The *onumber*, or the *register*, or both, may be omitted. Either the *number* or the *apc* may be omitted, but not both.

Preindex ‘([*number*, *apc*, *register*.*size***scale*], *onumber*)’

The *number*, or the *apc*, or the *register*, or any two of them, may be omitted. The *onumber* may be omitted. The *register* and the *apc* may appear in either

order. If both *apc* and *register* are address registers, and the *size* and *scale* are omitted, then the first register is taken as the base register, and the second as the index register.

8.12.4 Floating Point

Packed decimal (P) format floating literals are not supported. Feel free to add the code! The floating point formats generated by directives are these.

```
.float      Single precision floating point constants.
.double    Double precision floating point constants.
.extend
.ldouble   Extended precision (long double) floating point constants.
```

8.12.5 680x0 Machine Directives

In order to be compatible with the Sun assembler the 680x0 assembler understands the following directives.

```
.data1     This directive is identical to a .data 1 directive.
.data2     This directive is identical to a .data 2 directive.
.even      This directive is a special case of the .align directive; it aligns the output to
           an even byte boundary.
.skip      This directive is identical to a .space directive.
```

8.12.6 Opcodes

8.12.6.1 Branch Improvement

Certain pseudo opcodes are permitted for branch instructions. They expand to the shortest branch instruction that reach the target. Generally these mnemonics are made by substituting ‘j’ for ‘b’ at the start of a Motorola mnemonic.

The following table summarizes the pseudo-operations. A * flags cases that are more fully described after the table:

Pseudo-Op	Displacement				
	BYTE	WORD	LONG	LONG	non-PC relative
jbsr	bsrs	bsr	bsrl	jsr	jsr
jra	bras	bra	bral	jmp	jmp
* jXX	bXXs	bXX	bXXl	bNXs;jmpl	bNXs;jmp
* dbXX	dbXX	dbXX		dbXX; bra;	jmp
* fjXX	fbXXw	fbXXw	fbXXl		fbNXw;jmp

XX: condition
 NX: negative of condition XX

*—see full description below

jbsr

bra These are the simplest jump pseudo-operations; they always map to one particular machine instruction, depending on the displacement to the branch target.

jXX Here, ‘jXX’ stands for an entire family of pseudo-operations, where XX is a conditional branch or condition-code test. The full list of pseudo-ops in this family is:

```

    jhi   jls   jcc   jcs   jne   jeq   jvc
    jvs   jpl   jmi   jge   jlt   jgt   jle
  
```

For the cases of non-PC relative displacements and long displacements on the 68000 or 68010, as issues a longer code fragment in terms of NX, the opposite condition to XX. For example, for the non-PC relative case:

```
jXX foo
```

gives

```

    bNXs oof
    jmp foo
oof:
  
```

dbXX The full family of pseudo-operations covered here is

```

    dbhi  dbls  dbcc  dbcs  dbne  dbeq  dbvc
    dbvs  dbpl  dbmi  dbge  dblt  dbgt  dble
    dbf   dbra  dbt
  
```

Other than for word and byte displacements, when the source reads ‘dbXX foo’, as emits

```

    dbXX oo1
    bra oo2
oo1: jmpl foo
oo2:
  
```

fjXX This family includes

```

    fjne  fjeq  fjge  fjlt  fjgt  fjle  fjf
    fjt   fjgl  fjgle  fjnge  fjngl  fjngle  fjngt
    fjnle fjnlt fjoge  fjogl  fjogt  fjole  fjolt
    fjor  fjseq fjsf  fjsne fjst  fjueq  fjuge
    fjugt fjule fjult  fjun
  
```

For branch targets that are not PC relative, as emits

```

    fbNX oof
    jmp foo
oof:
  
```

when it encounters ‘fjXX foo’.

8.12.6.2 Special Characters

The immediate character is '#' for Sun compatibility. The line-comment character is '|' (unless the '--bitwise-or' option is used). If a '#' appears at the beginning of a line, it is treated as a comment unless it looks like '# line file', in which case it is treated normally.

8.13 MIPS Dependent Features

GNU `as` for MIPS architectures supports several different MIPS processors, and MIPS ISA levels I through IV. For information about the MIPS instruction set, see *MIPS RISC Architecture*, by Kane and Heindrich (Prentice-Hall). For an overview of MIPS assembly conventions, see “Appendix D: Assembly Language Programming” in the same work.

8.13.1 Assembler options

The MIPS configurations of GNU `as` support these special options:

- `-G num` This option sets the largest size of an object that can be referenced implicitly with the `gp` register. It is only accepted for targets that use ECOFF format. The default value is 8.
- `-EB`
- `-EL` Any MIPS configuration of `as` can select big-endian or little-endian output at run time (unlike the other GNU development tools, which must be configured for one or the other). Use `-EB` to select big-endian output, and `-EL` for little-endian.
- `-mips1`
- `-mips2`
- `-mips3`
- `-mips4` Generate code for a particular MIPS Instruction Set Architecture level. `-mips1` corresponds to the R2000 and R3000 processors, `-mips2` to the R6000 processor, `-mips3` to the R4000 processor, and `-mips4` to the R8000 and R10000 processors. You can also switch instruction sets during the assembly; see [Section 8.13.4 \[Directives to override the ISA level\]](#), page 100.
- `-mips16`
- `-no-mips16` Generate code for the MIPS 16 processor. This is equivalent to putting `.set mips16` at the start of the assembly file. `-no-mips16` turns off this option.
- `-m4010`
- `-no-m4010` Generate code for the LSI R4010 chip. This tells the assembler to accept the R4010 specific instructions (`addciu`, `ffc`, etc.), and to not schedule `nop` instructions around accesses to the `HI` and `LO` registers. `-no-m4010` turns off this option.

`-m4650`

`-no-m4650`

Generate code for the MIPS R4650 chip. This tells the assembler to accept the ‘`mad`’ and ‘`madu`’ instruction, and to not schedule ‘`nop`’ instructions around accesses to the ‘`HI`’ and ‘`LO`’ registers. ‘`-no-m4650`’ turns off this option.

`-m3900`

`-no-m3900`

`-m4100`

`-no-m4100`

For each option ‘`-mnnnn`’, generate code for the MIPS *Rnnnn* chip. This tells the assembler to accept instructions specific to that chip, and to schedule for that chip’s hazards.

`-mcpu=cpu`

Generate code for a particular MIPS *cpu*. It is exactly equivalent to ‘`-mcpu`’, except that there are more value of *cpu* understood. Valid *cpu* value are:

2000, 3000, 3900, 4000, 4010, 4100, 4111, 4300, 4400, 4600, 4650,
5000, 6000, 8000, 10000

`-nocpp` This option is ignored. It is accepted for command-line compatibility with other assemblers, which use it to turn off C style preprocessing. With GNU `as`, there is no need for ‘`-nocpp`’, because the GNU assembler itself never runs the C preprocessor.

`--trap`

`--no-break`

`as` automatically macro expands certain division and multiplication instructions to check for overflow and division by zero. This option causes `as` to generate code to take a trap exception rather than a break exception when an error is detected. The trap instructions are only supported at Instruction Set Architecture level 2 and higher.

`--break`

`--no-trap`

Generate code to take a break exception rather than a trap exception when an error is detected. This is the default.

8.13.2 MIPS ECOFF object code

Assembling for a MIPS ECOFF target supports some additional sections besides the usual `.text`, `.data` and `.bss`. The additional sections are `.rdata`, used for read-only data, `.sdata`, used for small data, and `.sbss`, used for small common objects.

When assembling for ECOFF, the assembler uses the `$gp` (`$28`) register to form the address of a “small object”. Any object in the `.sdata` or `.sbss` sections is considered “small” in this sense. For external objects, or for objects in the `.bss` section, you can use the `gcc` ‘`-G`’ option to control the size of objects addressed via `$gp`; the default value is 8, meaning that a reference to any object eight bytes or smaller uses `$gp`. Passing ‘`-G 0`’ to `as` prevents it from using the `$gp` register on the basis of object size (but the assembler uses `$gp` for objects in `.sdata` or `sbss` in any case). The size of an object in the `.bss` section

is set by the `.comm` or `.lcomm` directive that defines it. The size of an external object may be set with the `.extern` directive. For example, `'extern sym,4'` declares that the object at `sym` is 4 bytes in length, while leaving `sym` otherwise undefined.

Using small ECOFF objects requires linker support, and assumes that the `$gp` register is correctly initialized (normally done automatically by the startup code). MIPS ECOFF assembly code must not modify the `$gp` register.

8.13.3 Directives for debugging information

MIPS ECOFF `as` supports several directives used for generating debugging information which are not supported by traditional MIPS assemblers. These are `.def`, `.undef`, `.dim`, `.file`, `.scl`, `.size`, `.tag`, `.type`, `.val`, `.stabd`, `.stabsn`, and `.stabs`. The debugging information generated by the three `.stab` directives can only be read by GDB, not by traditional MIPS debuggers (this enhancement is required to fully support C++ debugging). These directives are primarily used by compilers, not assembly language programmers!

8.13.4 Directives to override the ISA level

GNU `as` supports an additional directive to change the MIPS Instruction Set Architecture level on the fly: `.set mipsn`. `n` should be a number from 0 to 4. A value from 1 to 4 makes the assembler accept instructions for the corresponding ISA level, from that point on in the assembly. `.set mipsn` affects not only which instructions are permitted, but also how certain macros are expanded. `.set mips0` restores the ISA level to its original level: either the level you selected with command line options, or the default for your configuration. You can use this feature to permit specific R4000 instructions while assembling in 32 bit mode. Use this directive with care!

The directive `'set mips16'` puts the assembler into MIPS 16 mode, in which it will assemble instructions for the MIPS 16 processor. Use `'set nomips16'` to return to normal 32 bit mode.

Traditional MIPS assemblers do not support this directive.

8.13.5 Directives for extending MIPS 16 bit instructions

By default, MIPS 16 instructions are automatically extended to 32 bits when necessary. The directive `'set noautoextend'` will turn this off. When `'set noautoextend'` is in effect, any 32 bit instruction must be explicitly extended with the `'e'` modifier (e.g., `'li.e $4,1000'`). The directive `'set autoextend'` may be used to once again automatically extend instructions when necessary.

This directive is only meaningful when in MIPS 16 mode. Traditional MIPS assemblers do not support this directive.

8.13.6 Directive to mark data as an instruction

The `.insn` directive tells `as` that the following data is actually instructions. This makes a difference in MIPS 16 mode: when loading the address of a label which precedes instructions, `as` automatically adds 1 to the value, so that jumping to the loaded address will do the right thing.

8.13.7 Directives to save and restore options

The directives `.set push` and `.set pop` may be used to save and restore the current settings for all the options which are controlled by `.set`. The `.set push` directive saves the current settings on a stack. The `.set pop` directive pops the stack and restores the settings.

These directives can be useful inside an macro which must change an option such as the ISA level or instruction reordering but does not want to change the state of the code which invoked the macro.

Traditional MIPS assemblers do not support these directives.

8.14 picoJava Dependent Features

8.14.1 Options

as has two additional command-line options for the picoJava architecture.

-ml This option selects little endian data output.

-mb This option selects big endian data output.

8.15 Hitachi SH Dependent Features

8.15.1 Options

`as` has no additional command-line options for the Hitachi SH family.

8.15.2 Syntax

8.15.2.1 Special Characters

'!' is the line comment character.

You can use ';' instead of a newline to separate statements.

Since '\$' has no special meaning, you may use it in symbol names.

8.15.2.2 Register Names

You can use the predefined symbols 'r0', 'r1', 'r2', 'r3', 'r4', 'r5', 'r6', 'r7', 'r8', 'r9', 'r10', 'r11', 'r12', 'r13', 'r14', and 'r15' to refer to the SH registers.

The SH also has these control registers:

<code>pr</code>	procedure register (holds return address)
<code>pc</code>	program counter
<code>mach</code>	
<code>macl</code>	high and low multiply accumulator registers
<code>sr</code>	status register
<code>gbr</code>	global base register
<code>vbr</code>	vector base register (for interrupt vectors)

8.15.2.3 Addressing Modes

`as` understands the following addressing modes for the SH. *Rn* in the following refers to any of the numbered registers, but *not* the control registers.

<i>Rn</i>	Register direct
@ <i>Rn</i>	Register indirect
@- <i>Rn</i>	Register indirect with pre-decrement
@ <i>Rn</i> +	Register indirect with post-increment
@(<i>disp</i> , <i>Rn</i>)	Register indirect with displacement
@(R0, <i>Rn</i>)	Register indexed
@(<i>disp</i> , GBR)	GBR offset

`@(R0, GBR)`

GBR indexed

`addr`

`@(disp, PC)`

PC relative address (for branch or for addressing memory). The `as` implementation allows you to use the simpler form `addr` anywhere a PC relative address is called for; the alternate form is supported for compatibility with other assemblers.

`#imm`

Immediate data

8.15.3 Floating Point

The SH family has no hardware floating point, but the `.float` directive generates IEEE floating-point numbers for compatibility with other development tools.

8.15.4 SH Machine Directives

`uaword`

`ualong` `as` will issue a warning when a misaligned `.word` or `.long` directive is used. You may use `.uaword` or `.ualong` to indicate that the value is intentionally misaligned.

8.15.5 Opcodes

For detailed information on the SH machine instruction set, see *SH-Microcomputer User's Manual* (Hitachi Micro Systems, Inc.).

`as` implements all the standard SH opcodes. No additional pseudo-instructions are needed on this family. Note, however, that because `as` supports a simpler form of PC-relative addressing, you may simply write (for example)

```
mov.l  bar,r0
```

where other assemblers might require an explicit displacement to `bar` from the program counter:

```
mov.l  @(disp, PC)
```

8.16 SPARC Dependent Features

8.16.1 Options

The SPARC chip family includes several successive levels, using the same core instruction set, but including a few additional instructions at each level. There are exceptions to this however. For details on what instructions each variant supports, please see the chip's architecture reference manual.

By default, `as` assumes the core instruction set (SPARC v6), but “bumps” the architecture level as needed: it switches to successively higher architectures as it encounters instructions that only exist in the higher levels.

If not configured for SPARC v9 (`sparc64-***`) GAS will not bump passed `sparclite` by default, an option must be passed to enable the v9 instructions.

GAS treats `sparclite` as being compatible with v8, unless an architecture is explicitly requested. SPARC v9 is always incompatible with `sparclite`.

`-Av6` | `-Av7` | `-Av8` | `-Asparclet` | `-Asparclite`
`-Av8plus` | `-Av8plusa` | `-Av9` | `-Av9a`

Use one of the ‘-A’ options to select one of the SPARC architectures explicitly. If you select an architecture explicitly, `as` reports a fatal error if it encounters an instruction or feature requiring an incompatible or higher level.

‘-Av8plus’ and ‘-Av8plusa’ select a 32 bit environment.

‘-Av9’ and ‘-Av9a’ select a 64 bit environment and are not available unless GAS is explicitly configured with 64 bit environment support.

‘-Av8plusa’ and ‘-Av9a’ enable the SPARC V9 instruction set with Ultra-SPARC extensions.

`-xarch=v8plus` | `-xarch=v8plusa`

For compatibility with the Solaris v9 assembler. These options are equivalent to `-Av8plus` and `-Av8plusa`, respectively.

`-bump` Warn whenever it is necessary to switch to another level. If an architecture level is explicitly requested, GAS will not issue warnings until that level is reached, and will then bump the level as required (except between incompatible levels).

`-32` | `-64` Select the word size, either 32 bits or 64 bits. These options are only available with the ELF object file format, and require that the necessary BFD support has been included.

8.16.2 Enforcing aligned data

SPARC GAS normally permits data to be misaligned. For example, it permits the `.long` pseudo-op to be used on a byte boundary. However, the native SunOS and Solaris assemblers issue an error when they see misaligned data.

You can use the `--enforce-aligned-data` option to make SPARC GAS also issue an error about misaligned data, just as the SunOS and Solaris assemblers do.

The `--enforce-aligned-data` option is not the default because `gcc` issues misaligned data pseudo-ops when it initializes certain packed data structures (structures defined using the `packed` attribute). You may have to assemble with `GAS` in order to initialize packed data structures in your own code.

8.16.3 Floating Point

The Sparc uses IEEE floating-point numbers.

8.16.4 Sparc Machine Directives

The Sparc version of `as` supports the following additional machine directives:

- `.align` This must be followed by the desired alignment in bytes.
- `.common` This must be followed by a symbol name, a positive number, and `"bss"`. This behaves somewhat like `.comm`, but the syntax is different.
- `.half` This is functionally identical to `.short`.
- `.nword` On the Sparc, the `.nword` directive produces native word sized value, ie. if assembling with `-32` it is equivalent to `.word`, if assembling with `-64` it is equivalent to `.xword`.
- `.proc` This directive is ignored. Any text following it on the same line is also ignored.
- `.register` This directive declares use of a global application or system register. It must be followed by a register name `%g2`, `%g3`, `%g6` or `%g7`, comma and the symbol name for that register. If symbol name is `#scratch`, it is a scratch register, if it is `#ignore`, it just suppresses any errors about using undeclared global register, but does not emit any information about it into the object file. This can be useful e.g. if you save the register before use and restore it after.
- `.reserve` This must be followed by a symbol name, a positive number, and `"bss"`. This behaves somewhat like `.lcomm`, but the syntax is different.
- `.seg` This must be followed by `"text"`, `"data"`, or `"data1"`. It behaves like `.text`, `.data`, or `.data 1`.
- `.skip` This is functionally identical to the `.space` directive.
- `.word` On the Sparc, the `.word` directive produces 32 bit values, instead of the 16 bit values it produces on many other machines.
- `.xword` On the Sparc V9 processor, the `.xword` directive produces 64 bit values.

8.17 Z8000 Dependent Features

The Z8000 as supports both members of the Z8000 family: the unsegmented Z8002, with 16 bit addresses, and the segmented Z8001 with 24 bit addresses.

When the assembler is in unsegmented mode (specified with the `unsegm` directive), an address takes up one word (16 bit) sized register. When the assembler is in segmented mode (specified with the `segm` directive), a 24-bit address takes up a long (32 bit) register. See [Section 8.17.3 \[Assembler Directives for the Z8000\], page 108](#), for a list of other Z8000 specific assembler directives.

8.17.1 Options

`as` has no additional command-line options for the Zilog Z8000 family.

8.17.2 Syntax

8.17.2.1 Special Characters

'!' is the line comment character.

You can use ';' instead of a newline to separate statements.

8.17.2.2 Register Names

The Z8000 has sixteen 16 bit registers, numbered 0 to 15. You can refer to different sized groups of registers by register number, with the prefix 'r' for 16 bit registers, 'rr' for 32 bit registers and 'rq' for 64 bit registers. You can also refer to the contents of the first eight (of the sixteen 16 bit registers) by bytes. They are named 'rnh' and 'rnl'.

byte registers

```
r0l r0h r1l r1h r2l r2h r3l r3h
r4l r4h r5l r5h r6l r6h r7l r7h
```

word registers

```
r0 r1 r2 r3 r4 r5 r6 r7 r8 r9 r10 r11 r12 r13 r14 r15
```

long word registers

```
rr0 rr2 rr4 rr6 rr8 rr10 rr12 rr14
```

quad word registers

```
rq0 rq4 rq8 rq12
```

8.17.2.3 Addressing Modes

`as` understands the following addressing modes for the Z8000:

<code>rn</code>	Register direct
<code>@rn</code>	Indirect register
<code>addr</code>	Direct: the 16 bit or 24 bit address (depending on whether the assembler is in segmented or unsegmented mode) of the operand is in the instruction.

<code>address(rn)</code>	Indexed: the 16 or 24 bit address is added to the 16 bit register to produce the final address in memory of the operand.
<code>rn(#imm)</code>	Base Address: the 16 or 24 bit register is added to the 16 bit sign extended immediate displacement to produce the final address in memory of the operand.
<code>rn(rm)</code>	Base Index: the 16 or 24 bit register <code>rn</code> is added to the sign extended 16 bit index register <code>rm</code> to produce the final address in memory of the operand.
<code>#xx</code>	Immediate data <code>xx</code> .

8.17.3 Assembler Directives for the Z8000

The Z8000 port of `as` includes these additional assembler directives, for compatibility with other Z8000 assemblers. As shown, these do not begin with `'.'` (unlike the ordinary `as` directives).

<code>segm</code>	Generates code for the segmented Z8001.
<code>unsegm</code>	Generates code for the unsegmented Z8002.
<code>name</code>	Synonym for <code>.file</code>
<code>global</code>	Synonym for <code>.global</code>
<code>wval</code>	Synonym for <code>.word</code>
<code>lval</code>	Synonym for <code>.long</code>
<code>bval</code>	Synonym for <code>.byte</code>
<code>sval</code>	Assemble a string. <code>sval</code> expects one string literal, delimited by single quotes. It assembles each byte of the string into consecutive addresses. You can use the escape sequence <code>'%xx'</code> (where <code>xx</code> represents a two-digit hexadecimal number) to represent the character whose ASCII value is <code>xx</code> . Use this feature to describe single quote and other characters that may not appear in string literals as themselves. For example, the C statement <code>'char *a = "he said \"it's 50% off\"";'</code> is represented in Z8000 assembly language (shown with the assembler output in hex at the left) as <pre>68652073 sval 'he said %22it%27s 50%25 off%22%00' 61696420 22697427 73203530 25206F66 662200</pre>
<code>rsect</code>	synonym for <code>.section</code>
<code>block</code>	synonym for <code>.space</code>
<code>even</code>	special case of <code>.align</code> ; aligns output to even byte boundary.

8.17.4 Opcodes

For detailed information on the Z8000 machine instruction set, see *Z8000 Technical Manual*.

8.18 VAX Dependent Features

8.18.1 VAX Command-Line Options

The Vax version of `as` accepts any of the following options, gives a warning message that the option was ignored and proceeds. These options are for compatibility with scripts designed for other people's assemblers.

`-D` (Debug)

`-S` (Symbol Table)

`-T` (Token Trace)

These are obsolete options used to debug old assemblers.

`-d` (Displacement size for JUMPs)

This option expects a number following the `'-d'`. Like options that expect file names, the number may immediately follow the `'-d'` (old standard) or constitute the whole of the command line argument that follows `'-d'` (GNU standard).

`-V` (Virtualize Interpass Temporary File)

Some other assemblers use a temporary file. This option commanded them to keep the information in active memory rather than in a disk file. `as` always does this, so this option is redundant.

`-J` (JUMPify Longer Branches)

Many 32-bit computers permit a variety of branch instructions to do the same job. Some of these instructions are short (and fast) but have a limited range; others are long (and slow) but can branch anywhere in virtual memory. Often there are 3 flavors of branch: short, medium and long. Some other assemblers would emit short and medium branches, unless told by this option to emit short and long branches.

`-t` (Temporary File Directory)

Some other assemblers may use a temporary file, and this option takes a filename being the directory to site the temporary file. Since `as` does not use a temporary disk file, this option makes no difference. `'-t'` needs exactly one filename.

The Vax version of the assembler accepts additional options when compiled for VMS:

`'-h n'` External symbol or section (used for global variables) names are not case sensitive on VAX/VMS and always mapped to upper case. This is contrary to the C language definition which explicitly distinguishes upper and lower case. To implement a standard conforming C compiler, names must be changed (mapped) to preserve the case information. The default mapping is to convert all lower case characters to uppercase and adding an underscore followed by a 6 digit hex value, representing a 24 digit binary value. The one digits in the binary value represent which characters are uppercase in the original symbol name.

The `'-h n'` option determines how we map names. This takes several values. No `'-h'` switch at all allows case hacking as described above. A value of zero (`'-h0'`) implies names should be upper case, and inhibits the case hack. A value

of 2 (`-h2`) implies names should be all lower case, with no case hack. A value of 3 (`-h3`) implies that case should be preserved. The value 1 is unused. The `-H` option directs `as` to display every mapped symbol during assembly.

Symbols whose names include a dollar sign '\$' are exceptions to the general name mapping. These symbols are normally only used to reference VMS library names. Such symbols are always mapped to upper case.

- '-+' The '-+' option causes `as` to truncate any symbol name larger than 31 characters. The '-+' option also prevents some code following the `'_main'` symbol normally added to make the object file compatible with Vax-11 "C".
- '-1' This option is ignored for backward compatibility with `as` version 1.x.
- '-H' The '-H' option causes `as` to print every symbol which was changed by case mapping.

8.18.2 VAX Floating Point

Conversion of flonums to floating point is correct, and compatible with previous assemblers. Rounding is towards zero if the remainder is exactly half the least significant bit.

D, F, G and H floating point formats are understood.

Immediate floating literals (*e.g.* `'S'$6.9'`) are rendered correctly. Again, rounding is towards zero in the boundary case.

The `.float` directive produces `f` format numbers. The `.double` directive produces `d` format numbers.

8.18.3 Vax Machine Directives

The Vax version of the assembler supports four directives for generating Vax floating point constants. They are described in the table below.

- `.dfloat` This expects zero or more flonums, separated by commas, and assembles Vax `d` format 64-bit floating point constants.
- `.ffloat` This expects zero or more flonums, separated by commas, and assembles Vax `f` format 32-bit floating point constants.
- `.gfloat` This expects zero or more flonums, separated by commas, and assembles Vax `g` format 64-bit floating point constants.
- `.hfloat` This expects zero or more flonums, separated by commas, and assembles Vax `h` format 128-bit floating point constants.

8.18.4 VAX Opcodes

All DEC mnemonics are supported. Beware that `case...` instructions have exactly 3 operands. The dispatch table that follows the `case...` instruction should be made with `.word` statements. This is compatible with all unix assemblers we know of.

8.18.5 VAX Branch Improvement

Certain pseudo opcodes are permitted. They are for branch instructions. They expand to the shortest branch instruction that reaches the target. Generally these mnemonics are made by substituting ‘j’ for ‘b’ at the start of a DEC mnemonic. This feature is included both for compatibility and to help compilers. If you do not need this feature, avoid these opcodes. Here are the mnemonics, and the code they can expand into.

jbsb ‘**Jsb**’ is already an instruction mnemonic, so we chose ‘jbsb’.

 (byte displacement)
 bsbb ...

 (word displacement)
 bsbw ...

 (long displacement)
 jsb ...

jbr
jr Unconditional branch.

 (byte displacement)
 brb ...

 (word displacement)
 brw ...

 (long displacement)
 jmp ...

jCOND *COND* may be any one of the conditional branches **neq**, **nequ**, **eql**, **eqlu**, **gtr**, **geq**, **lss**, **gtru**, **lequ**, **vc**, **vs**, **gequ**, **cc**, **lssu**, **cs**. *COND* may also be one of the bit tests **bs**, **bc**, **bss**, **bcs**, **bsc**, **bcc**, **bssi**, **bcci**, **lbs**, **lbc**. *NOTCOND* is the opposite condition to *COND*.

 (byte displacement)
 bCOND ...

 (word displacement)
 bNOTCOND foo ; brw ... ; foo:

 (long displacement)
 bNOTCOND foo ; jmp ... ; foo:

jacbX *X* may be one of **b d f g h l w**.

 (word displacement)
 OPCODE ...

 (long displacement)
 OPCODE ..., foo ;
 brb bar ;
 foo: jmp ... ;
 bar:

jaobYYY YYY may be one of lss leq.

jsobZZZ ZZZ may be one of geq gtr.

(byte displacement)

OPCODE ...

(word displacement)

OPCODE ..., foo ;
brb bar ;
foo: brw *destination* ;
bar:

(long displacement)

OPCODE ..., foo ;
brb bar ;
foo: jmp *destination* ;
bar:

aobleq

aoblss

sobgeq

sobgtr

(byte displacement)

OPCODE ...

(word displacement)

OPCODE ..., foo ;
brb bar ;
foo: brw *destination* ;
bar:

(long displacement)

OPCODE ..., foo ;
brb bar ;
foo: jmp *destination* ;
bar:

8.18.6 VAX Operands

The immediate character is '\$' for Unix compatibility, not '#' as DEC writes it.

The indirect character is '*' for Unix compatibility, not '@' as DEC writes it.

The displacement sizing character is `` (an accent grave) for Unix compatibility, not '^' as DEC writes it. The letter preceding `` may have either case. 'G' is not understood, but all other letters (b i l s w) are understood.

Register names understood are r0 r1 r2 ... r15 ap fp sp pc. Upper and lower case letters are equivalent.

For instance

tstb *w'\$4(r5)

Any expression is permitted in an operand. Operands are comma separated.

8.18.7 Not Supported on VAX

Vax bit fields can not be assembled with `as`. Someone can add the required code if they really need it.

8.19 v850 Dependent Features

8.19.1 Options

`as` supports the following additional command-line options for the V850 processor family:

-wsigned_overflow

Causes warnings to be produced when signed immediate values overflow the space available for them within their opcodes. By default this option is disabled as it is possible to receive spurious warnings due to using exact bit patterns as immediate constants.

-wunsigned_overflow

Causes warnings to be produced when unsigned immediate values overflow the space available for them within their opcodes. By default this option is disabled as it is possible to receive spurious warnings due to using exact bit patterns as immediate constants.

-mv850

Specifies that the assembled code should be marked as being targeted at the V850 processor. This allows the linker to detect attempts to link such code with code assembled for other processors.

-mv850e

Specifies that the assembled code should be marked as being targeted at the V850E processor. This allows the linker to detect attempts to link such code with code assembled for other processors.

-mv850any

Specifies that the assembled code should be marked as being targeted at the V850 processor but support instructions that are specific to the extended variants of the process. This allows the production of binaries that contain target specific code, but which are also intended to be used in a generic fashion. For example `libgcc.a` contains generic routines used by the code produced by GCC for all versions of the v850 architecture, together with support routines only used by the V850E architecture.

8.19.2 Syntax

8.19.2.1 Special Characters

`#` is the line comment character.

8.19.2.2 Register Names

as supports the following names for registers:

general register 0

r0, zero

general register 1

r1

general register 2

r2, hp

general register 3

r3, sp

general register 4

r4, gp

general register 5

r5, tp

general register 6

r6

general register 7

r7

general register 8

r8

general register 9

r9

general register 10

r10

general register 11

r11

general register 12

r12

general register 13

r13

general register 14

r14

general register 15

r15

general register 16

r16

general register 17

r17

general register 18
r18

general register 19
r19

general register 20
r20

general register 21
r21

general register 22
r22

general register 23
r23

general register 24
r24

general register 25
r25

general register 26
r26

general register 27
r27

general register 28
r28

general register 29
r29

general register 30
r30, ep

general register 31
r31, lp

system register 0
eipc

system register 1
eipsw

system register 2
fepc

system register 3
fepsw

system register 4
ecr

```

system register 5
    psw
system register 16
    ctpc
system register 17
    ctpsw
system register 18
    dbpc
system register 19
    dbpsw
system register 20
    ctbp

```

8.19.3 Floating Point

The V850 family uses IEEE floating-point numbers.

8.19.4 V850 Machine Directives

```

.offset <expression>
    Moves the offset into the current section to the specified amount.

.section "name", <type>
    This is an extension to the standard .section directive. It sets the current section
    to be <type> and creates an alias for this section called "name".

.v850
    Specifies that the assembled code should be marked as being targeted at the
    V850 processor. This allows the linker to detect attempts to link such code
    with code assembled for other processors.

.v850e
    Specifies that the assembled code should be marked as being targeted at the
    V850E processor. This allows the linker to detect attempts to link such code
    with code assembled for other processors.

```

8.19.5 Opcodes

as implements all the standard V850 opcodes.

as also implements the following pseudo ops:

```

hi0()
    Computes the higher 16 bits of the given expression and stores it into the
    immediate operand field of the given instruction. For example:
    'mulhi hi0(here - there), r5, r6'
    computes the difference between the address of labels 'here' and 'there', takes
    the upper 16 bits of this difference, shifts it down 16 bits and then multiplies
    it by the lower 16 bits in register 5, putting the result into register 6.

```

- lo()** Computes the lower 16 bits of the given expression and stores it into the immediate operand field of the given instruction. For example:
`'addi lo(here - there), r5, r6'`
 computes the difference between the address of labels 'here' and 'there', takes the lower 16 bits of this difference and adds it to register 5, putting the result into register 6.
- hi()** Computes the higher 16 bits of the given expression and then adds the value of the most significant bit of the lower 16 bits of the expression and stores the result into the immediate operand field of the given instruction. For example the following code can be used to compute the address of the label 'here' and store it into register 6:
`'movhi hi(here), r0, r6' 'movea lo(here), r6, r6'`
 The reason for this special behaviour is that `movea` performs a sign extension on its immediate operand. So for example if the address of 'here' was `0xFFFFFFFF` then without the special behaviour of the `hi()` pseudo-op the `movhi` instruction would put `0xFFFF0000` into `r6`, then the `movea` instruction would take its immediate operand, `0xFFFF`, sign extend it to 32 bits, `0xFFFFFFFF`, and then add it into `r6` giving `0xFFFEFFFF` which is wrong (the fifth nibble is E). With the `hi()` pseudo op adding in the top bit of the `lo()` pseudo op, the `movhi` instruction actually stores 0 into `r6` (`0xFFFF + 1 = 0x0000`), so that the `movea` instruction stores `0xFFFFFFFF` into `r6` - the right value.
- hilo()** Computes the 32 bit value of the given expression and stores it into the immediate operand field of the given instruction (which must be a `mov` instruction). For example:
`'mov hilo(here), r6'`
 computes the absolute address of label 'here' and puts the result into register 6.
- sdaoff()** Computes the offset of the named variable from the start of the Small Data Area (whose address is held in register 4, the GP register) and stores the result as a 16 bit signed value in the immediate operand field of the given instruction. For example:
`'ld.w sdaoff(_a_variable)[gp], r6'`
 loads the contents of the location pointed to by the label '_a_variable' into register 6, provided that the label is located somewhere within +/- 32K of the address held in the GP register. [Note the linker assumes that the GP register contains a fixed address set to the address of the label called '_gp'. This can either be set up automatically by the linker, or specifically set by using the '`--defsym __gp=<value>`' command line option].
- tdaoff()** Computes the offset of the named variable from the start of the Tiny Data Area (whose address is held in register 30, the EP register) and stores the result as a 4,5, 7 or 8 bit unsigned value in the immediate operand field of the given instruction. For example:

```
'sld.w tdaoff(_a_variable)[ep],r6'
```

loads the contents of the location pointed to by the label '_a_variable' into register 6, provided that the label is located somewhere within +256 bytes of the address held in the EP register. [Note the linker assumes that the EP register contains a fixed address set to the address of the label called '_ep'. This can either be set up automatically by the linker, or specifically set by using the '--defsym __ep=<value>' command line option].

zdaoff() Computes the offset of the named variable from address 0 and stores the result as a 16 bit signed value in the immediate operand field of the given instruction. For example:

```
'movea zdaoff(_a_variable),zero,r6'
```

puts the address of the label '_a_variable' into register 6, assuming that the label is somewhere within the first 32K of memory. (Strictly speaking it is also possible to access the last 32K of memory as well, as the offsets are signed).

ctoff() Computes the offset of the named variable from the start of the Call Table Area (whose address is held in system register 20, the CTBP register) and stores the result as a 6 or 16 bit unsigned value in the immediate field of the given instruction or piece of data. For example:

```
'callt ctoff(table_func1)'
```

will put the call table function whose address is held in the call table at the location labeled 'table_func1'.

For information on the V850 instruction set, see *V850 Family 32-/16-Bit single-Chip Microcontroller Architecture Manual* from NEC. Ltd.

9 Reporting Bugs

Your bug reports play an essential role in making `as` 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 `as` work better. Bug reports are your contribution to the maintenance of `as`.

In order for a bug report to serve its purpose, you must include the information that enables us to fix the bug.

9.1 Have you found a bug?

If you are not sure whether you have found a bug, here are some guidelines:

- If the assembler gets a fatal signal, for any input whatever, that is a `as` bug. Reliable assemblers never crash.
- If `as` produces an error message for valid input, that is a bug.
- If `as` 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 assemblers, your suggestions for improvement of `as` are welcome in any case.

9.2 How to report bugs

A number of companies and individuals offer support for GNU products. If you obtained `as` 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 `as` to `bug-gnu-utils@gnu.org`.

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 a symbol 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 assembler 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 `as`. `as` announces it if you start it with the ‘`--version`’ argument. Without this, we will not know whether there is any point in looking for the bug in the current version of `as`.
- Any patches you may have applied to the `as` source.
- The type of machine you are using, and the operating system name and version number.
- What compiler (and its version) was used to compile `as`—e.g. “`gcc-2.7`”.
- The command arguments you gave the assembler to assemble your example and observe the bug. 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 file that will reproduce the bug. If the bug is observed when the assembler is invoked via a compiler, send the assembler source, not the high level language source. Most compilers will produce the assembler source when run with the ‘`-S`’ option. If you are using `gcc`, use the options ‘`-v --save-temps`’; this will save the assembler source in a file with an extension of ‘`.s`’, and also show you exactly how `as` is being run.
- 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 `as` 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 `as` 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 `as` source, send us context diffs, as generated by `diff` with the ‘`-u`’, ‘`-c`’, or ‘`-p`’ option. Always send diffs from the old file to the new file. If you even discuss something in the `as` 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 as 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.

10 Acknowledgements

If you have contributed to `as` and your name isn't listed here, it is not meant as a slight. We just don't know about it. Send mail to the maintainer, and we'll correct the situation. Currently the maintainer is Ken Raeburn (email address `raeburn@cygnus.com`).

Dean Elsner wrote the original GNU assembler for the VAX.¹

Jay Fenlason maintained GAS for a while, adding support for GDB-specific debug information and the 68k series machines, most of the preprocessing pass, and extensive changes in `messages.c`, `input-file.c`, `write.c`.

K. Richard Pixley maintained GAS for a while, adding various enhancements and many bug fixes, including merging support for several processors, breaking GAS up to handle multiple object file format back ends (including heavy rewrite, testing, an integration of the `coff` and `b.out` back ends), adding configuration including heavy testing and verification of cross assemblers and file splits and renaming, converted GAS to strictly ANSI C including full prototypes, added support for `m680[34]0` and `cpu32`, did considerable work on `i960` including a COFF port (including considerable amounts of reverse engineering), a SPARC opcode file rewrite, DECstation, `rs6000`, and `hp300hpux` host ports, updated “know” assertions and made them work, much other reorganization, cleanup, and lint.

Ken Raeburn wrote the high-level BFD interface code to replace most of the code in format-specific I/O modules.

The original VMS support was contributed by David L. Kashtan. Eric Youngdale has done much work with it since.

The Intel 80386 machine description was written by Eliot Dresselhaus.

Minh Tran-Le at IntelliCorp contributed some AIX 386 support.

The Motorola 88k machine description was contributed by Devon Bowen of Buffalo University and Torbjorn Granlund of the Swedish Institute of Computer Science.

Keith Knowles at the Open Software Foundation wrote the original MIPS back end (`tc-mips.c`, `tc-mips.h`), and contributed Rose format support (which hasn't been merged in yet). Ralph Campbell worked with the MIPS code to support `a.out` format.

Support for the Zilog Z8k and Hitachi H8/300 and H8/500 processors (`tc-z8k`, `tc-h8300`, `tc-h8500`), and IEEE 695 object file format (`obj-ieee`), was written by Steve Chamberlain of Cygnus Support. Steve also modified the COFF back end to use BFD for some low-level operations, for use with the H8/300 and AMD 29k targets.

John Gilmore built the AMD 29000 support, added `.include` support, and simplified the configuration of which versions accept which directives. He updated the 68k machine description so that Motorola's opcodes always produced fixed-size instructions (e.g. `jsr`), while synthetic instructions remained shrinkable (`jbsr`). John fixed many bugs, including true tested cross-compilation support, and one bug in relaxation that took a week and required the proverbial one-bit fix.

Ian Lance Taylor of Cygnus Support merged the Motorola and MIT syntax for the 68k, completed support for some COFF targets (68k, `i386 SVR3`, and `SCO Unix`), added support for MIPS ECOFF and ELF targets, wrote the initial `RS/6000` and `PowerPC` assembler, and made a few other minor patches.

¹ Any more details?

Steve Chamberlain made `as` able to generate listings.

Hewlett-Packard contributed support for the HP9000/300.

Jeff Law wrote GAS and BFD support for the native HPPA object format (SOM) along with a fairly extensive HPPA testsuite (for both SOM and ELF object formats). This work was supported by both the Center for Software Science at the University of Utah and Cygnus Support.

Support for ELF format files has been worked on by Mark Eichen of Cygnus Support (original, incomplete implementation for SPARC), Pete Hoogenboom and Jeff Law at the University of Utah (HPPA mainly), Michael Meissner of the Open Software Foundation (i386 mainly), and Ken Raeburn of Cygnus Support (sparc, and some initial 64-bit support).

Linus Vepstas added GAS support for the ESA/390 "Instruction 370" architecture.

Richard Henderson rewrote the Alpha assembler. Klaus Kaempf wrote GAS and BFD support for openVMS/Alpha.

Several engineers at Cygnus Support have also provided many small bug fixes and configuration enhancements.

Many others have contributed large or small bugfixes and enhancements. If you have contributed significant work and are not mentioned on this list, and want to be, let us know. Some of the history has been lost; we are not intentionally leaving anyone out.

Index

#

#	16
#APP	15
#NO_APP	15

\$

\$ in symbol names	62, 66, 72, 103
--------------------	-----------------

-

--	7
'--base-size-default-16'	91
'--base-size-default-32'	91
'--bitwise-or' option, M680x0	91
'--disp-size-default-16'	91
'--disp-size-default-32'	91
--enforce-aligned-data	105
'--fatal-warnings'	13
--MD	12
'--no-warn'	13
'--register-prefix-optional' option, M680x0	91
--statistics	13
--traditional-format	13
'--warn'	13
'+' option, VAX/VMS	110
'-1' option, VAX/VMS	110
-a	9
-A options, i960	87
-ac	9
-ad	9
-ah	9
-al	9
-an	9
-as	9
-Asparclet	105
-Asparclite	105
-Av6	105
-Av8	105
-Av9	105
-Av9a	105
-b option, i960	87
-D	9
-D, ignored on VAX	109
-d, VAX option	109
-EB command line option, ARM	57
-EB option (MIPS)	98
-EL command line option, ARM	57
-EL option (MIPS)	98
-f	9
-G option (MIPS)	98
'-h' option, VAX/VMS	109
'-H' option, VAX/VMS	110
-I <i>path</i>	10
-J, ignored on VAX	109
-K	10
-k command line option, ARM	57
-L	10
'-l' option, M680x0	91
-M	10
'-m68000' and related options	91
-mall command line option, ARM	57
-mapcs command line option, ARM	57
-marm command line option, ARM	57
-marmv command line option, ARM	57
-mbig-endian option (ARC)	54
-mfpa command line option, ARM	57
-mfpe-old command line option, ARM	57
-mlittle-endian option (ARC)	54
-mno-fpu command line option, ARM	57
-mthumb command line option, ARM	57
-mthumb-interwork command line option, ARM	57
-mv850 command line option, V850	113
-mv850any command line option, V850	113
-mv850e command line option, V850	113
-no-relax option, i960	88
-nocpp ignored (MIPS)	99
-o	12
-R	12
-S, ignored on VAX	109
-t, ignored on VAX	109
-T, ignored on VAX	109
-v	13
-V, redundant on VAX	109
-version	13
'-W'	13
-wsigned_overflow command line option, V850	113
-wunsigned_overflow command line option, V850	113
.	
. (symbol)	28
.insn	100
.ltorg directive, ARM	59
.o	8
.param on HPPA	75
.pool directive, ARM	59
.set autoextend	100
.set mipsn	100
.set noautoextend	100

<code>.set pop</code>	101	addressing modes, M680x0	93
<code>.set push</code>	101	addressing modes, SH	103
<code>.v850</code> directive, V850	116	addressing modes, Z8000	107
<code>.v850e</code> directive, V850	116	<code>ADR reg,<label></code> pseudo op, ARM	59
:		<code>ADRL reg,<label></code> pseudo op, ARM	59
: (label)	17	advancing location counter	44
@		<code>align</code> directive	33
@word modifier, D10V	64	<code>align</code> directive, SPARC	106
\		altered difference tables	52
\" (doublequote character)	18	alternate syntax for the 680x0	94
\\ ('\ character)	18	AMD 29K floating point (IEEE)	55
\b (backspace character)	17	AMD 29K identifiers	55
\ddd (octal character code)	18	AMD 29K line comment character	55
\f (formfeed character)	18	AMD 29K machine directives	56
\n (newline character)	18	AMD 29K macros	55
\r (carriage return character)	18	AMD 29K opcodes	56
\t (tab)	18	AMD 29K options (none)	55
\xd... (hex character code)	18	AMD 29K protected registers	55
1		AMD 29K register names	55
16-bit code, i386	85	AMD 29K special purpose registers	55
2		AMD 29K support	55
29K support	55	ARC architectures	54
3		ARC big-endian output	54
3DNow!, i386	85	ARC endianness	3
A		ARC floating point (IEEE)	54
a.out	8	ARC little-endian output	54
a.out symbol attributes	29	ARC machine directives	54
abort directive	33	ARC options	54
ABORT directive	33	ARC support	54
absolute section	22	architecture options, i960	87
addition, permitted arguments	32	architecture options, M680x0	91
addresses	31	architectures, ARC	54
addresses, format of	22	architectures, SPARC	105
addressing modes, D10V	63	arguments for addition	32
addressing modes, D30V	68	arguments for subtraction	32
addressing modes, H8/300	69	arguments in expressions	31
addressing modes, H8/500	72	arithmetic functions	31
		arithmetic operands	31
		<code>arm</code> directive, ARM	58
		ARM floating point (IEEE)	58
		ARM identifiers	58
		ARM immediate character	58
		ARM line comment character	58
		ARM line separator	58
		ARM machine directives	58
		ARM opcodes	59
		ARM options (none)	57
		ARM register names	58
		ARM support	57
		<code>ascii</code> directive	34
		<code>asciz</code> directive	34
		assembler bugs, reporting	119

- assembler crash 119
 - assembler internal logic error 23
 - assembler version 13
 - assembler, and linker 21
 - assembly listings, enabling 9
 - assigning values to symbols 27, 36
 - attributes, symbol 28
 - auxiliary attributes, COFF symbols 29
 - auxiliary symbol information, COFF 35
 - Av7 105
- B**
- backslash (\) 18
 - backspace (\b) 17
 - balgn** directive 34
 - balgnl** directive 34
 - balgnw** directive 34
 - big endian output, ARC 3
 - big endian output, MIPS 5
 - big endian output, PJ 4
 - big-endian output, ARC 54
 - big-endian output, MIPS 98
 - bignums 19
 - binary integers 19
 - bitfields, not supported on VAX 113
 - block** 108
 - block** directive, AMD 29K 56
 - branch improvement, M680x0 95
 - branch improvement, VAX 111
 - branch recording, i960 87
 - branch statistics table, i960 87
 - bss** directive, i960 88
 - bss section 22, 24
 - bug criteria 119
 - bug reports 119
 - bugs in assembler 119
 - bus lock prefixes, i386 83
 - bval** 108
 - byte** directive 34
- C**
- call instructions, i386 82
 - callj**, i960 pseudo-opcode 89
 - carriage return (\r) 18
 - character constants 17
 - character escape codes 17
 - character, single 18
 - characters used in symbols 16
 - code** directive, ARM 58
 - code16** directive, i386 85
 - code16gcc** directive, i386 85
 - code32** directive, i386 85
 - COFF auxiliary symbol information 35
 - COFF structure debugging 50
 - COFF symbol attributes 29
 - COFF symbol descriptor 35
 - COFF symbol storage class 46
 - COFF symbol type 51
 - COFF symbols, debugging 35
 - COFF value attribute 51
 - COMDAT 41
 - comm** directive 35
 - command line conventions 7
 - command line options, V850 113
 - command-line options ignored, VAX 109
 - comments 15
 - comments, M680x0 97
 - comments, removed by preprocessor 15
 - common** directive, SPARC 106
 - common sections 41
 - common variable storage 24
 - compare and jump expansions, i960 89
 - compare/branch instructions, i960 89
 - conditional assembly 39
 - constant, single character 18
 - constants 17
 - constants, bignum 19
 - constants, character 17
 - constants, converted by preprocessor 15
 - constants, floating point 19
 - constants, integer 19
 - constants, number 18
 - constants, string 17
 - conversion instructions, i386 82
 - coprocessor wait, i386 83
 - cpu** directive, SPARC 54
 - cputype** directive, AMD 29K 56
 - crash of assembler 119
 - ctbp** register, V850 116
 - ctoff** pseudo-op, V850 118
 - ctpc** register, V850 116
 - ctpsw** register, V850 116
 - current address 28
 - current address, advancing 44

D

D10V @word modifier	64
D10V addressing modes	63
D10V floating point	64
D10V line comment character	62
D10V opcode summary	64
D10V optimization	3
D10V options	61
D10V registers	62
D10V size modifiers	61
D10V sub-instruction ordering	62
D10V sub-instructions	61
D10V support	61
D10V syntax	61
D30V addressing modes	68
D30V floating point	68
D30V Guarded Execution	67
D30V line comment character	66
D30V nops	3
D30V nops after 32-bit multiply	4
D30V opcode summary	68
D30V optimization	3
D30V options	65
D30V registers	67
D30V size modifiers	65
D30V sub-instruction ordering	66
D30V sub-instructions	65
D30V support	65
D30V syntax	65
data alignment on SPARC	105
data and text sections, joining	12
data directive	35
data section	22
data1 directive, M680x0	95
data2 directive, M680x0	95
dbpc register, V850	116
dbpsw register, V850	116
debuggers, and symbol order	27
debugging COFF symbols	35
decimal integers	19
def directive	35
dependency tracking	12
deprecated directives	52
desc directive	35
descriptor, of a.out symbol	29
dfloat directive, VAX	110
difference tables altered	52
difference tables, warning	10
dim directive	35
directives and instructions	17
directives, M680x0	95
directives, machine independent	33

directives, Z8000	108
displacement sizing character, VAX	112
dot (symbol)	28
double directive	36
double directive, i386	84
double directive, M680x0	95
double directive, VAX	110
doublequote (\")	18

E

ECOFF sections	99
ecr register, V850	115
eight-byte integer	45
eipc register, V850	115
eipsw register, V850	115
eject directive	36
else directive	36
empty expressions	31
emulation	5
end directive	36
endif directive	36
endfunc directive	36
endianness, ARC	3
endianness, MIPS	5
endianness, PJ	4
endif directive	36
endm directive	43
EOF, newline must precede	16
ep register, V850	115
equ directive	36
equiv directive	37
err directive	37
error messages	8
error on valid input	119
errors, caused by warnings	13
errors, continuing after	13
ESA/390 floating point (IEEE)	79
ESA/390 support	78
ESA/390 Syntax	78
ESA/390-only directives	79
escape codes, character	17
even	108
even directive, M680x0	95
exitm directive	43
expr (internal section)	23
expression arguments	31
expressions	31
expressions, empty	31
expressions, integer	31
extend directive M680x0	95
extended directive, i960	88

extern directive 37

F

fail directive 37
faster processing (-f) 9
fatal signal 119
fepc register, V850 115
fepsw register, V850 115
ffloat directive, VAX 110
file directive 37
file directive, AMD 29K 56
file name, logical 37
files, including 40
files, input 7
fill directive 38
filling memory 48
float directive 38
float directive, i386 84
float directive, M680x0 95
float directive, VAX 110
floating point numbers 19
floating point numbers (double) 36
floating point numbers (single) 38, 48
floating point, AMD 29K (IEEE) 55
floating point, ARC (IEEE) 54
floating point, ARM (IEEE) 58
floating point, D10V 64
floating point, D30V 68
floating point, ESA/390 (IEEE) 79
floating point, H8/300 (IEEE) 70
floating point, H8/500 (IEEE) 73
floating point, HPPA (IEEE) 74
floating point, i386 84
floating point, i960 (IEEE) 88
floating point, M680x0 95
floating point, SH (IEEE) 104
floating point, SPARC (IEEE) 106
floating point, V850 (IEEE) 116
floating point, VAX 110
flonums 19
force_thumb directive, ARM 58
format of error messages 8
format of warning messages 8
formfeed (\f) 18
func directive 38
functions, in expressions 31

G

gbr960, i960 postprocessor 87
gfloat directive, VAX 110
global 108
global directive 38
gp register, MIPS 99
gp register, V850 114
grouping data 23

H

H8/300 addressing modes 69
H8/300 floating point (IEEE) 70
H8/300 line comment character 69
H8/300 line separator 69
H8/300 machine directives (none) 71
H8/300 opcode summary 71
H8/300 options (none) 69
H8/300 registers 69
H8/300 size suffixes 71
H8/300 support 69
H8/300H, assembling for 71
H8/500 addressing modes 72
H8/500 floating point (IEEE) 73
H8/500 line comment character 72
H8/500 line separator 72
H8/500 machine directives (none) 73
H8/500 opcode summary 73
H8/500 options (none) 72
H8/500 registers 72
H8/500 support 72
half directive, SPARC 106
hex character code (\xd...) 18
hexadecimal integers 19
hfloat directive, VAX 110
hi pseudo-op, V850 117
hi0 pseudo-op, V850 116
hidden directive 51
hilo pseudo-op, V850 117
HPPA directives not supported 75
HPPA floating point (IEEE) 74
HPPA Syntax 74
HPPA-only directives 75
hword directive 38

I

i370 support 78
i386 16-bit code 85
i386 conversion instructions 82
i386 floating point 84
i386 immediate operands 81

- `lflags` directive (ignored) 41
 - line comment character 16
 - line comment character, AMD 29K 55
 - line comment character, ARM 58
 - line comment character, D10V 62
 - line comment character, D30V 66
 - line comment character, H8/300 69
 - line comment character, H8/500 72
 - line comment character, M680x0 97
 - line comment character, SH 103
 - line comment character, V850 113
 - line comment character, Z8000 107
 - `line` directive 41
 - `line` directive, AMD 29K 56
 - line numbers, in input files 7
 - line numbers, in warnings/errors 8
 - line separator character 16
 - line separator, ARM 58
 - line separator, H8/300 69
 - line separator, H8/500 72
 - line separator, SH 103
 - line separator, Z8000 107
 - lines starting with # 16
 - linker 8
 - linker, and assembler 21
 - `linkonce` directive 41
 - `list` directive 42
 - listing control, turning off 44
 - listing control, turning on 42
 - listing control: new page 36
 - listing control: paper size 45
 - listing control: subtitle 46
 - listing control: title line 51
 - listings, enabling 9
 - little endian output, ARC 3
 - little endian output, MIPS 5
 - little endian output, PJ 4
 - little-endian output, ARC 54
 - little-endian output, MIPS 98
 - `ln` directive 42
 - `lo` pseudo-op, V850 116
 - local common symbols 41
 - local labels, retaining in output 10
 - local symbol names 27
 - location counter 28
 - location counter, advancing 44
 - logical file name 37
 - logical line number 41
 - logical line numbers 16
 - `long` directive 42
 - `long` directive, i386 84
 - `lp` register, V850 115
 - `lval` 108
- ## M
- M680x0 addressing modes 93
 - M680x0 architecture options 91
 - M680x0 branch improvement 95
 - M680x0 directives 95
 - M680x0 floating point 95
 - M680x0 immediate character 97
 - M680x0 line comment character 97
 - M680x0 opcodes 95
 - M680x0 options 91
 - M680x0 pseudo-opcodes 95
 - M680x0 size modifiers 93
 - M680x0 support 91
 - M680x0 syntax 93
 - machine dependencies 53
 - machine directives, AMD 29K 56
 - machine directives, ARC 54
 - machine directives, ARM 58
 - machine directives, H8/300 (none) 71
 - machine directives, H8/500 (none) 73
 - machine directives, i960 88
 - machine directives, SH 104
 - machine directives, SPARC 106
 - machine directives, V850 116
 - machine directives, VAX 110
 - machine independent directives 33
 - machine instructions (not covered) 6
 - machine-independent syntax 15
 - `macro` directive 43
 - macros 43
 - Macros, AMD 29K 55
 - macros, count executed 43
 - make rules 12
 - manual, structure and purpose 6
 - memory references, i386 83
 - merging text and data sections 12
 - messages from assembler 8
 - minus, permitted arguments 32
 - MIPS architecture options 98
 - MIPS big-endian output 98
 - MIPS debugging directives 100
 - MIPS ECOFF sections 99
 - MIPS endianness 5
 - MIPS ISA 5
 - MIPS ISA override 100
 - MIPS little-endian output 98
 - MIPS option stack 101
 - MIPS processor 98
 - MIT 93

MMX, i386	85
mnemonic suffixes, i386	81
mnemonics for opcodes, VAX	110
mnemonics, D10V	64
mnemonics, D30V	68
mnemonics, H8/300	71
mnemonics, H8/500	73
mnemonics, SH	104
mnemonics, Z8000	108
Motorola syntax for the 680x0	94
MRI compatibility mode	10
<code>mri</code> directive	42
MRI mode, temporarily	42
<code>mul</code> instruction, i386	86

N

<code>name</code>	108
named section	46
named sections	22
names, symbol	27
naming object file	12
new page, in listings	36
newline (<code>\n</code>)	18
newline, required at file end	16
<code>nolist</code> directive	44
NOP pseudo op, ARM	59
null-terminated strings	34
number constants	18
number of macros executed	43
numbered subsections	23
numbers, 16-bit	38
numeric values	31
<code>nword</code> directive, SPARC	106

O

object file	8
object file format	7
object file name	12
object file, after errors	13
obsolescent directives	52
<code>octa</code> directive	44
octal character code (<code>\ddd</code>)	18
octal integers	19
<code>offset</code> directive, V850	116
opcode mnemonics, VAX	110
opcode summary, D10V	64
opcode summary, D30V	68
opcode summary, H8/300	71
opcode summary, H8/500	73
opcode summary, SH	104

opcode summary, Z8000	108
opcodes for AMD 29K	56
opcodes for ARM	59
opcodes for V850	116
opcodes, i960	89
opcodes, M680x0	95
operand delimiters, i386	81
operand notation, VAX	112
operands in expressions	31
operator precedence	32
operators, in expressions	31
operators, permitted arguments	32
optimization, D10V	3
optimization, D30V	3
option summary	1
options for AMD29K (none)	55
options for ARC	54
options for ARM (none)	57
options for i386 (none)	81
options for SPARC	105
options for V850 (none)	113
options for VAX/VMS	109
options, all versions of assembler	9
options, command line	7
options, D10V	61
options, D30V	65
options, H8/300 (none)	69
options, H8/500 (none)	72
options, i960	87
options, M680x0	91
options, PJ	102
options, SH (none)	103
options, Z8000	107
<code>org</code> directive	44
other attribute, of <code>a.out</code> symbol	29
output file	8

P

<code>p2align</code> directive	44
<code>p2align1</code> directive	45
<code>p2alignw</code> directive	45
padding the location counter	33
padding the location counter given a power of two	44
padding the location counter given number of bytes	34
page, in listings	36
paper size, for listings	45
paths for <code>.include</code>	10
patterns, writing in memory	38
PIC code generation for ARM	57

- PJ endianness 4
 - PJ options 102
 - PJ support 102
 - plus, permitted arguments 32
 - precedence of operators 32
 - precision, floating point 19
 - prefix operators 32
 - prefixes, i386 82
 - preprocessing 15
 - preprocessing, turning on and off 15
 - primary attributes, COFF symbols 29
 - print directive 45
 - proc directive, SPARC 106
 - protected directive 51
 - protected registers, AMD 29K 55
 - pseudo-opcodes, M680x0 95
 - pseudo-ops for branch, VAX 111
 - pseudo-ops, machine independent 33
 - psize directive 45
 - psw register, V850 115
 - purgem directive 45
 - purpose of GNU assembler 6
- Q**
- quad directive 45
 - quad directive, i386 84
- R**
- real-mode code, i386 85
 - register directive, SPARC 106
 - register names, AMD 29K 55
 - register names, ARM 58
 - register names, H8/300 69
 - register names, V850 114
 - register names, VAX 112
 - register operands, i386 81
 - registers, D10V 62
 - registers, D30V 67
 - registers, H8/500 72
 - registers, i386 82
 - registers, SH 103
 - registers, Z8000 107
 - relocation 21
 - relocation example 23
 - repeat prefixes, i386 83
 - reporting bugs in assembler 119
 - rept directive 46
 - req directive, ARM 58
 - reserve directive, SPARC 106
 - return instructions, i386 81
 - rsect 108
- S**
- sbt1 directive 46
 - scl directive 46
 - sdaoff pseudo-op, V850 117
 - search path for .include 10
 - sect directive, AMD 29K 56
 - section directive 46
 - section directive, V850 116
 - section override prefixes, i386 82
 - section-relative addressing 22
 - sections 21
 - sections in messages, internal 23
 - sections, i386 81
 - sections, named 22
 - seg directive, SPARC 106
 - segm 108
 - set directive 47
 - SH addressing modes 103
 - SH floating point (IEEE) 104
 - SH line comment character 103
 - SH line separator 103
 - SH machine directives 104
 - SH opcode summary 104
 - SH options (none) 103
 - SH registers 103
 - SH support 103
 - short directive 48
 - SIMD, i386 85
 - single character constant 18
 - single directive 48
 - single directive, i386 84
 - sixteen bit integers 38
 - sixteen byte integer 44
 - size directive 48
 - size modifiers, D10V 61
 - size modifiers, D30V 65
 - size modifiers, M680x0 93
 - size prefixes, i386 83
 - size suffixes, H8/300 71
 - sizes operands, i386 81
 - skip directive 48
 - skip directive, M680x0 95
 - skip directive, SPARC 106
 - sleb128 directive 48
 - small objects, MIPS ECOFF 99
 - SOM symbol attributes 29
 - source program 7
 - source, destination operands; i386 81
 - sp register, V850 114

space directive	48
space used, maximum for assembly	13
SPARC architectures	105
SPARC data alignment	105
SPARC floating point (IEEE)	106
SPARC machine directives	106
SPARC options	105
SPARC support	105
special characters, M680x0	97
special purpose registers, AMD 29K	55
stabd directive	49
stabn directive	49
stabs directive	49
stabx directives	49
standard assembler sections	21
standard input, as input file	7
statement separator character	16
statement separator, ARM	58
statement separator, H8/300	69
statement separator, H8/500	72
statement separator, SH	103
statement separator, Z8000	107
statements, structure of	16
statistics, about assembly	13
stopping the assembly	33
string constants	17
string directive	49
string directive on HPPA	76
string literals	34
string, copying to object file	49
struct directive	50
structure debugging, COFF	50
sub-instruction ordering, D10V	62
sub-instruction ordering, D30V	66
sub-instructions, D10V	61
sub-instructions, D30V	65
subexpressions	31
subtitles for listings	46
subtraction, permitted arguments	32
summary of options	1
support	74
supporting files, including	40
suppressing warnings	13
sval	108
symbol attributes	28
symbol attributes, a.out	29
symbol attributes, COFF	29
symbol attributes, SOM	29
symbol descriptor, COFF	35
symbol names	27
symbol names, '\$' in	62, 66, 72, 103
symbol names, local	27
symbol names, temporary	27
symbol storage class (COFF)	46
symbol type	29
symbol type, COFF	51
symbol value	28
symbol value, setting	47
symbol values, assigning	27
symbol versioning	50
symbol visibility	51
symbol, common	35
symbol, making visible to linker	38
symbolic debuggers, information for	49
symbols	27
symbols with uppercase, VAX/VMS	109
symbols, assigning values to	36
symbols, local common	41
symver directive	50
syntax compatibility, i386	81
syntax, D10V	61
syntax, D30V	65
syntax, M680x0	93
syntax, machine-independent	15
sysproc directive, i960	89
T	
tab (\t)	18
tag directive	50
tdaoff pseudo-op, V850	117
temporary symbol names	27
text and data sections, joining	12
text directive	51
text section	22
tfloat directive, i386	84
thumb directive, ARM	58
Thumb support	57
thumb_func directive, ARM	58
thumb_set directive, ARM	58
time, total for assembly	13
title directive	51
tp register, V850	114
trusted compiler	9
turning preprocessing on and off	15
type directive	51
type of a symbol	29
U	
ualong directive, SH	104
uaword directive, SH	104
uleb128 directive	51
undefined section	23

`unsegm` 108
`use` directive, AMD 29K 56

V

V850 command line options 113
V850 floating point (IEEE) 116
V850 line comment character 113
V850 machine directives 116
V850 opcodes 116
V850 options (none) 113
V850 register names 114
V850 support 113
`val` directive 51
value attribute, COFF 51
value of a symbol 28
VAX bitfields not supported 113
VAX branch improvement 111
VAX command-line options ignored 109
VAX displacement sizing character 112
VAX floating point 110
VAX immediate character 112
VAX indirect character 112
VAX machine directives 110
VAX opcode mnemonics 110
VAX operand notation 112
VAX register names 112
VAX support 109
Vax-11 C compatibility 109
VAX/VMS options 109
version of assembler 13
versions of symbols 50
VMS (VAX) options 109

W

warning for altered difference tables 10
warning messages 8
warnings, causing error 13
warnings, suppressing 13
warnings, switching on 13
whitespace 15
whitespace, removed by preprocessor 15
wide floating point directives, VAX 110
`word` directive 52
`word` directive, H8/300 71
`word` directive, H8/500 73
`word` directive, i386 84
`word` directive, SPARC 106
writing patterns in memory 38
`wval` 108

X

`xword` directive, SPARC 106

Z

Z800 addressing modes 107
Z8000 directives 108
Z8000 line comment character 107
Z8000 line separator 107
Z8000 opcode summary 108
Z8000 options 107
Z8000 registers 107
Z8000 support 107
`zdaoff` pseudo-op, V850 118
`zero` register, V850 114
zero-terminated strings 34

Table of Contents

1	Overview	1
1.1	Structure of this Manual	6
1.2	The GNU Assembler	6
1.3	Object File Formats	7
1.4	Command Line	7
1.5	Input Files	7
1.6	Output (Object) File	8
1.7	Error and Warning Messages	8
2	Command-Line Options	9
2.1	Enable Listings: <code>-a[cdhlms]</code>	9
2.2	<code>-D</code>	9
2.3	Work Faster: <code>-f</code>	9
2.4	.include search path: <code>-I path</code>	10
2.5	Difference Tables: <code>-K</code>	10
2.6	Include Local Labels: <code>-L</code>	10
2.7	Assemble in MRI Compatibility Mode: <code>-M</code>	10
2.8	Dependency tracking: <code>--MD</code>	12
2.9	Name the Object File: <code>-o</code>	12
2.10	Join Data and Text Sections: <code>-R</code>	12
2.11	Display Assembly Statistics: <code>--statistics</code>	13
2.12	Compatible output: <code>--traditional-format</code>	13
2.13	Announce Version: <code>-v</code>	13
2.14	Control Warnings: <code>-W</code> , <code>--warn</code> , <code>--no-warn</code> , <code>--fatal-warnings</code>	13
2.15	Generate Object File in Spite of Errors: <code>-Z</code>	13
3	Syntax	15
3.1	Preprocessing	15
3.2	Whitespace	15
3.3	Comments	15
3.4	Symbols	16
3.5	Statements	16
3.6	Constants	17
3.6.1	Character Constants	17
3.6.1.1	Strings	17
3.6.1.2	Characters	18
3.6.2	Number Constants	18
3.6.2.1	Integers	19
3.6.2.2	Bignums	19
3.6.2.3	Flonums	19

4	Sections and Relocation	21
4.1	Background	21
4.2	Linker Sections	22
4.3	Assembler Internal Sections	23
4.4	Sub-Sections	23
4.5	bss Section	24
5	Symbols	27
5.1	Labels	27
5.2	Giving Symbols Other Values	27
5.3	Symbol Names	27
5.4	The Special Dot Symbol	28
5.5	Symbol Attributes	28
5.5.1	Value	28
5.5.2	Type	29
5.5.3	Symbol Attributes: <code>a.out</code>	29
5.5.3.1	Descriptor	29
5.5.3.2	Other	29
5.5.4	Symbol Attributes for COFF	29
5.5.4.1	Primary Attributes	29
5.5.4.2	Auxiliary Attributes	29
5.5.5	Symbol Attributes for SOM	29
6	Expressions	31
6.1	Empty Expressions	31
6.2	Integer Expressions	31
6.2.1	Arguments	31
6.2.2	Operators	31
6.2.3	Prefix Operator	32
6.2.4	Infix Operators	32
7	Assembler Directives	33
7.1	<code>.abort</code>	33
7.2	<code>.ABORT</code>	33
7.3	<code>.align <i>abs-expr</i>, <i>abs-expr</i>, <i>abs-expr</i></code>	33
7.4	<code>.ascii "string"</code>	34
7.5	<code>.asciz "string"</code>	34
7.6	<code>.balign[wl] <i>abs-expr</i>, <i>abs-expr</i>, <i>abs-expr</i></code>	34
7.7	<code>.byte expressions</code>	34
7.8	<code>.comm <i>symbol</i>, <i>length</i></code>	35
7.9	<code>.data subsection</code>	35
7.10	<code>.def <i>name</i></code>	35
7.11	<code>.desc <i>symbol</i>, <i>abs-expression</i></code>	35
7.12	<code>.dim</code>	35
7.13	<code>.double <i>flonums</i></code>	36
7.14	<code>.eject</code>	36
7.15	<code>.else</code>	36

7.16	<code>.end</code>	36
7.17	<code>.endif</code>	36
7.18	<code>.endfunc</code>	36
7.19	<code>.endif</code>	36
7.20	<code>.equ symbol, expression</code>	36
7.21	<code>.equiv symbol, expression</code>	37
7.22	<code>.err</code>	37
7.23	<code>.exitm</code>	37
7.24	<code>.extern</code>	37
7.25	<code>.fail expression</code>	37
7.26	<code>.file string</code>	37
7.27	<code>.fill repeat, size, value</code>	38
7.28	<code>.float flonums</code>	38
7.29	<code>.func name[,label]</code>	38
7.30	<code>.global symbol, .globl symbol</code>	38
7.31	<code>.hword expressions</code>	38
7.32	<code>.ident</code>	39
7.33	<code>.if absolute expression</code>	39
7.34	<code>.include "file"</code>	40
7.35	<code>.int expressions</code>	40
7.36	<code>.irp symbol, values</code>	40
7.37	<code>.irpc symbol, values</code>	40
7.38	<code>.lcomm symbol, length</code>	41
7.39	<code>.lflags</code>	41
7.40	<code>.line line-number</code>	41
7.41	<code>.linkonce [type]</code>	41
7.42	<code>.ln line-number</code>	42
7.43	<code>.mri val</code>	42
7.44	<code>.list</code>	42
7.45	<code>.long expressions</code>	42
7.46	<code>.macro</code>	43
7.47	<code>.nolist</code>	44
7.48	<code>.octa bignums</code>	44
7.49	<code>.org new-lc, fill</code>	44
7.50	<code>.p2align[w1] abs-expr, abs-expr, abs-expr</code>	44
7.51	<code>.print string</code>	45
7.52	<code>.psize lines, columns</code>	45
7.53	<code>.purge name</code>	45
7.54	<code>.quad bignums</code>	45
7.55	<code>.rept count</code>	46
7.56	<code>.sbttl "subheading"</code>	46
7.57	<code>.scl class</code>	46
7.58	<code>.section name</code>	46
7.59	<code>.set symbol, expression</code>	47
7.60	<code>.short expressions</code>	48
7.61	<code>.single flonums</code>	48
7.62	<code>.size</code>	48
7.63	<code>.sleb128 expressions</code>	48

7.64	<code>.skip size , fill</code>	48
7.65	<code>.space size , fill</code>	48
7.66	<code>.stabd, .stabn, .stabs</code>	49
7.67	<code>.string "str"</code>	49
7.68	<code>.struct expression</code>	50
7.69	<code>.symver</code>	50
7.70	<code>.tag structname</code>	50
7.71	<code>.text subsection</code>	51
7.72	<code>.title "heading"</code>	51
7.73	<code>.type int</code>	51
7.74	<code>.val addr</code>	51
7.75	<code>.uleb128 expressions</code>	51
7.76	<code>.internal, .hidden, .protected</code>	51
7.77	<code>.word expressions</code>	52
7.78	Deprecated Directives.....	52

8 Machine Dependent Features 53

8.1	ARC Dependent Features.....	54
8.1.1	Options.....	54
8.1.2	Floating Point.....	54
8.1.3	ARC Machine Directives.....	54
8.2	AMD 29K Dependent Features.....	55
8.2.1	Options.....	55
8.2.2	Syntax.....	55
8.2.2.1	Macros.....	55
8.2.2.2	Special Characters.....	55
8.2.2.3	Register Names.....	55
8.2.3	Floating Point.....	55
8.2.4	AMD 29K Machine Directives.....	56
8.2.5	Opcodes.....	56
8.3	ARM Dependent Features.....	57
8.3.1	Options.....	57
8.3.2	Syntax.....	58
8.3.2.1	Special Characters.....	58
8.3.2.2	Register Names.....	58
8.3.3	Floating Point.....	58
8.3.4	ARM Machine Directives.....	58
8.3.5	Opcodes.....	59
8.4	D10V Dependent Features.....	61
8.4.1	D10V Options.....	61
8.4.2	Syntax.....	61
8.4.2.1	Size Modifiers.....	61
8.4.2.2	Sub-Instructions.....	61
8.4.2.3	Special Characters.....	62
8.4.2.4	Register Names.....	62
8.4.2.5	Addressing Modes.....	63
8.4.2.6	@WORD Modifier.....	64
8.4.3	Floating Point.....	64

8.4.4	Opcodes	64
8.5	D30V Dependent Features	65
8.5.1	D30V Options	65
8.5.2	Syntax	65
8.5.2.1	Size Modifiers	65
8.5.2.2	Sub-Instructions	65
8.5.2.3	Special Characters	66
8.5.2.4	Guarded Execution	67
8.5.2.5	Register Names	67
8.5.2.6	Addressing Modes	68
8.5.3	Floating Point	68
8.5.4	Opcodes	68
8.6	H8/300 Dependent Features	69
8.6.1	Options	69
8.6.2	Syntax	69
8.6.2.1	Special Characters	69
8.6.2.2	Register Names	69
8.6.2.3	Addressing Modes	69
8.6.3	Floating Point	70
8.6.4	H8/300 Machine Directives	71
8.6.5	Opcodes	71
8.7	H8/500 Dependent Features	72
8.7.1	Options	72
8.7.2	Syntax	72
8.7.2.1	Special Characters	72
8.7.2.2	Register Names	72
8.7.2.3	Addressing Modes	72
8.7.3	Floating Point	73
8.7.4	H8/500 Machine Directives	73
8.7.5	Opcodes	73
8.8	HPPA Dependent Features	74
8.8.1	Notes	74
8.8.2	Options	74
8.8.3	Syntax	74
8.8.4	Floating Point	74
8.8.5	HPPA Assembler Directives	75
8.8.6	Opcodes	77
8.9	ESA/390 Dependent Features	78
8.9.1	Notes	78
8.9.2	Options	78
8.9.3	Syntax	78
8.9.4	Floating Point	79
8.9.5	ESA/390 Assembler Directives	79
8.9.6	Opcodes	80
8.10	80386 Dependent Features	81
8.10.1	Options	81
8.10.2	AT&T Syntax versus Intel Syntax	81
8.10.3	Instruction Naming	81

8.10.4	Register Naming	82
8.10.5	Instruction Prefixes	82
8.10.6	Memory References	83
8.10.7	Handling of Jump Instructions	84
8.10.8	Floating Point	84
8.10.9	Intel's MMX and AMD's 3DNow! SIMD Operations	85
8.10.10	Writing 16-bit Code	85
8.10.11	AT&T Syntax bugs	86
8.10.12	Notes	86
8.11	Intel 80960 Dependent Features	87
8.11.1	i960 Command-line Options	87
8.11.2	Floating Point	88
8.11.3	i960 Machine Directives	88
8.11.4	i960 Opcodes	89
8.11.4.1	callj	89
8.11.4.2	Compare-and-Branch	89
8.12	M680x0 Dependent Features	91
8.12.1	M680x0 Options	91
8.12.2	Syntax	93
8.12.3	Motorola Syntax	94
8.12.4	Floating Point	95
8.12.5	680x0 Machine Directives	95
8.12.6	Opcodes	95
8.12.6.1	Branch Improvement	95
8.12.6.2	Special Characters	97
8.13	MIPS Dependent Features	98
8.13.1	Assembler options	98
8.13.2	MIPS ECOFF object code	99
8.13.3	Directives for debugging information	100
8.13.4	Directives to override the ISA level	100
8.13.5	Directives for extending MIPS 16 bit instructions	100
8.13.6	Directive to mark data as an instruction	100
8.13.7	Directives to save and restore options	101
8.14	picoJava Dependent Features	102
8.14.1	Options	102
8.15	Hitachi SH Dependent Features	103
8.15.1	Options	103
8.15.2	Syntax	103
8.15.2.1	Special Characters	103
8.15.2.2	Register Names	103
8.15.2.3	Addressing Modes	103
8.15.3	Floating Point	104
8.15.4	SH Machine Directives	104
8.15.5	Opcodes	104
8.16	SPARC Dependent Features	105
8.16.1	Options	105

8.16.2	Enforcing aligned data	105
8.16.3	Floating Point	106
8.16.4	Sparc Machine Directives	106
8.17	Z8000 Dependent Features	107
8.17.1	Options	107
8.17.2	Syntax	107
8.17.2.1	Special Characters	107
8.17.2.2	Register Names	107
8.17.2.3	Addressing Modes	107
8.17.3	Assembler Directives for the Z8000	108
8.17.4	Opcodes	108
8.18	VAX Dependent Features	109
8.18.1	VAX Command-Line Options	109
8.18.2	VAX Floating Point	110
8.18.3	Vax Machine Directives	110
8.18.4	VAX Opcodes	110
8.18.5	VAX Branch Improvement	111
8.18.6	VAX Operands	112
8.18.7	Not Supported on VAX	113
8.19	v850 Dependent Features	113
8.19.1	Options	113
8.19.2	Syntax	113
8.19.2.1	Special Characters	113
8.19.2.2	Register Names	114
8.19.3	Floating Point	116
8.19.4	V850 Machine Directives	116
8.19.5	Opcodes	116
9	Reporting Bugs	119
9.1	Have you found a bug?	119
9.2	How to report bugs	119
10	Acknowledgements	123
Index	125

