

Using `as`

The GNU Assembler

(GNU Binutils)

Version 2.21.1

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, 1992, 1993, 1994, 1995, 1996, 1997, 1998, 1999, 2000, 2001, 2002, 2006, 2007, 2008, 2009, 2010, 2011 Free Software Foundation, Inc.

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

Table of Contents

| | | |
|----------|--|-----------|
| 1 | Overview | 1 |
| 1.1 | Structure of this Manual | 16 |
| 1.2 | The GNU Assembler | 17 |
| 1.3 | Object File Formats | 17 |
| 1.4 | Command Line | 17 |
| 1.5 | Input Files | 17 |
| 1.6 | Output (Object) File | 18 |
| 1.7 | Error and Warning Messages | 18 |
| 2 | Command-Line Options | 21 |
| 2.1 | Enable Listings: ‘-a[cdghlms]’ | 21 |
| 2.2 | ‘--alternate’ | 21 |
| 2.3 | ‘-D’ | 22 |
| 2.4 | Work Faster: ‘-f’ | 22 |
| 2.5 | .include Search Path: ‘-I’ <i>path</i> | 22 |
| 2.6 | Difference Tables: ‘-K’ | 22 |
| 2.7 | Include Local Symbols: ‘-L’ | 22 |
| 2.8 | Configuring listing output: ‘--listing’ | 22 |
| 2.9 | Assemble in MRI Compatibility Mode: ‘-M’ | 23 |
| 2.10 | Dependency Tracking: ‘--MD’ | 25 |
| 2.11 | Name the Object File: ‘-o’ | 25 |
| 2.12 | Join Data and Text Sections: ‘-R’ | 25 |
| 2.13 | Display Assembly Statistics: ‘--statistics’ | 25 |
| 2.14 | Compatible Output: ‘--traditional-format’ | 25 |
| 2.15 | Announce Version: ‘-v’ | 25 |
| 2.16 | Control Warnings: ‘-W’, ‘--warn’, ‘--no-warn’, ‘--fatal-warnings’ | 26 |
| 2.17 | Generate Object File in Spite of Errors: ‘-Z’ | 26 |
| 3 | Syntax | 27 |
| 3.1 | Preprocessing | 27 |
| 3.2 | Whitespace | 27 |
| 3.3 | Comments | 27 |
| 3.4 | Symbols | 28 |
| 3.5 | Statements | 28 |
| 3.6 | Constants | 29 |
| 3.6.1 | Character Constants | 29 |
| 3.6.1.1 | Strings | 29 |
| 3.6.1.2 | Characters | 30 |
| 3.6.2 | Number Constants | 30 |
| 3.6.2.1 | Integers | 30 |
| 3.6.2.2 | Bignums | 31 |
| 3.6.2.3 | Flonums | 31 |

| | | |
|----------|---|-----------|
| 4 | Sections and Relocation | 33 |
| 4.1 | Background | 33 |
| 4.2 | Linker Sections | 34 |
| 4.3 | Assembler Internal Sections | 35 |
| 4.4 | Sub-Sections | 35 |
| 4.5 | bss Section | 36 |
| 5 | Symbols | 39 |
| 5.1 | Labels | 39 |
| 5.2 | Giving Symbols Other Values | 39 |
| 5.3 | Symbol Names | 39 |
| 5.4 | The Special Dot Symbol | 41 |
| 5.5 | Symbol Attributes | 41 |
| 5.5.1 | Value | 41 |
| 5.5.2 | Type | 41 |
| 5.5.3 | Symbol Attributes: <code>a.out</code> | 41 |
| 5.5.3.1 | Descriptor | 42 |
| 5.5.3.2 | Other | 42 |
| 5.5.4 | Symbol Attributes for COFF | 42 |
| 5.5.4.1 | Primary Attributes | 42 |
| 5.5.4.2 | Auxiliary Attributes | 42 |
| 5.5.5 | Symbol Attributes for SOM | 42 |
| 6 | Expressions | 43 |
| 6.1 | Empty Expressions | 43 |
| 6.2 | Integer Expressions | 43 |
| 6.2.1 | Arguments | 43 |
| 6.2.2 | Operators | 43 |
| 6.2.3 | Prefix Operator | 43 |
| 6.2.4 | Infix Operators | 44 |
| 7 | Assembler Directives | 47 |
| 7.1 | <code>.abort</code> | 47 |
| 7.2 | <code>.ABORT (COFF)</code> | 47 |
| 7.3 | <code>.align abs-expr, abs-expr, abs-expr</code> | 47 |
| 7.4 | <code>.altmacro</code> | 48 |
| 7.5 | <code>.ascii "string"</code> | 48 |
| 7.6 | <code>.asciz "string"</code> | 48 |
| 7.7 | <code>.balign[wl] abs-expr, abs-expr, abs-expr</code> | 48 |
| 7.8 | <code>.byte expressions</code> | 49 |
| 7.9 | <code>.cfi_sections section_list</code> | 49 |
| 7.10 | <code>.cfi_startproc [simple]</code> | 49 |
| 7.11 | <code>.cfi_endproc</code> | 49 |
| 7.12 | <code>.cfi_personality encoding [, exp]</code> | 49 |
| 7.13 | <code>.cfi_lsdas encoding [, exp]</code> | 49 |
| 7.14 | <code>.cfi_def_cfa register, offset</code> | 50 |
| 7.15 | <code>.cfi_def_cfa_register register</code> | 50 |

| | | |
|------|---|----|
| 7.16 | .cfi_def_cfa_offset offset | 50 |
| 7.17 | .cfi_adjust_cfa_offset offset | 50 |
| 7.18 | .cfi_offset register, offset | 50 |
| 7.19 | .cfi_rel_offset register, offset | 50 |
| 7.20 | .cfi_register register1, register2 | 50 |
| 7.21 | .cfi_restore register | 50 |
| 7.22 | .cfi_undefined register | 50 |
| 7.23 | .cfi_same_value register | 50 |
| 7.24 | .cfi_remember_state, | 50 |
| 7.25 | .cfi_return_column register | 51 |
| 7.26 | .cfi_signal_frame | 51 |
| 7.27 | .cfi_window_save | 51 |
| 7.28 | .cfi_escape expression[, ...] | 51 |
| 7.29 | .cfi_val_encoded_addr register, encoding, label | 51 |
| 7.30 | .comm symbol , length | 51 |
| 7.31 | .data subsection | 52 |
| 7.32 | .def name | 52 |
| 7.33 | .desc symbol, abs-expression | 52 |
| 7.34 | .dim | 52 |
| 7.35 | .double flonums | 52 |
| 7.36 | .eject | 52 |
| 7.37 | .else | 52 |
| 7.38 | .elseif | 52 |
| 7.39 | .end | 53 |
| 7.40 | .endif | 53 |
| 7.41 | .endfunc | 53 |
| 7.42 | .endif | 53 |
| 7.43 | .equ symbol, expression | 53 |
| 7.44 | .equiv symbol, expression | 53 |
| 7.45 | .eqv symbol, expression | 53 |
| 7.46 | .err | 53 |
| 7.47 | .error "string" | 54 |
| 7.48 | .exitm | 54 |
| 7.49 | .extern | 54 |
| 7.50 | .fail expression | 54 |
| 7.51 | .file | 54 |
| 7.52 | .fill repeat , size , value | 55 |
| 7.53 | .float flonums | 55 |
| 7.54 | .func name [,label] | 55 |
| 7.55 | .global symbol, .globl symbol | 55 |
| 7.56 | .gnu_attribute tag,value | 55 |
| 7.57 | .hidden names | 55 |
| 7.58 | .hword expressions | 56 |
| 7.59 | .ident | 56 |
| 7.60 | .if absolute expression | 56 |
| 7.61 | .incbin "file" [,skip[,count]] | 57 |
| 7.62 | .include "file" | 57 |
| 7.63 | .int expressions | 58 |

| | | |
|-------|---|----|
| 7.64 | .internal names | 58 |
| 7.65 | .irp symbol, values | 58 |
| 7.66 | .irpc symbol, values | 58 |
| 7.67 | .lcomm symbol , length | 59 |
| 7.68 | .lflags | 59 |
| 7.69 | .line line-number | 59 |
| 7.70 | .linkonce [type] | 59 |
| 7.71 | .list | 60 |
| 7.72 | .ln line-number | 60 |
| 7.73 | .loc fileno lineno [column] [options] | 60 |
| 7.74 | .loc_mark_labels enable | 61 |
| 7.75 | .local names | 61 |
| 7.76 | .long expressions | 61 |
| 7.77 | .macro | 61 |
| 7.78 | .mri val | 64 |
| 7.79 | .noaltmacro | 64 |
| 7.80 | .nolist | 64 |
| 7.81 | .octa bignums | 64 |
| 7.82 | .org new-lc , fill | 64 |
| 7.83 | .p2align[wl] abs-expr, abs-expr, abs-expr | 64 |
| 7.84 | .popsection | 65 |
| 7.85 | .previous | 65 |
| 7.86 | .print string | 66 |
| 7.87 | .protected names | 66 |
| 7.88 | .psize lines , columns | 66 |
| 7.89 | .purgem name | 66 |
| 7.90 | .pushsection name [, subsection] [, "flags"[, @type[, arguments]]] | 66 |
| 7.91 | .quad bignums | 67 |
| 7.92 | .reloc offset, reloc_name[, expression] | 67 |
| 7.93 | .rept count | 67 |
| 7.94 | .sbttl "subheading" | 67 |
| 7.95 | .scl class | 68 |
| 7.96 | .section name | 68 |
| 7.97 | .set symbol, expression | 70 |
| 7.98 | .short expressions | 71 |
| 7.99 | .single flonums | 71 |
| 7.100 | .size | 71 |
| 7.101 | .skip size , fill | 71 |
| 7.102 | .sleb128 expressions | 71 |
| 7.103 | .space size , fill | 71 |
| 7.104 | .stabd, .stabn, .stabs | 72 |
| 7.105 | .string "str", .string8 "str", .string16 | 72 |
| 7.106 | .struct expression | 73 |
| 7.107 | .subsection name | 73 |
| 7.108 | .symver | 73 |
| 7.109 | .tag structname | 74 |
| 7.110 | .text subsection | 74 |

| | | |
|----------|---|-----------|
| 7.111 | .title "heading" | 74 |
| 7.112 | .type | 74 |
| 7.113 | .uleb128 expressions | 75 |
| 7.114 | .val addr | 76 |
| 7.115 | .version "string" | 76 |
| 7.116 | .vtable_entry table, offset | 76 |
| 7.117 | .vtable_inherit child, parent | 76 |
| 7.118 | .warning "string" | 76 |
| 7.119 | .weak names | 76 |
| 7.120 | .weakref alias, target | 76 |
| 7.121 | .word expressions | 77 |
| 7.122 | Deprecated Directives | 77 |
| 8 | Object Attributes | 79 |
| 8.1 | GNU Object Attributes | 79 |
| 8.1.1 | Common GNU attributes | 79 |
| 8.1.2 | MIPS Attributes | 79 |
| 8.1.3 | PowerPC Attributes | 80 |
| 8.2 | Defining New Object Attributes | 80 |
| 9 | Machine Dependent Features | 81 |
| 9.1 | Alpha Dependent Features | 82 |
| 9.1.1 | Notes | 82 |
| 9.1.2 | Options | 82 |
| 9.1.3 | Syntax | 83 |
| 9.1.3.1 | Special Characters | 83 |
| 9.1.3.2 | Register Names | 83 |
| 9.1.3.3 | Relocations | 83 |
| 9.1.4 | Floating Point | 85 |
| 9.1.5 | Alpha Assembler Directives | 85 |
| 9.1.6 | Opcodes | 88 |
| 9.2 | ARC Dependent Features | 89 |
| 9.2.1 | Options | 89 |
| 9.2.2 | Syntax | 89 |
| 9.2.2.1 | Special Characters | 89 |
| 9.2.2.2 | Register Names | 89 |
| 9.2.3 | Floating Point | 89 |
| 9.2.4 | ARC Machine Directives | 89 |
| 9.2.5 | Opcodes | 92 |
| 9.3 | ARM Dependent Features | 93 |
| 9.3.1 | Options | 93 |
| 9.3.2 | Syntax | 95 |
| 9.3.2.1 | Instruction Set Syntax | 95 |
| 9.3.2.2 | Special Characters | 96 |
| 9.3.2.3 | Register Names | 96 |
| 9.3.2.4 | NEON Alignment Specifiers | 96 |
| 9.3.3 | Floating Point | 96 |
| 9.3.3.1 | ARM relocation generation | 96 |

| | | |
|---------|--|-----|
| 9.3.4 | ARM Machine Directives | 97 |
| 9.3.5 | Opcodes | 102 |
| 9.3.6 | Mapping Symbols | 102 |
| 9.3.7 | Unwinding | 103 |
| 9.4 | AVR Dependent Features | 106 |
| 9.4.1 | Options | 106 |
| 9.4.2 | Syntax | 107 |
| 9.4.2.1 | Special Characters | 107 |
| 9.4.2.2 | Register Names | 107 |
| 9.4.2.3 | Relocatable Expression Modifiers | 107 |
| 9.4.3 | Opcodes | 108 |
| 9.5 | Blackfin Dependent Features | 112 |
| 9.5.1 | Options | 112 |
| 9.5.2 | Syntax | 112 |
| 9.5.3 | Directives | 114 |
| 9.6 | CR16 Dependent Features | 115 |
| 9.6.1 | CR16 Operand Qualifiers | 115 |
| 9.7 | CRIS Dependent Features | 117 |
| 9.7.1 | Command-line Options | 117 |
| 9.7.2 | Instruction expansion | 118 |
| 9.7.3 | Symbols | 118 |
| 9.7.4 | Syntax | 118 |
| 9.7.4.1 | Special Characters | 119 |
| 9.7.4.2 | Symbols in position-independent code | 119 |
| 9.7.4.3 | Register names | 120 |
| 9.7.4.4 | Assembler Directives | 120 |
| 9.8 | D10V Dependent Features | 122 |
| 9.8.1 | D10V Options | 122 |
| 9.8.2 | Syntax | 122 |
| 9.8.2.1 | Size Modifiers | 122 |
| 9.8.2.2 | Sub-Instructions | 122 |
| 9.8.2.3 | Special Characters | 123 |
| 9.8.2.4 | Register Names | 123 |
| 9.8.2.5 | Addressing Modes | 124 |
| 9.8.2.6 | @WORD Modifier | 125 |
| 9.8.3 | Floating Point | 125 |
| 9.8.4 | Opcodes | 125 |
| 9.9 | D30V Dependent Features | 126 |
| 9.9.1 | D30V Options | 126 |
| 9.9.2 | Syntax | 126 |
| 9.9.2.1 | Size Modifiers | 126 |
| 9.9.2.2 | Sub-Instructions | 126 |
| 9.9.2.3 | Special Characters | 126 |
| 9.9.2.4 | Guarded Execution | 127 |
| 9.9.2.5 | Register Names | 128 |
| 9.9.2.6 | Addressing Modes | 129 |
| 9.9.3 | Floating Point | 129 |
| 9.9.4 | Opcodes | 129 |

| | | |
|----------|---|-----|
| 9.10 | H8/300 Dependent Features..... | 130 |
| 9.10.1 | Options..... | 130 |
| 9.10.2 | Syntax..... | 130 |
| 9.10.2.1 | Special Characters..... | 130 |
| 9.10.2.2 | Register Names..... | 130 |
| 9.10.2.3 | Addressing Modes..... | 130 |
| 9.10.3 | Floating Point..... | 131 |
| 9.10.4 | H8/300 Machine Directives..... | 132 |
| 9.10.5 | Opcodes..... | 132 |
| 9.11 | HPPA Dependent Features..... | 133 |
| 9.11.1 | Notes..... | 133 |
| 9.11.2 | Options..... | 133 |
| 9.11.3 | Syntax..... | 133 |
| 9.11.4 | Floating Point..... | 133 |
| 9.11.5 | HPPA Assembler Directives..... | 133 |
| 9.11.6 | Opcodes..... | 137 |
| 9.12 | ESA/390 Dependent Features..... | 138 |
| 9.12.1 | Notes..... | 138 |
| 9.12.2 | Options..... | 138 |
| 9.12.3 | Syntax..... | 138 |
| 9.12.4 | Floating Point..... | 139 |
| 9.12.5 | ESA/390 Assembler Directives..... | 139 |
| 9.12.6 | Opcodes..... | 140 |
| 9.13 | 80386 Dependent Features..... | 141 |
| 9.13.1 | Options..... | 141 |
| 9.13.2 | x86 specific Directives..... | 142 |
| 9.13.3 | AT&T Syntax versus Intel Syntax..... | 142 |
| 9.13.4 | Instruction Naming..... | 143 |
| 9.13.5 | AT&T Mnemonic versus Intel Mnemonic..... | 144 |
| 9.13.6 | Register Naming..... | 144 |
| 9.13.7 | Instruction Prefixes..... | 145 |
| 9.13.8 | Memory References..... | 146 |
| 9.13.9 | Handling of Jump Instructions..... | 147 |
| 9.13.10 | Floating Point..... | 147 |
| 9.13.11 | Intel's MMX and AMD's 3DNow! SIMD Operations... | 147 |
| 9.13.12 | AMD's Lightweight Profiling Instructions..... | 148 |
| 9.13.13 | Writing 16-bit Code..... | 148 |
| 9.13.14 | AT&T Syntax bugs..... | 149 |
| 9.13.15 | Specifying CPU Architecture..... | 149 |
| 9.13.16 | Notes..... | 150 |
| 9.14 | Intel i860 Dependent Features..... | 151 |
| 9.14.1 | i860 Notes..... | 151 |
| 9.14.2 | i860 Command-line Options..... | 151 |
| 9.14.2.1 | SVR4 compatibility options..... | 151 |
| 9.14.2.2 | Other options..... | 151 |
| 9.14.3 | i860 Machine Directives..... | 151 |
| 9.14.4 | i860 Opcodes..... | 152 |
| 9.14.4.1 | Other instruction support (pseudo-instructions).... | 152 |

| | | |
|----------|---|-----|
| 9.15 | Intel 80960 Dependent Features | 154 |
| 9.15.1 | i960 Command-line Options | 154 |
| 9.15.2 | Floating Point | 155 |
| 9.15.3 | i960 Machine Directives | 155 |
| 9.15.4 | i960 Opcodes | 156 |
| 9.15.4.1 | callj | 156 |
| 9.15.4.2 | Compare-and-Branch | 156 |
| 9.16 | IA-64 Dependent Features | 158 |
| 9.16.1 | Options | 158 |
| 9.16.2 | Syntax | 159 |
| 9.16.2.1 | Special Characters | 159 |
| 9.16.2.2 | Register Names | 159 |
| 9.16.2.3 | IA-64 Processor-Status-Register (PSR) Bit Names | 159 |
| 9.16.2.4 | Relocations | 159 |
| 9.16.3 | Opcodes | 160 |
| 9.17 | IP2K Dependent Features | 161 |
| 9.17.1 | IP2K Options | 161 |
| 9.18 | LM32 Dependent Features | 162 |
| 9.18.1 | Options | 162 |
| 9.18.2 | Syntax | 162 |
| 9.18.2.1 | Register Names | 162 |
| 9.18.2.2 | Relocatable Expression Modifiers | 163 |
| 9.18.3 | Opcodes | 164 |
| 9.19 | M32C Dependent Features | 165 |
| 9.19.1 | M32C Options | 165 |
| 9.19.2 | Symbolic Operand Modifiers | 165 |
| 9.20 | M32R Dependent Features | 167 |
| 9.20.1 | M32R Options | 167 |
| 9.20.2 | M32R Directives | 168 |
| 9.20.3 | M32R Warnings | 169 |
| 9.21 | M680x0 Dependent Features | 171 |
| 9.21.1 | M680x0 Options | 171 |
| 9.21.2 | Syntax | 174 |
| 9.21.3 | Motorola Syntax | 175 |
| 9.21.4 | Floating Point | 176 |
| 9.21.5 | 680x0 Machine Directives | 176 |
| 9.21.6 | Opcodes | 177 |
| 9.21.6.1 | Branch Improvement | 177 |
| 9.21.6.2 | Special Characters | 178 |
| 9.22 | M68HC11 and M68HC12 Dependent Features | 179 |
| 9.22.1 | M68HC11 and M68HC12 Options | 179 |
| 9.22.2 | Syntax | 180 |
| 9.22.3 | Symbolic Operand Modifiers | 181 |
| 9.22.4 | Assembler Directives | 182 |
| 9.22.5 | Floating Point | 182 |
| 9.22.6 | Opcodes | 183 |
| 9.22.6.1 | Branch Improvement | 183 |

| | | |
|----------|--|-----|
| 9.23 | MicroBlaze Dependent Features | 184 |
| 9.23.1 | Directives | 184 |
| 9.24 | MIPS Dependent Features | 185 |
| 9.24.1 | Assembler options | 185 |
| 9.24.2 | MIPS ECOFF object code | 189 |
| 9.24.3 | Directives for debugging information | 190 |
| 9.24.4 | Directives to override the size of symbols | 190 |
| 9.24.5 | Directives to override the ISA level | 191 |
| 9.24.6 | Directives for extending MIPS 16 bit instructions | 191 |
| 9.24.7 | Directive to mark data as an instruction | 191 |
| 9.24.8 | Directives to save and restore options | 192 |
| 9.24.9 | Directives to control generation of MIPS ASE instructions | 192 |
| 9.24.10 | Directives to override floating-point options | 192 |
| 9.25 | MMIX Dependent Features | 194 |
| 9.25.1 | Command-line Options | 194 |
| 9.25.2 | Instruction expansion | 195 |
| 9.25.3 | Syntax | 195 |
| 9.25.3.1 | Special Characters | 195 |
| 9.25.3.2 | Symbols | 196 |
| 9.25.3.3 | Register names | 196 |
| 9.25.3.4 | Assembler Directives | 197 |
| 9.25.4 | Differences to <code>mmixal</code> | 199 |
| 9.26 | MSP 430 Dependent Features | 201 |
| 9.26.1 | Options | 201 |
| 9.26.2 | Syntax | 201 |
| 9.26.2.1 | Macros | 201 |
| 9.26.2.2 | Special Characters | 201 |
| 9.26.2.3 | Register Names | 201 |
| 9.26.2.4 | Assembler Extensions | 201 |
| 9.26.3 | Floating Point | 202 |
| 9.26.4 | MSP 430 Machine Directives | 202 |
| 9.26.5 | Opcodes | 203 |
| 9.26.6 | Profiling Capability | 203 |
| 9.27 | PDP-11 Dependent Features | 205 |
| 9.27.1 | Options | 205 |
| 9.27.1.1 | Code Generation Options | 205 |
| 9.27.1.2 | Instruction Set Extension Options | 205 |
| 9.27.1.3 | CPU Model Options | 206 |
| 9.27.1.4 | Machine Model Options | 206 |
| 9.27.2 | Assembler Directives | 207 |
| 9.27.3 | PDP-11 Assembly Language Syntax | 207 |
| 9.27.4 | Instruction Naming | 207 |
| 9.27.5 | Synthetic Instructions | 208 |
| 9.28 | picoJava Dependent Features | 209 |
| 9.28.1 | Options | 209 |
| 9.29 | PowerPC Dependent Features | 210 |
| 9.29.1 | Options | 210 |

| | | |
|----------|--|-----|
| 9.29.2 | PowerPC Assembler Directives | 212 |
| 9.30 | RX Dependent Features | 213 |
| 9.30.1 | RX Options | 213 |
| 9.30.2 | Symbolic Operand Modifiers | 213 |
| 9.30.3 | Assembler Directives | 214 |
| 9.30.4 | Floating Point | 214 |
| 9.31 | IBM S/390 Dependent Features | 215 |
| 9.31.1 | Options | 215 |
| 9.31.2 | Special Characters | 215 |
| 9.31.3 | Instruction syntax | 215 |
| 9.31.3.1 | Register naming | 216 |
| 9.31.3.2 | Instruction Mnemonics | 216 |
| 9.31.3.3 | Instruction Operands | 217 |
| 9.31.3.4 | Instruction Formats | 219 |
| 9.31.3.5 | Instruction Aliases | 222 |
| 9.31.3.6 | Instruction Operand Modifier | 225 |
| 9.31.3.7 | Instruction Marker | 227 |
| 9.31.3.8 | Literal Pool Entries | 227 |
| 9.31.4 | Assembler Directives | 228 |
| 9.31.5 | Floating Point | 229 |
| 9.32 | SCORE Dependent Features | 230 |
| 9.32.1 | Options | 230 |
| 9.32.2 | SCORE Assembler Directives | 230 |
| 9.33 | Renesas / SuperH SH Dependent Features | 232 |
| 9.33.1 | Options | 232 |
| 9.33.2 | Syntax | 232 |
| 9.33.2.1 | Special Characters | 232 |
| 9.33.2.2 | Register Names | 232 |
| 9.33.2.3 | Addressing Modes | 233 |
| 9.33.3 | Floating Point | 233 |
| 9.33.4 | SH Machine Directives | 234 |
| 9.33.5 | Opcodes | 234 |
| 9.34 | SuperH SH64 Dependent Features | 235 |
| 9.34.1 | Options | 235 |
| 9.34.2 | Syntax | 235 |
| 9.34.2.1 | Special Characters | 235 |
| 9.34.2.2 | Register Names | 236 |
| 9.34.2.3 | Addressing Modes | 236 |
| 9.34.3 | SH64 Machine Directives | 236 |
| 9.34.4 | Opcodes | 237 |
| 9.35 | SPARC Dependent Features | 238 |
| 9.35.1 | Options | 238 |
| 9.35.2 | Enforcing aligned data | 238 |
| 9.35.3 | Sparc Syntax | 239 |
| 9.35.3.1 | Special Characters | 239 |
| 9.35.3.2 | Register Names | 239 |
| 9.35.3.3 | Constants | 241 |
| 9.35.3.4 | Relocations | 242 |

| | | |
|----------|--|-----|
| 9.35.3.5 | Size Translations | 244 |
| 9.35.4 | Floating Point | 245 |
| 9.35.5 | Sparc Machine Directives | 245 |
| 9.36 | TIC54X Dependent Features | 247 |
| 9.36.1 | Options | 247 |
| 9.36.2 | Blocking | 247 |
| 9.36.3 | Environment Settings | 247 |
| 9.36.4 | Constants Syntax | 247 |
| 9.36.5 | String Substitution | 247 |
| 9.36.6 | Local Labels | 248 |
| 9.36.7 | Math Builtins | 248 |
| 9.36.8 | Extended Addressing | 250 |
| 9.36.9 | Directives | 250 |
| 9.36.10 | Macros | 255 |
| 9.36.11 | Memory-mapped Registers | 256 |
| 9.37 | TIC6X Dependent Features | 257 |
| 9.37.1 | TIC6X Options | 257 |
| 9.37.2 | TIC6X Syntax | 258 |
| 9.37.3 | TIC6X Directives | 258 |
| 9.38 | Z80 Dependent Features | 259 |
| 9.38.1 | Options | 259 |
| 9.38.2 | Syntax | 259 |
| 9.38.2.1 | Special Characters | 259 |
| 9.38.2.2 | Register Names | 260 |
| 9.38.2.3 | Case Sensitivity | 260 |
| 9.38.3 | Floating Point | 260 |
| 9.38.4 | Z80 Assembler Directives | 260 |
| 9.38.5 | Opcodes | 261 |
| 9.39 | Z8000 Dependent Features | 262 |
| 9.39.1 | Options | 262 |
| 9.39.2 | Syntax | 262 |
| 9.39.2.1 | Special Characters | 262 |
| 9.39.2.2 | Register Names | 262 |
| 9.39.2.3 | Addressing Modes | 262 |
| 9.39.3 | Assembler Directives for the Z8000 | 263 |
| 9.39.4 | Opcodes | 264 |
| 9.40 | VAX Dependent Features | 264 |
| 9.40.1 | VAX Command-Line Options | 264 |
| 9.40.2 | VAX Floating Point | 265 |
| 9.40.3 | Vax Machine Directives | 265 |
| 9.40.4 | VAX Opcodes | 265 |
| 9.40.5 | VAX Branch Improvement | 266 |
| 9.40.6 | VAX Operands | 267 |
| 9.40.7 | Not Supported on VAX | 268 |
| 9.41 | v850 Dependent Features | 268 |
| 9.41.1 | Options | 268 |
| 9.41.2 | Syntax | 269 |
| 9.41.2.1 | Special Characters | 269 |

| | | |
|-------------------|--|------------|
| 9.41.2.2 | Register Names | 269 |
| 9.41.3 | Floating Point | 271 |
| 9.41.4 | V850 Machine Directives | 271 |
| 9.41.5 | Opcodes | 272 |
| 9.42 | Xtensa Dependent Features | 275 |
| 9.42.1 | Command Line Options | 275 |
| 9.42.2 | Assembler Syntax | 276 |
| 9.42.2.1 | Opcode Names | 276 |
| 9.42.2.2 | Register Names | 277 |
| 9.42.3 | Xtensa Optimizations | 277 |
| 9.42.3.1 | Using Density Instructions | 277 |
| 9.42.3.2 | Automatic Instruction Alignment | 277 |
| 9.42.4 | Xtensa Relaxation | 278 |
| 9.42.4.1 | Conditional Branch Relaxation | 278 |
| 9.42.4.2 | Function Call Relaxation | 279 |
| 9.42.4.3 | Other Immediate Field Relaxation | 279 |
| 9.42.5 | Directives | 280 |
| 9.42.5.1 | schedule | 281 |
| 9.42.5.2 | longcalls | 281 |
| 9.42.5.3 | transform | 281 |
| 9.42.5.4 | literal | 281 |
| 9.42.5.5 | literal_position | 282 |
| 9.42.5.6 | literal_prefix | 283 |
| 9.42.5.7 | absolute-literals | 283 |
| 10 | Reporting Bugs | 285 |
| 10.1 | Have You Found a Bug? | 285 |
| 10.2 | How to Report Bugs | 285 |
| 11 | Acknowledgements | 289 |
| Appendix A | GNU Free Documentation License | |
| | | 291 |
| AS Index | | 299 |

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 \[Command-Line Options\]](#), [page 21](#).

```
as [-a[cdghlms][=file]] [-alternate] [-D]
  [-compress-debug-sections] [-nocompress-debug-sections]
  [-debug-prefix-map old=new]
  [-defsym sym=val] [-f] [-g] [-gstabs]
  [-gstabs+] [-gdwarf-2] [-help] [-I dir] [-J]
  [-K] [-L] [-listing-lhs-width=NUM]
  [-listing-lhs-width2=NUM] [-listing-rhs-width=NUM]
  [-listing-cont-lines=NUM] [-keep-locals] [-o
objfile] [-R] [-reduce-memory-overheads] [-statistics]
  [-v] [-version] [-version] [-W] [-warn]
  [-fatal-warnings] [-w] [-x] [-Z] [@FILE]
  [-size-check=[error|warning]]
  [-target-help] [target-options]
  [-|files ...]
```

Target Alpha options:

```
[-mcpu]
[-mdebug | -no-mdebug]
[-replace | -noreplace]
[-relax] [-g] [-Gsize]
[-F] [-32addr]
```

Target ARC options:

```
[-marc[5|6|7|8]]
[-EB|-EL]
```

Target ARM options:

```
[-mcpu=processor[+extension...]]
[-march=architecture[+extension...]]
[-mfpu=floating-point-format]
[-mfloat-abi=abi]
[-meabi=ver]
[-mthumb]
[-EB|-EL]
[-mapcs-32|-mapcs-26|-mapcs-float |
-mapcs-reentrant]
[-mthumb-interwork] [-k]
```

Target Blackfin options:

```
[-mcpu=processor[-sirevision]]
[-mfdpic]
[-mno-fdpic]
[-mnopic]
```

Target CRIS options:

```
[-underscore | -no-underscore]
[-pic] [-N]
[-emulation=criself | -emulation=crisaout]
[-march=v0_v10 | -march=v10 | -march=v32 | -march=common_v10_v32]
```

Target D10V options:

```
[-O]
```

Target D30V options:
 [-O|-n|-N]

Target H8/300 options:
 [-h-tick-hex]

Target i386 options:
 [-32|-64] [-n]
 [-march=CPU[+EXTENSION...]] [-mtune=CPU]

Target i960 options:
 [-ACA|-ACA_A|-ACB|-ACC|-AKA|-AKB|
 -AKC|-AMC]
 [-b] [-no-relax]

Target IA-64 options:
 [-mconstant-gp|-mauto-pic]
 [-milp32|-milp64|-mlp64|-mp64]
 [-mle|mbe]
 [-mtune=itanium1|-mtune=itanium2]
 [-munwind-check=warning|-munwind-check=error]
 [-mhint.b=ok|-mhint.b=warning|-mhint.b=error]
 [-x|-xexplicit] [-xauto] [-xdebug]

Target IP2K options:
 [-mip2022|-mip2022ext]

Target M32C options:
 [-m32c|-m16c] [-relax] [-h-tick-hex]

Target M32R options:
 [-m32rx|-no-warn-explicit-parallel-conflicts|
 -W[n]p]

Target M680X0 options:
 [-l] [-m68000|-m68010|-m68020|...]

Target M68HC11 options:
 [-m68hc11|-m68hc12|-m68hcs12]
 [-mshort|-mlong]
 [-mshort-double|-mlong-double]
 [-force-long-branches] [-short-branches]
 [-strict-direct-mode] [-print-insn-syntax]
 [-print-opcodes] [-generate-example]

Target MCORE options:
 [-jsri2bsr] [-sifilter] [-relax]
 [-mcpu=[210|340]]

Target MICROBLAZE options:

Target MIPS options:
 [-nocpp] [-EL] [-EB] [-O[optimization level]]
 [-g[debug level]] [-G num] [-KPIC] [-call.shared]
 [-non.shared] [-xgot [-mvxworks-pic]
 [-mabi=ABI] [-32] [-n32] [-64] [-mfp32] [-mfp32]
 [-march=CPU] [-mtune=CPU] [-mips1] [-mips2]
 [-mips3] [-mips4] [-mips5] [-mips32] [-mips32r2]


```

[-mips64] [-mips64r2]
[-construct-floats] [-no-construct-floats]
[-trap] [-no-break] [-break] [-no-trap]
[-mips16] [-no-mips16]
[-msmartmips] [-mno-smartmips]
[-mips3d] [-no-mips3d]
[-mdmx] [-no-mdmx]
[-mdsp] [-mno-dsp]
[-mdspr2] [-mno-dspr2]
[-mmt] [-mno-mt]
[-mfix7000] [-mno-fix7000]
[-mfix-vr4120] [-mno-fix-vr4120]
[-mfix-vr4130] [-mno-fix-vr4130]
[-mdebug] [-no-mdebug]
[-mpdr] [-mno-pdr]

```

Target MMIX options:

```

[-fixed-special-register-names] [-globalize-symbols]
[-gnu-syntax] [-relax] [-no-predefined-symbols]
[-no-expand] [-no-merge-gregs] [-x]
[-linker-allocated-gregs]

```

Target PDP11 options:

```

[-mpic|-mno-pic] [-mall] [-mno-extensions]
[-mextension|-mno-extension]
[-mcpu] [-mmachine]

```

Target picoJava options:

```

[-mb|-me]

```

Target PowerPC options:

```

[-a32|-a64]
[-mpwrx|-mpwr2|-mpwr|-m601|-mppc|-mppc32|-m603|-m604|-m403|-m405|
-m440|-m464|-m476|-m7400|-m7410|-m7450|-m7455|-m750cl|-mppc64|
-m620|-me500|-e500x2|-me500mc|-me500mc64|-mppc64bridge|-mbooke|
-mpower4|-mpr4|-mpower5|-mpwr5|-mpwr5x|-mpower6|-mpwr6|
-mpower7|-mpw7|-ma2|-mcell|-mspe|-mtitan|-me300|-mcom]
[-many] [-maltivec|-mvsx]
[-mregnames|-mno-regnames]
[-mrelocatable|-mrelocatable-lib|-K PIC] [-memb]
[-mlittle|-mlittle-endian|-le|-mbig|-mbig-endian|-be]
[-msolaris|-mno-solaris]
[-nops=count]

```

Target RX options:

```

[-mlittle-endian|-mbig-endian]
[-m32bit-ints|-m16bit-ints]
[-m32bit-doubles|-m64bit-doubles]

```

Target s390 options:

```

[-m31|-m64] [-mesa|-mzarch] [-march=CPU]
[-mregnames|-mno-regnames]
[-mwarn-areg-zero]

```

Target SCORE options:

```

[-EB] [-EL] [-FIXDD] [-NWARN]
[-SCORE5] [-SCORE5U] [-SCORE7] [-SCORE3]
[-march=score7] [-march=score3]

```

`[-USE_R1] [-KPIC] [-O0] [-G num] [-V]`

Target SPARC options:

`[-Av6|-Av7|-Av8|-Asparclet|-Asparclite
-Av8plus|-Av8plusa|-Av9|-Av9a]
[-xarch=v8plus|-xarch=v8plusa] [-bump]
[-32|-64]`

Target TIC54X options:

`[-mcpu=54[123589]|-mcpu=54[56]lp] [-mfarm-mode|-mf]
[-merrors-to-file <filename>|-me <filename>]`

Target TIC6X options:

`[-march=arch] [-matomic|-mno-atomic]
[-mbig-endian|-mlittle-endian] [-mdsbt|-mno-dsbt]
[-mpid=no|-mpid=near|-mpid=far] [-mpic|-mno-pic]`

Target Z80 options:

`[-z80] [-r800]
[-ignore-undocumented-instructions] [-Wnud]
[-ignore-unportable-instructions] [-Wnup]
[-warn-undocumented-instructions] [-Wud]
[-warn-unportable-instructions] [-Wup]
[-forbid-undocumented-instructions] [-Fud]
[-forbid-unportable-instructions] [-Fup]`

Target Xtensa options:

`[-[no-]text-section-literals] [-[no-]absolute-literals]
[-[no-]target-align] [-[no-]longcalls]
[-[no-]transform]
[-rename-section oldname=newname]`

@file Read command-line options from *file*. The options read are inserted in place of the original **@file** option. If *file* does not exist, or cannot be read, then the option will be treated literally, and not removed.

Options in *file* are separated by whitespace. A whitespace character may be included in an option by surrounding the entire option in either single or double quotes. Any character (including a backslash) may be included by prefixing the character to be included with a backslash. The *file* may itself contain additional **@file** options; any such options will be processed recursively.

-a[cdghlmns]

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

| | |
|------------|---|
| -ac | omit false conditionals |
| -ad | omit debugging directives |
| -ag | include general information, like as version and options passed |
| -ah | include high-level source |
| -al | 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'`.

--alternate

Begin in alternate macro mode. See [Section 7.4 \[.altmacro\]](#), page 48.

--compress-debug-sections

Compress DWARF debug sections using zlib. The debug sections are renamed to begin with `'.zdebug'`, and the resulting object file may not be compatible with older linkers and object file utilities.

--nocompress-debug-sections

Do not compress DWARF debug sections. This is the default.

-D Ignored. This option is accepted for script compatibility with calls to other assemblers.

--debug-prefix-map *old=new*

When assembling files in directory `'old'`, record debugging information describing them as in `'new'` instead.

--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. The value of the symbol can be overridden inside a source file via the use of a `.set` pseudo-op.

-f “fast”—skip whitespace and comment preprocessing (assume source is compiler output).

-g

--gen-debug

Generate debugging information for each assembler source line using whichever debug format is preferred by the target. This currently means either STABS, ECOFF or DWARF2.

--gstabs Generate stabs debugging information for each assembler line. This may help debugging assembler code, if the debugger can handle it.

--gstabs+

Generate stabs debugging information for each assembler line, with GNU extensions that probably only gdb can handle, and that could make other debuggers crash or refuse to read your program. This may help debugging assembler code. Currently the only GNU extension is the location of the current working directory at assembling time.

--gdwarf-2
Generate DWARF2 debugging information for each assembler line. This may help debugging assembler code, if the debugger can handle it. Note—this option is only supported by some targets, not all of them.

--size-check=error
--size-check=warning
Issue an error or warning for invalid ELF `.size` directive.

--help
Print a summary of the command line options and exit.

--target-help
Print a summary of all target specific 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. These symbols start with system-specific local label prefixes, typically `'.L'` for ELF systems or `'L'` for traditional a.out systems. See [Section 5.3 \[Symbol Names\]](#), page 39.

--listing-lhs-width=*number*
Set the maximum width, in words, of the output data column for an assembler listing to *number*.

--listing-lhs-width2=*number*
Set the maximum width, in words, of the output data column for continuation lines in an assembler listing to *number*.

--listing-rhs-width=*number*
Set the maximum width of an input source line, as displayed in a listing, to *number* bytes.

--listing-cont-lines=*number*
Set the maximum number of lines printed in a listing for a single line of input to *number* + 1.

-o *objfile*
Name the object-file output from `as` *objfile*.

-R
Fold the data section into the text section.
Set the default size of GAS's hash tables to a prime number close to *number*. Increasing this value can reduce the length of time it takes the assembler to perform its tasks, at the expense of increasing the assembler's memory requirements. Similarly reducing this value can reduce the memory requirements at the expense of speed.

--reduce-memory-overheads
This option reduces GAS's memory requirements, at the expense of making the assembly processes slower. Currently this switch is a synonym for `--hash-size=4051`, but in the future it may have other effects as well.

```

--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.

```

-marc[5|6|7|8]
    This option selects the core processor variant.

-EB | -EL   Select either big-endian (-EB) or little-endian (-EL) output.

```

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

```

-mcpu=processor[+extension...]
    Specify which ARM processor variant is the target.

-march=architecture[+extension...]
    Specify which ARM architecture variant is used by the target.

-mfpu=floating-point-format
    Select which Floating Point architecture is the target.

-mfloat-abi=abi
    Select which floating point ABI is in use.

-mthumb     Enable Thumb only instruction decoding.

-mapcs-32 | -mapcs-26 | -mapcs-float | -mapcs-reentrant
    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 the Blackfin processor family.

-mcpu=processor**[-sirevision]**

This option specifies the target processor. The optional *sirevision* is not used in assembler.

-mfdpic Assemble for the FDPIC ABI.

-mno-fdpic

-mnopic Disable -mfdpic.

See the info pages for documentation of the CRIS-specific options.

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 Uvicom IP2K series.

-mip2022ext

Specifies that the extended IP2022 instructions are allowed.

-mip2022 Restores the default behaviour, which restricts the permitted instructions to just the basic IP2022 ones.

The following options are available when as is configured for the Renesas M32C and M16C processors.

-m32c Assemble M32C instructions.

-m16c Assemble M16C instructions (the default).

-relax Enable support for link-time relaxations.

-h-tick-hex

Support H'00 style hex constants in addition to 0x00 style.

The following options are available when as is configured for the Renesas M32R (formerly Mitsubishi M32R) series.

--m32rx Specify which processor in the M32R family is the target. The default is normally the M32R, but this option changes it to the M32RX.

--warn-explicit-parallel-conflicts or --Wp

Produce warning messages when questionable parallel constructs are encountered.

--no-warn-explicit-parallel-conflicts or --Wnp

Do not produce warning messages when questionable parallel constructs are encountered.

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.

For details about the PDP-11 machine dependent features options, see [Section 9.27.1 \[PDP-11-Options\]](#), page 205.

-mpic | -mno-pic

Generate position-independent (or position-dependent) code. The default is '-mpic'.

-mall

-mall-extensions

Enable all instruction set extensions. This is the default.

-mno-extensions

Disable all instruction set extensions.

-mextension | -mno-extension

Enable (or disable) a particular instruction set extension.

-mcpu Enable the instruction set extensions supported by a particular CPU, and disable all other extensions.

-mmachine Enable the instruction set extensions supported by a particular machine model, and disable all other extensions.

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 Motorola 68HC11 or 68HC12 series.

-m68hc11 | -m68hc12 | -m68hcs12
Specify what processor is the target. The default is defined by the configuration option when building the assembler.

-mshort Specify to use the 16-bit integer ABI.

-mlong Specify to use the 32-bit integer ABI.

-mshort-double
Specify to use the 32-bit double ABI.

-mlong-double
Specify to use the 64-bit double ABI.

--force-long-branches
Relative branches are turned into absolute ones. This concerns conditional branches, unconditional branches and branches to a sub routine.

-S | --short-branches
Do not turn relative branches into absolute ones when the offset is out of range.

--strict-direct-mode
Do not turn the direct addressing mode into extended addressing mode when the instruction does not support direct addressing mode.

--print-insn-syntax
Print the syntax of instruction in case of error.

--print-opcodes
print the list of instructions with syntax and then exit.

--generate-example
print an example of instruction for each possible instruction and then exit. This option is only useful for testing `as`.

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 the ‘c54x architecture.

-mfar-mode

Enable extended addressing mode. All addresses and relocations will assume extended addressing (usually 23 bits).

-mcpu=CPU_VERSION

Sets the CPU version being compiled for.

-merrors-to-file FILENAME

Redirect error output to a file, for broken systems which don’t support such behaviour in the shell.

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

-mips4

-mips5

-mips32

-mips32r2

-mips64

-mips64r2

Generate code for a particular MIPS Instruction Set Architecture level. ‘-mips1’ is an alias for ‘-march=r3000’, ‘-mips2’ is an alias for ‘-march=r6000’, ‘-mips3’ is an alias for ‘-march=r4000’ and ‘-mips4’ is an alias for ‘-march=r8000’. ‘-mips5’, ‘-mips32’, ‘-mips32r2’, ‘-mips64’, and ‘-mips64r2’ correspond to generic ‘MIPS V’, ‘MIPS32’, ‘MIPS32 Release 2’, ‘MIPS64’, and ‘MIPS64 Release 2’ ISA processors, respectively.

-march=CPU

Generate code for a particular MIPS cpu.

-mtune=cpu

Schedule and tune for a particular MIPS cpu.

`-mfix7000`
`-mno-fix7000` Cause nops to be inserted if the read of the destination register of an `mfhi` or `mflo` instruction occurs in the following two instructions.

`-mdebug`
`-no-mdebug` Cause stabs-style debugging output to go into an ECOFF-style `.mdebug` section instead of the standard ELF `.stabs` sections.

`-mpdr`
`-mno-pdr` Control generation of `.pdr` sections.

`-mgp32`
`-mfp32` The register sizes are normally inferred from the ISA and ABI, but these flags force a certain group of registers to be treated as 32 bits wide at all times. ‘`-mgp32`’ controls the size of general-purpose registers and ‘`-mfp32`’ controls the size of floating-point registers.

`-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.

`-msmartmips`
`-mno-smartmips` Enables the SmartMIPS extension to the MIPS32 instruction set. This is equivalent to putting `.set smartmips` at the start of the assembly file. ‘`-mno-smartmips`’ turns off this option.

`-mips3d`
`-no-mips3d` Generate code for the MIPS-3D Application Specific Extension. This tells the assembler to accept MIPS-3D instructions. ‘`-no-mips3d`’ turns off this option.

`-mdmx`
`-no-mdmx` Generate code for the MDMX Application Specific Extension. This tells the assembler to accept MDMX instructions. ‘`-no-mdmx`’ turns off this option.

`-mdsp`
`-mno-dsp` Generate code for the DSP Release 1 Application Specific Extension. This tells the assembler to accept DSP Release 1 instructions. ‘`-mno-dsp`’ turns off this option.

`-mdspr2`
`-mno-dspr2` Generate code for the DSP Release 2 Application Specific Extension. This option implies `-mdsp`. This tells the assembler to accept DSP Release 2 instructions. ‘`-mno-dspr2`’ turns off this option.

`-mmt`
`-mno-mt` Generate code for the MT Application Specific Extension. This tells the assembler to accept MT instructions. ‘`-mno-mt`’ turns off this option.

`--construct-floats`

`--no-construct-floats`

The ‘`--no-construct-floats`’ option disables the construction of double width floating point constants by loading the two halves of the value into the two single width floating point registers that make up the double width register. By default ‘`--construct-floats`’ is selected, allowing construction of these floating point constants.

`--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.

`-n` When this option is used, `as` will issue a warning every time it generates a nop instruction from a macro.

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.
- `-mcpu=[210|340]`
 Select the cpu type on the target hardware. This controls which instructions can be assembled.
- `-EB` Assemble for a big endian target.
- `-EL` Assemble for a little endian target.

See the info pages for documentation of the MMIX-specific options.

See the info pages for documentation of the RX-specific options.

The following options are available when as is configured for the s390 processor family.

- `-m31`
- `-m64` Select the word size, either 31/32 bits or 64 bits.
- `-mesa`
- `-mzarch` Select the architecture mode, either the Enterprise System Architecture (esa) or the z/Architecture mode (zarch).
- `-march=processor`
 Specify which s390 processor variant is the target, ‘g6’, ‘g6’, ‘z900’, ‘z990’, ‘z9-109’, ‘z9-ec’, or ‘z10’.
- `-mregnames`
- `-mno-regnames`
 Allow or disallow symbolic names for registers.
- `-mwarn-areg-zero`
 Warn whenever the operand for a base or index register has been specified but evaluates to zero.

The following options are available when as is configured for a TMS320C6000 processor.

- `-march=arch`
 Enable (only) instructions from architecture *arch*. By default, all instructions are permitted.
 The following values of *arch* are accepted: c62x, c64x, c64x+, c67x, c67x+, c674x.
- `-matomic`
- `-mno-atomic`
 Enable or disable the optional C64x+ atomic operation instructions. By default, they are enabled if no ‘-march’ option is given, or if an architecture is specified with ‘-march’ that implies these instructions are present (currently, there are no such architectures); they are disabled if an architecture is specified with ‘-march’ on which the instructions are optional or not present. This option overrides such a default from the architecture, independent of the order in which the ‘-march’ or ‘-matomic’ or ‘-mno-atomic’ options are passed.

`-mdsbt`
`-mno-dsbt`

The ‘`-mdsbt`’ option causes the assembler to generate the `Tag_ABI_DSBT` attribute with a value of 1, indicating that the code is using DSBT addressing. The ‘`-mno-dsbt`’ option, the default, causes the tag to have a value of 0, indicating that the code does not use DSBT addressing. The linker will emit a warning if objects of different type (DSBT and non-DSBT) are linked together.

`-mpid=no`
`-mpid=near`
`-mpid=far`

The ‘`-mpid=`’ option causes the assembler to generate the `Tag_ABI_PID` attribute with a value indicating the form of data addressing used by the code. ‘`-mpid=no`’, the default, indicates position-dependent data addressing, ‘`-mpid=near`’ indicates position-independent addressing with GOT accesses using near DP addressing, and ‘`-mpid=far`’ indicates position-independent addressing with GOT accesses using far DP addressing. The linker will emit a warning if objects built with different settings of this option are linked together.

`-mpic`
`-mno-pic`

The ‘`-mpic`’ option causes the assembler to generate the `Tag_ABI_PIC` attribute with a value of 1, indicating that the code is using position-independent code addressing. The `-mno-pic` option, the default, causes the tag to have a value of 0, indicating position-dependent code addressing. The linker will emit a warning if objects of different type (position-dependent and position-independent) are linked together.

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

Generate code for the specified endianness. The default is little-endian.

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

`--text-section-literals` | `--no-text-section-literals`

With ‘`--text-section-literals`’, literal pools are interspersed in the text section. The default is ‘`--no-text-section-literals`’, which places literals in a separate section in the output file. These options only affect literals referenced via PC-relative L32R instructions; literals for absolute mode L32R instructions are handled separately.

`--absolute-literals` | `--no-absolute-literals`

Indicate to the assembler whether L32R instructions use absolute or PC-relative addressing. The default is to assume absolute addressing if the Xtensa processor includes the absolute L32R addressing option. Otherwise, only the PC-relative L32R mode can be used.

`--target-align` | `--no-target-align`

Enable or disable automatic alignment to reduce branch penalties at the expense of some code density. The default is ‘`--target-align`’.

--longcalls | --no-longcalls
 Enable or disable transformation of call instructions to allow calls across a greater range of addresses. The default is ‘--no-longcalls’.

--transform | --no-transform
 Enable or disable all assembler transformations of Xtensa instructions. The default is ‘--transform’; ‘--no-transform’ should be used only in the rare cases when the instructions must be exactly as specified in the assembly source.

--rename-section *oldname=newname*
 When generating output sections, rename the *oldname* section to *newname*.

The following options are available when as is configured for a Z80 family processor.

-z80 Assemble for Z80 processor.

-r800 Assemble for R800 processor.

-ignore-undocumented-instructions
-Wnud Assemble undocumented Z80 instructions that also work on R800 without warning.

-ignore-unportable-instructions
-Wnup Assemble all undocumented Z80 instructions without warning.

-warn-undocumented-instructions
-Wud Issue a warning for undocumented Z80 instructions that also work on R800.

-warn-unportable-instructions
-Wup Issue a warning for undocumented Z80 instructions that do not work on R800.

-forbid-undocumented-instructions
-Fud Treat all undocumented instructions as errors.

-forbid-unportable-instructions
-Fup Treat undocumented Z80 instructions that do not work on R800 as errors.

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 9 \[Machine Dependencies\]](#), page 81). 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 7.82 \[.org\]](#), page 64).

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 41.

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 CTL-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.

Filenames 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 18.

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 7.51 \[.file\]](#), page 54.

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`. 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 7.51 \[.file\]](#), page 54) 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 7.69 \[.line\]](#), page 59) 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 9 \[Machine Dependencies\]](#), page 81, for options specific to particular machine architectures.

If you are invoking `as` via the GNU C compiler, 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[cdghlns]’`

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 `‘-ag’` option to print a first section with general assembly information, like as version, switches passed, or time stamp.

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.*, `‘-aln’`.

Note if the assembler source is coming from the standard input (*e.g.*, because it is being created by `gcc` and the `‘-pipe’` command line switch is being used) then the listing will not contain any comments or preprocessor directives. This is because the listing code buffers input source lines from `stdin` only after they have been preprocessed by the assembler. This reduces memory usage and makes the code more efficient.

2.2 `‘--alternate’`

Begin in alternate macro mode, see [Section 7.4 \[.altmacro\]](#), page 48.

2.3 ‘-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`.

2.4 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 27](#).

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

2.5 .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 7.62 \[.include\], page 57](#)). 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.

2.6 Difference Tables: ‘-K’

`as` sometimes alters the code emitted for directives of the form ‘.word *sym1-sym2*’. See [Section 7.121 \[.word\], page 77](#). You can use the ‘-K’ option if you want a warning issued when this is done.

2.7 Include Local Symbols: ‘-L’

Symbols beginning with system-specific local label prefixes, typically ‘.L’ for ELF systems or ‘L’ for traditional a.out systems, are called *local symbols*. See [Section 5.3 \[Symbol Names\], page 39](#). Normally you do not see such symbols 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 symbols, so you do not normally debug with them.

This option tells `as` to retain those local symbols in the object file. Usually if you do this you also tell the linker `ld` to preserve those symbols.

2.8 Configuring listing output: ‘--listing’

The listing feature of the assembler can be enabled via the command line switch ‘-a’ (see [Section 2.1 \[a\], page 21](#)). This feature combines the input source file(s) with a hex dump of the corresponding locations in the output object file, and displays them as a listing file. The format of this listing can be controlled by directives inside the assembler source (i.e., `.list` (see [Section 7.71 \[List\], page 60](#)), `.title` (see [Section 7.111 \[Title\], page 74](#)), `.sbttl` (see [Section 7.94 \[Sbttl\], page 67](#)), `.psize` (see [Section 7.88 \[Psize\], page 66](#)), and `.eject` (see [Section 7.36 \[Eject\], page 52](#)) and also by the following switches:

`--listing-lhs-width='number'`

Sets the maximum width, in words, of the first line of the hex byte dump. This dump appears on the left hand side of the listing output.

`--listing-lhs-width2='number'`

Sets the maximum width, in words, of any further lines of the hex byte dump for a given input source line. If this value is not specified, it defaults to being the same as the value specified for `'--listing-lhs-width'`. If neither switch is used the default is to one.

`--listing-rhs-width='number'`

Sets the maximum width, in characters, of the source line that is displayed alongside the hex dump. The default value for this parameter is 100. The source line is displayed on the right hand side of the listing output.

`--listing-cont-lines='number'`

Sets the maximum number of continuation lines of hex dump that will be displayed for a given single line of source input. The default value is 4.

2.9 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.

- `N00BJ` pseudo-op

The m68k `N00BJ` 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.

2.10 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.

2.11 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.

2.12 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 33](#).)

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 or ELF 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`.

2.13 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).

2.14 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.

2.15 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.

2.16 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.

2.17 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 7.62 \[.include\]](#), page 57). 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 29), 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 ‘;’ on the ARC; ‘@’ on the ARM; ‘;’ for the H8/300 family; ‘;’ for the HPPA; ‘#’ on the i386 and x86-64; ‘#’ on the i960; ‘;’ for the PDP-11; ‘;’ for picoJava; ‘#’ for Motorola PowerPC; ‘#’ for IBM S/390; ‘#’ for the Sunplus SCORE; ‘!’ for the Renesas / SuperH SH; ‘!’ on the SPARC; ‘#’ on the ip2k; ‘#’ on the m32c; ‘#’ on the m32r; ‘|’ on the 680x0; ‘#’ on the 68HC11 and 68HC12; ‘#’ on the RX; ‘;’ on the TMS320C6X; ‘#’ on the Vax; ‘;’ for the Z80; ‘!’ for the Z8000; ‘#’ on the V850; ‘#’ for Xtensa systems; see [Chapter 9 \[Machine Dependencies\]](#), page 81.

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 43): the logical line number of the *next* line. Then a string (see [Section 3.6.1.1 \[Strings\]](#), page 29) 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.)

```

                                # This is an ordinary comment.
# 42-6 "new_file_name"         # 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 9 \[Machine Dependencies\]](#), page 81. 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 39.

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 9 \[Machine Dependencies\]](#), page 81.) 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 39](#).

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 43).

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, Renesas / SuperH 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 for the Renesas / SuperH 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 or ELF output, `as` can also generate whatever other named sections you specify using the `‘.section’` directive (see [Section 7.96 \[‘.section’\]](#), page 68). 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
 $(\text{address}) - (\text{start-address of section})$?

- 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 of 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`. When generating ELF output, you can also use the `.subsection` directive (see [Section 7.107 \[SubSection\], page 73](#)) to specify a subsection: `.subsection expression`. *Expression* should be an absolute expression (see [Chapter 6 \[Expressions\], page 43](#)). 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 7.67 \[.lcomm\], page 59](#).

The `.comm` pseudo-op may be used to declare a common symbol, which is another form of uninitialized symbol; see [Section 7.30 \[.comm\], page 51](#).

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 7.96](#) [`.section`], page 68. You may only assemble zero values into the section. Typically the section will only contain symbol definitions and `.skip` directives (see [Section 7.101](#) [`.skip`], page 71).

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 43). This is equivalent to using the `.set` directive. See [Section 7.97 \[.set\]](#), page 70. In the same way, using a double equals sign `==` here represents an equivalent of the `.eqv` directive. See [Section 7.45 \[.eqv\]](#), page 53.

Blackfin does not support symbol assignment with `=`.

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 9 \[Machine Dependencies\]](#), page 81. That character may be followed by any string of digits, letters, dollar signs (unless otherwise noted for a particular target machine), and underscores.

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

A local symbol is any symbol beginning with certain local label prefixes. By default, the local label prefix is `._L` for ELF systems or `._L` for traditional a.out systems, but each target may have its own set of local label prefixes. On the HPPA local symbols begin with `._L$`.

Local symbols are defined and used within the assembler, but they are normally not saved in object files. Thus, they are not visible when debugging. You may use the `-L` option (see [Section 2.7 \[Include Local Symbols: '-L'\]](#), page 22) to retain the local symbols in the object files.

Local Labels

Local labels help compilers and programmers use names temporarily. They create symbols which are guaranteed to be unique over the entire scope of the input source code and which can be referred to by a simple notation. To define a local label, write a label of the form ‘**N:**’ (where **N** represents any positive integer). To refer to the most recent previous definition of that label write ‘**Nb**’, using the same number as when you defined the label. To refer to the next definition of a local label, write ‘**Nf**’—the ‘**b**’ stands for “backwards” and the ‘**f**’ stands for “forwards”.

There is no restriction on how you can use these labels, and you can reuse them too. So that it is possible to repeatedly define the same local label (using the same number ‘**N**’), although you can only refer to the most recently defined local label of that number (for a backwards reference) or the next definition of a specific local label for a forward reference. It is also worth noting that the first 10 local labels (‘**0:**’ . . . ‘**9:**’) are implemented in a slightly more efficient manner than the others.

Here is an example:

```
1:      branch 1f
2:      branch 1b
1:      branch 2f
2:      branch 1b
```

Which is the equivalent of:

```
label_1: branch label_3
label_2: branch label_1
label_3: branch label_4
label_4: branch label_3
```

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

local label prefix

All local symbols begin with the system-specific local label prefix. Normally both **as** and **ld** forget symbols that start with the local label prefix. 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.

number This is the number that was used in the local label definition. So if the label is written ‘**55:**’ then the number is ‘**55**’.

C-B This unusual character is included so you do not accidentally invent a symbol of the same name. The character has ASCII value of ‘\002’ (control-B).

ordinal number

This is a serial number to keep the labels distinct. The first definition of ‘**0:**’ gets the number ‘**1**’. The 15th definition of ‘**0:**’ gets the number ‘**15**’, and so on. Likewise the first definition of ‘**1:**’ gets the number ‘**1**’ and its 15th definition gets ‘**15**’ as well.

So for example, the first **1:** may be named **.L1C-B1**, and the 44th **3:** may be named **.L3C-B44**.

Dollar Local Labels

as also supports an even more local form of local labels called dollar labels. These labels go out of scope (i.e., they become undefined) as soon as a non-local label is defined. Thus they remain valid for only a small region of the input source code. Normal local labels, by contrast, remain in scope for the entire file, or until they are redefined by another occurrence of the same local label.

Dollar labels are defined in exactly the same way as ordinary local labels, except that they have a dollar sign suffix to their numeric value, e.g., `'55$'`.

They can also be distinguished from ordinary local labels by their transformed names which use ASCII character `'\001'` (control-A) as the magic character to distinguish them from ordinary labels. For example, the fifth definition of `'6$'` may be named `'L6C-A5'`.

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 7.33 \[.desc\]](#), page 52). 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`, `.tag`, and `.weak` 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

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

2. Intermediate precedence

| | |
|---|------------------------------|
| | <i>Bitwise Inclusive Or.</i> |
| & | <i>Bitwise And.</i> |
| ^ | <i>Bitwise Exclusive Or.</i> |
| ! | <i>Bitwise Or Not.</i> |

3. Low Precedence

| | |
|----|--|
| + | <i>Addition.</i> If either argument is absolute, the result has the section of the other argument. You may not add together arguments from different sections. |
| - | <i>Subtraction.</i> 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. |
| == | <i>Is Equal To</i> |
| <> | |
| != | <i>Is Not Equal To</i> |
| < | <i>Is Less Than</i> |
| > | <i>Is Greater Than</i> |
| >= | <i>Is Greater Than Or Equal To</i> |
| <= | <i>Is Less Than Or Equal To</i> |

The comparison operators can be used as infix operators. A true results has a value of -1 whereas a false result has a value of 0. Note, these operators perform signed comparisons.

4. Lowest Precedence

| | |
|----|---------------------|
| && | <i>Logical And.</i> |
|----|---------------------|

|| *Logical Or.*

These two logical operations can be used to combine the results of sub expressions. Note, unlike the comparison operators a true result returns a value of 1 but a false results does still return 0. Also note that the logical or operator has a slightly lower precedence than logical and.

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 9 \[Machine Dependencies\]](#), page 81.

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 (COFF)`

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

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 arc, hppa, i386 using ELF, i860, iq2000, m68k, or32, s390, sparc, tic4x, tic80 and xtensa, 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 the tic54x, the first expression is the alignment request in words.

For other systems, including ppc, i386 using a.out format, arm and strongarm, 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 .altmacro

Enable alternate macro mode, enabling:

LOCAL *name* [, ...]

One additional directive, **LOCAL**, is available. It is used to generate a string replacement for each of the *name* arguments, and replace any instances of *name* in each macro expansion. The replacement string is unique in the assembly, and different for each separate macro expansion. **LOCAL** allows you to write macros that define symbols, without fear of conflict between separate macro expansions.

String delimiters

You can write strings delimited in these other ways besides "*string*":

'*string*' You can delimit strings with single-quote characters.

<*string*> You can delimit strings with matching angle brackets.

single-character string escape

To include any single character literally in a string (even if the character would otherwise have some special meaning), you can prefix the character with '!' (an exclamation mark). For example, you can write '<4.3 !> 5.4!!>' to get the literal text '4.3 > 5.4!'.

Expression results as strings

You can write '%*expr*' to evaluate the expression *expr* and use the result as a string.

7.5 .ascii "*string*"...

.ascii expects zero or more string literals (see [Section 3.6.1.1 \[Strings\], page 29](#)) separated by commas. It assembles each string (with no automatic trailing zero byte) into consecutive addresses.

7.6 .asciz "*string*"...

.asciz is just like **.ascii**, but each string is followed by a zero byte. The "z" in '**.asciz**' stands for "zero".

7.7 .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.

7.8 `.byte expressions`

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

7.9 `.cfi_sections section_list`

`.cfi_sections` may be used to specify whether CFI directives should emit `.eh_frame` section and/or `.debug_frame` section. If `section_list` is `.eh_frame`, `.eh_frame` is emitted, if `section_list` is `.debug_frame`, `.debug_frame` is emitted. To emit both use `.eh_frame, .debug_frame`. The default if this directive is not used is `.cfi_sections .eh_frame`.

7.10 `.cfi_startproc [simple]`

`.cfi_startproc` is used at the beginning of each function that should have an entry in `.eh_frame`. It initializes some internal data structures. Don't forget to close the function by `.cfi_endproc`.

Unless `.cfi_startproc` is used along with parameter `simple` it also emits some architecture dependent initial CFI instructions.

7.11 `.cfi_endproc`

`.cfi_endproc` is used at the end of a function where it closes its unwind entry previously opened by `.cfi_startproc`, and emits it to `.eh_frame`.

7.12 `.cfi_personality encoding [, exp]`

`.cfi_personality` defines personality routine and its encoding. `encoding` must be a constant determining how the personality should be encoded. If it is 255 (`DW_EH_PE_omit`), second argument is not present, otherwise second argument should be a constant or a symbol name. When using indirect encodings, the symbol provided should be the location where personality can be loaded from, not the personality routine itself. The default after `.cfi_startproc` is `.cfi_personality 0xff`, no personality routine.

7.13 `.cfi_extern encoding [, exp]`

`.cfi_extern` defines LSDA and its encoding. `encoding` must be a constant determining how the LSDA should be encoded. If it is 255 (`DW_EH_PE_omit`), second argument is not present, otherwise second argument should be a constant or a symbol name. The default after `.cfi_startproc` is `.cfi_extern 0xff`, no LSDA.

7.14 `.cfi_def_cfa register, offset`

`.cfi_def_cfa` defines a rule for computing CFA as: *take address from register and add offset to it.*

7.15 `.cfi_def_cfa_register register`

`.cfi_def_cfa_register` modifies a rule for computing CFA. From now on *register* will be used instead of the old one. Offset remains the same.

7.16 `.cfi_def_cfa_offset offset`

`.cfi_def_cfa_offset` modifies a rule for computing CFA. Register remains the same, but *offset* is new. Note that it is the absolute offset that will be added to a defined register to compute CFA address.

7.17 `.cfi_adjust_cfa_offset offset`

Same as `.cfi_def_cfa_offset` but *offset* is a relative value that is added/subtracted from the previous offset.

7.18 `.cfi_offset register, offset`

Previous value of *register* is saved at offset *offset* from CFA.

7.19 `.cfi_rel_offset register, offset`

Previous value of *register* is saved at offset *offset* from the current CFA register. This is transformed to `.cfi_offset` using the known displacement of the CFA register from the CFA. This is often easier to use, because the number will match the code it's annotating.

7.20 `.cfi_register register1, register2`

Previous value of *register1* is saved in register *register2*.

7.21 `.cfi_restore register`

`.cfi_restore` says that the rule for *register* is now the same as it was at the beginning of the function, after all initial instruction added by `.cfi_startproc` were executed.

7.22 `.cfi_undefined register`

From now on the previous value of *register* can't be restored anymore.

7.23 `.cfi_same_value register`

Current value of *register* is the same like in the previous frame, i.e. no restoration needed.

7.24 `.cfi_remember_state,`

First save all current rules for all registers by `.cfi_remember_state`, then totally screw them up by subsequent `.cfi_*` directives and when everything is hopelessly bad, use `.cfi_restore_state` to restore the previous saved state.

7.25 `.cfi_return_column register`

Change return column *register*, i.e. the return address is either directly in *register* or can be accessed by rules for *register*.

7.26 `.cfi_signal_frame`

Mark current function as signal trampoline.

7.27 `.cfi_window_save`

SPARC register window has been saved.

7.28 `.cfi_escape expression[, ...]`

Allows the user to add arbitrary bytes to the unwind info. One might use this to add OS-specific CFI opcodes, or generic CFI opcodes that GAS does not yet support.

7.29 `.cfi_val_encoded_addr register, encoding, label`

The current value of *register* is *label*. The value of *label* will be encoded in the output file according to *encoding*; see the description of `.cfi_personality` for details on this encoding.

The usefulness of equating a register to a fixed label is probably limited to the return address register. Here, it can be useful to mark a code segment that has only one return address which is reached by a direct branch and no copy of the return address exists in memory or another register.

7.30 `.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 or (as a GNU extension) PE, the `.comm` directive takes an optional third argument. This is the desired alignment of the symbol, specified for ELF as a byte boundary (for example, an alignment of 16 means that the least significant 4 bits of the address should be zero), and for PE as a power of two (for example, an alignment of 5 means aligned to a 32-byte boundary). 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 on ELF, or the default section alignment of 4 on PE¹.

¹ This is not the same as the executable image file alignment controlled by `ld`'s `--section-alignment` option; image file sections in PE are aligned to multiples of 4096, which is far too large an alignment for ordinary variables. It is rather the default alignment for (non-debug) sections within object (`*.o`) files, which are less strictly aligned.

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

7.31 `.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.

7.32 `.def name`

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

7.33 `.desc symbol, abs-expression`

This directive sets the descriptor of the symbol (see [Section 5.5 \[Symbol Attributes\]](#), page 41) 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.

7.34 `.dim`

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

7.35 `.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 9 \[Machine Dependencies\]](#), page 81.

7.36 `.eject`

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

7.37 `.else`

`.else` is part of the `as` support for conditional assembly; see [Section 7.60 \[.if\]](#), page 56. It marks the beginning of a section of code to be assembled if the condition for the preceding `.if` was false.

7.38 `.elseif`

`.elseif` is part of the `as` support for conditional assembly; see [Section 7.60 \[.if\]](#), page 56. It is shorthand for beginning a new `.if` block that would otherwise fill the entire `.else` section.

7.39 .end

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

7.40 .endef

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

7.41 .endfunc

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

7.42 .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 7.60 \[.if\]](#), page 56.

7.43 .equ *symbol*, *expression*

This directive sets the value of *symbol* to *expression*. It is synonymous with `.set`; see [Section 7.97 \[.set\]](#), page 70.

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

The syntax for `equ` on the Z80 is `'symbol equ expression'`. On the Z80 it is an error if *symbol* is already defined, but the symbol is not protected from later redefinition. Compare [Section 7.44 \[Equiv\]](#), page 53.

7.44 .equiv *symbol*, *expression*

The `.equiv` directive is like `.equ` and `.set`, except that the assembler will signal an error if *symbol* is already defined. Note a symbol which has been referenced but not actually defined is considered to be undefined.

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

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

plus it protects the symbol from later redefinition.

7.45 .eqv *symbol*, *expression*

The `.eqv` directive is like `.equiv`, but no attempt is made to evaluate the expression or any part of it immediately. Instead each time the resulting symbol is used in an expression, a snapshot of its current value is taken.

7.46 .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 an error in conditionally compiled code.

7.47 `.error "string"`

Similarly to `.err`, this directive emits an error, but you can specify a string that will be emitted as the error message. If you don't specify the message, it defaults to `".error directive invoked in source file"`. See [Section 1.7 \[Error and Warning Messages\]](#), page 18.

```
.error "This code has not been assembled and tested."
```

7.48 `.exitm`

Exit early from the current macro definition. See [Section 7.77 \[Macro\]](#), page 61.

7.49 `.extern`

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

7.50 `.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.

7.51 `.file`

There are two different versions of the `.file` directive. Targets that support DWARF2 line number information use the DWARF2 version of `.file`. Other targets use the default version.

Default Version

This version of the `.file` directive tells `as` that we are about to start a new logical file. The syntax is:

```
.file string
```

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.

DWARF2 Version

When emitting DWARF2 line number information, `.file` assigns filenames to the `.debug_line` file name table. The syntax is:

```
.file fileno filename
```

The *fileno* operand should be a unique positive integer to use as the index of the entry in the table. The *filename* operand is a C string literal.

The detail of filename indices is exposed to the user because the filename table is shared with the `.debug_info` section of the DWARF2 debugging information, and thus the user must know the exact indices that table entries will have.

7.52 `.fill repeat , size , value`

`repeat`, `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.

7.53 `.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 9 \[Machine Dependencies\]](#), page 81.

7.54 `.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`.

7.55 `.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 9.11.5 \[HPPA Assembler Directives\]](#), page 133.

7.56 `.gnu_attribute tag,value`

Record a GNU object attribute for this file. See [Chapter 8 \[Object Attributes\]](#), page 79.

7.57 `.hidden names`

This is one of the ELF visibility directives. The other two are `.internal` (see [Section 7.64 \[.internal\]](#), page 58) and `.protected` (see [Section 7.87 \[.protected\]](#), page 66).

This directive overrides the named symbols default visibility (which is set by their binding: local, global or weak). The directive sets the visibility to `hidden` which means that

the symbols are not visible to other components. Such symbols are always considered to be `protected` as well.

7.58 `.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`’.

7.59 `.ident`

This directive is used by some assemblers to place tags in object files. The behavior of this directive varies depending on the target. When using the a.out object file format, `as` simply accepts the directive for source-file compatibility with existing assemblers, but does not emit anything for it. When using COFF, comments are emitted to the `.comment` or `.rdata` section, depending on the target. When using ELF, comments are emitted to the `.comment` section.

7.60 `.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 7.42 \[.endif\], page 53](#)); optionally, you may include code for the alternative condition, flagged by `.else` (see [Section 7.37 \[.else\], page 52](#)). If you have several conditions to check, `.elseif` may be used to avoid nesting blocks if/else within each subsequent `.else` block.

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

`.ifdef symbol`

Assembles the following section of code if the specified *symbol* has been defined. Note a symbol which has been referenced but not yet defined is considered to be undefined.

`.ifb text`

Assembles the following section of code if the operand is blank (empty).

`.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.

.ifnb *text*
Like **.ifb**, but the sense of the test is reversed: this assembles the following section of code if the operand is non-blank (non-empty).

.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. Note a symbol which has been referenced but not yet defined is considered to be undefined.

.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.

7.61 **.incbin** "*file*"[,*skip*[,*count*]]

The **incbin** directive includes *file* verbatim at the current location. You can control the search paths used with the **-I** command-line option (see [Chapter 2 \[Command-Line Options\]](#), page 21). Quotation marks are required around *file*.

The *skip* argument skips a number of bytes from the start of the *file*. The *count* argument indicates the maximum number of bytes to read. Note that the data is not aligned in any way, so it is the user's responsibility to make sure that proper alignment is provided both before and after the **incbin** directive.

7.62 **.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 21). Quotation marks are required around *file*.

7.63 `.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.

7.64 `.internal names`

This is one of the ELF visibility directives. The other two are `.hidden` (see [Section 7.57 \[.hidden\]](#), page 55) and `.protected` (see [Section 7.87 \[.protected\]](#), page 66).

This directive overrides the named symbols default visibility (which is set by their binding: local, global or weak). The directive sets the visibility to **internal** which means that the symbols are considered to be **hidden** (i.e., not visible to other components), and that some extra, processor specific processing must also be performed upon the symbols as well.

7.65 `.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@-
```

For some caveats with the spelling of *symbol*, see also [Section 7.77 \[Macro\]](#), page 61.

7.66 `.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@-
```

For some caveats with the spelling of *symbol*, see also the discussion at See [Section 7.77 \[Macro\]](#), page 61.

7.67 .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 7.55 \[.global\]](#), page 55), 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.

7.68 .lflags

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

7.69 .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.

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.

7.70 .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.

7.71 .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 21), the initial value of the listing counter is one.

7.72 .ln *line-number*

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

7.73 .loc *fileno lineno [column] [options]*

When emitting DWARF2 line number information, the `.loc` directive will add a row to the `.debug_line` line number matrix corresponding to the immediately following assembly instruction. The *fileno*, *lineno*, and optional *column* arguments will be applied to the `.debug_line` state machine before the row is added.

The *options* are a sequence of the following tokens in any order:

basic_block
This option will set the `basic_block` register in the `.debug_line` state machine to `true`.

prologue_end
This option will set the `prologue_end` register in the `.debug_line` state machine to `true`.

epilogue_begin
This option will set the `epilogue_begin` register in the `.debug_line` state machine to `true`.

is_stmt value
This option will set the `is_stmt` register in the `.debug_line` state machine to *value*, which must be either 0 or 1.

isa value
This directive will set the `isa` register in the `.debug_line` state machine to *value*, which must be an unsigned integer.

discriminator value

This directive will set the `discriminator` register in the `.debug_line` state machine to *value*, which must be an unsigned integer.

7.74 .loc_mark_labels enable

When emitting DWARF2 line number information, the `.loc_mark_labels` directive makes the assembler emit an entry to the `.debug_line` line number matrix with the `basic_block` register in the state machine set whenever a code label is seen. The *enable* argument should be either 1 or 0, to enable or disable this function respectively.

7.75 .local names

This directive, which is available for ELF targets, marks each symbol in the comma-separated list of *names* as a local symbol so that it will not be externally visible. If the symbols do not already exist, they will be created.

For targets where the `.lcomm` directive (see [Section 7.67 \[Lcomm\], page 59](#)) does not accept an alignment argument, which is the case for most ELF targets, the `.local` directive can be used in combination with `.comm` (see [Section 7.30 \[Comm\], page 51](#)) to define aligned local common data.

7.76 .long expressions

`.long` is the same as `‘.int’`. See [Section 7.63 \[.int\], page 58](#).

7.77 .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 qualify the macro argument to indicate whether all

invocations must specify a non-blank value (through ‘:req’), or whether it takes all of the remaining arguments (through ‘:vararg’). You can supply a default value for any macro argument by following the name with ‘=deflt’. You cannot define two macros with the same *macname* unless it has been subject to the `.purgem` directive (see [Section 7.89 \[Purgem\]](#), page 66) between the two definitions. 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*).

```
.macro m p1:req, p2=0, p3:vararg
```

Begin the definition of a macro called `m`, with at least three arguments. The first argument must always have a value specified, but not the second, which instead has a default value. The third formal will get assigned all remaining arguments specified at invocation time.

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`’.

Note that since each of the *macargs* can be an identifier exactly as any other one permitted by the target architecture, there may be occasional problems if the target hand-crafts special meanings to certain characters when they occur in a special position. For example, if the colon (:) is generally permitted to be part of a symbol name, but the architecture specific code special-cases it when occurring as the final character of a symbol (to denote a label), then the macro parameter replacement code will have no way of knowing that and consider the whole construct (including the colon) an identifier, and check only this identifier for being the subject to parameter substitution. So for example this macro definition:

```
.macro label l
```

```
\l:
```

```
.endm
```

might not work as expected. Invoking ‘`label foo`’ might not create a label called ‘`foo`’ but instead just insert the text ‘`\l:`’ into the assembler source, probably generating an error about an unrecognised identifier.

Similarly problems might occur with the period character (‘`.`’) which is often allowed inside opcode names (and hence identifier names). So for example constructing a macro to build an opcode from a base name and a length specifier like this:

```
.macro opcode base length
    \base.\length
.endm
```

and invoking it as ‘`opcode store 1`’ will not create a ‘`store.1`’ instruction but instead generate some kind of error as the assembler tries to interpret the text ‘`\base.\length`’.

There are several possible ways around this problem:

Insert white space

If it is possible to use white space characters then this is the simplest solution. eg:

```
.macro label l
    \l :
.endm
```

Use ‘`\()`’ The string ‘`\()`’ can be used to separate the end of a macro argument from the following text. eg:

```
.macro opcode base length
    \base\().\length
.endm
```

Use the alternate macro syntax mode

In the alternative macro syntax mode the ampersand character (‘`&`’) can be used as a separator. eg:

```
.altmacro
.macro label l
    l&:
.endm
```

Note: this problem of correctly identifying string parameters to pseudo ops also applies to the identifiers used in `.irp` (see [Section 7.65 \[Irp\]](#), page 58) and `.irpc` (see [Section 7.66 \[Irpc\]](#), page 58) as well.

`.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*.

`LOCAL name [, ...]`

Warning: `LOCAL` is only available if you select “alternate macro syntax” with ‘`--alternate`’ or `.altmacro`. See [Section 7.4 \[.altmacro\]](#), page 48.

7.78 `.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 2.9 \[MRI mode\]](#), page 23.

7.79 `.noaltmacro`

Disable alternate macro mode. See [Section 7.4 \[Altmacro\]](#), page 48.

7.80 `.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.

7.81 `.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.

7.82 `.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.

7.83 `.p2align[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 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.

7.84 .popsection

This is one of the ELF section stack manipulation directives. The others are `.section` (see [Section 7.96 \[Section\], page 68](#)), `.subsection` (see [Section 7.107 \[SubSection\], page 73](#)), `.pushsection` (see [Section 7.90 \[PushSection\], page 66](#)), and `.previous` (see [Section 7.85 \[Previous\], page 65](#)).

This directive replaces the current section (and subsection) with the top section (and subsection) on the section stack. This section is popped off the stack.

7.85 .previous

This is one of the ELF section stack manipulation directives. The others are `.section` (see [Section 7.96 \[Section\], page 68](#)), `.subsection` (see [Section 7.107 \[SubSection\], page 73](#)), `.pushsection` (see [Section 7.90 \[PushSection\], page 66](#)), and `.popsection` (see [Section 7.84 \[PopSection\], page 65](#)).

This directive swaps the current section (and subsection) with most recently referenced section/subsection pair prior to this one. Multiple `.previous` directives in a row will flip between two sections (and their subsections). For example:

```
.section A
.subsection 1
.word 0x1234
.subsection 2
.word 0x5678
.previous
.word 0x9abc
```

Will place 0x1234 and 0x9abc into subsection 1 and 0x5678 into subsection 2 of section A. Whilst:

```
.section A
.subsection 1
# Now in section A subsection 1
.word 0x1234
.section B
```

```
.subsection 0
# Now in section B subsection 0
.word 0x5678
.subsection 1
# Now in section B subsection 1
.word 0x9abc
.previous
# Now in section B subsection 0
.word 0xdef0
```

Will place 0x1234 into section A, 0x5678 and 0xdef0 into subsection 0 of section B and 0x9abc into subsection 1 of section B.

In terms of the section stack, this directive swaps the current section with the top section on the section stack.

7.86 `.print string`

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

7.87 `.protected names`

This is one of the ELF visibility directives. The other two are `.hidden` (see [Section 7.57 \[Hidden\]](#), page 55) and `.internal` (see [Section 7.64 \[Internal\]](#), page 58).

This directive overrides the named symbols default visibility (which is set by their binding: local, global or weak). The directive sets the visibility to `protected` which means that any references to the symbols from within the components that defines them must be resolved to the definition in that component, even if a definition in another component would normally preempt this.

7.88 `.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`.

7.89 `.purgem name`

Undefine the macro *name*, so that later uses of the string will not be expanded. See [Section 7.77 \[Macro\]](#), page 61.

7.90 `.pushsection name [, subsection] [, "flags" [, @type [, arguments]]]`

This is one of the ELF section stack manipulation directives. The others are `.section` (see [Section 7.96 \[Section\]](#), page 68), `.subsection` (see [Section 7.107 \[SubSection\]](#), page 73),

`.popsection` (see [Section 7.84 \[PopSection\]](#), page 65), and `.previous` (see [Section 7.85 \[Previous\]](#), page 65).

This directive pushes the current section (and subsection) onto the top of the section stack, and then replaces the current section and subsection with `name` and `subsection`. The optional `flags`, `type` and `arguments` are treated the same as in the `.section` (see [Section 7.96 \[Section\]](#), page 68) directive.

7.91 `.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.

7.92 `.reloc offset, reloc_name[, expression]`

Generate a relocation at *offset* of type *reloc_name* with value *expression*. If *offset* is a number, the relocation is generated in the current section. If *offset* is an expression that resolves to a symbol plus offset, the relocation is generated in the given symbol's section. *expression*, if present, must resolve to a symbol plus addend or to an absolute value, but note that not all targets support an addend. e.g. ELF REL targets such as i386 store an addend in the section contents rather than in the relocation. This low level interface does not support addends stored in the section.

7.93 `.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
```

7.94 `.sbtbl "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.

7.95 `.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.

7.96 `.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.

COFF Version

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

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

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:

| | |
|------------|---|
| b | bss section (uninitialized data) |
| n | section is not loaded |
| w | writable section |
| d | data section |
| r | read-only section |
| x | executable section |
| s | shared section (meaningful for PE targets) |
| a | ignored. (For compatibility with the ELF version) |
| y | section is not readable (meaningful for PE targets) |
| 0-9 | single-digit power-of-two section alignment (GNU extension) |

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. Note the **n** and **w** flags remove attributes from the section, rather than adding them, so if they are used on their own it will be as if no flags had been specified at all.

If the optional argument to the `.section` directive is not quoted, it is taken as a subsection number (see [Section 4.4 \[Sub-Sections\]](#), page 35).

ELF Version

This is one of the ELF section stack manipulation directives. The others are `.subsection` (see [Section 7.107 \[SubSection\]](#), page 73), `.pushsection` (see [Section 7.90 \[PushSection\]](#), page 66), `.popsection` (see [Section 7.84 \[PopSection\]](#), page 65), and `.previous` (see [Section 7.85 \[Previous\]](#), page 65).

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

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

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

| | |
|---|---|
| a | section is allocatable |
| e | section is excluded from executable and shared library. |
| w | section is writable |
| x | section is executable |
| M | section is mergeable |
| S | section contains zero terminated strings |
| G | section is a member of a section group |
| T | section is used for thread-local-storage |
| ? | section is a member of the previously-current section's group, if any |

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) |
| @note | section contains data which is used by things other than the program |
| @init_array | section contains an array of pointers to init functions |
| @fini_array | section contains an array of pointers to finish functions |
| @preinit_array | section contains an array of pointers to pre-init functions |

Many targets only support the first three section types.

Note on targets where the @ character is the start of a comment (eg ARM) then another character is used instead. For example the ARM port uses the % character.

If *flags* contains the M symbol then the *type* argument must be specified as well as an extra argument—*entsize*—like this:

```
.section name , "flags"M, @type, entsize
```

Sections with the M flag but not S flag must contain fixed size constants, each *entsize* octets long. Sections with both M and S must contain zero terminated strings where each character is *entsize* bytes long. The linker may remove duplicates within sections with the same name, same entity size and same flags. *entsize* must be an absolute expression. For sections with both M and S, a string which is a suffix of a larger string is considered a duplicate. Thus "def" will be merged with "abcdef"; A reference to the first "def" will be changed to a reference to "abcdef"+3.

If *flags* contains the G symbol then the *type* argument must be present along with an additional field like this:

```
.section name , "flags"G, @type, GroupName[, linkage]
```

The *GroupName* field specifies the name of the section group to which this particular section belongs. The optional linkage field can contain:

comdat indicates that only one copy of this section should be retained

.gnu.linkonce
an alias for **comdat**

Note: if both the *M* and *G* flags are present then the fields for the Merge flag should come first, like this:

```
.section name , "flags"MG, @type, entsize, GroupName[, linkage]
```

If *flags* contains the ? symbol then it may not also contain the *G* symbol and the *GroupName* or *linkage* fields should not be present. Instead, ? says to consider the section that's current before this directive. If that section used *G*, then the new section will use *G* with those same *GroupName* and *linkage* fields implicitly. If not, then the ? symbol has no effect.

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:

#alloc section is allocatable
#write section is writable
#execinstr
section is executable
#exclude section is excluded from executable and shared library.
#tls section is used for thread local storage

This directive replaces the current section and subsection. See the contents of the gas test suite directory `gas/testsuite/gas/elf` for some examples of how this directive and the other section stack directives work.

7.97 .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 41).

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.

On Z80 **set** is a real instruction, use '*symbol* **defl** *expression*' instead.

7.98 *.short expressions*

.short is normally the same as '*.word*'. See [Section 7.121 \[*.word*\], page 77](#).

In some configurations, however, *.short* and *.word* generate numbers of different lengths. See [Chapter 9 \[Machine Dependencies\], page 81](#).

7.99 *.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 9 \[Machine Dependencies\], page 81](#).

7.100 *.size*

This directive is used to set the size associated with a symbol.

COFF Version

For COFF targets, the *.size* directive is only permitted inside *.def*/*.endef* pairs. It is used like this:

```
.size expression
```

ELF Version

For ELF targets, the *.size* directive is used like this:

```
.size name , expression
```

This directive sets the size associated with a symbol *name*. The size in bytes is computed from *expression* which can make use of label arithmetic. This directive is typically used to set the size of function symbols.

7.101 *.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*'.

7.102 *.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 7.113 \[*.uleb128*\], page 75](#).

7.103 *.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 9.11.5 \[HPPA Assembler Directives\], page 133](#), for a summary.

7.104 .stabd, .stabn, .stabs

There are three directives that begin '`.stab`'. All emit symbols (see [Chapter 5 \[Symbols\]](#), [page 39](#)), 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.

7.105 .string "str", .string8 "str", .string16

`"str", .string32 "str", .string64 "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 29](#).

The variants `string16`, `string32` and `string64` differ from the `string` pseudo opcode in that each 8-bit character from *str* is copied and expanded to 16, 32 or 64 bits respectively. The expanded characters are stored in target endianness byte order.

Example:

```
.string32 "BYE"
expands to:
.string    "B\0\0\0Y\0\0\0E\0\0\0" /* On little endian targets. */
.string    "\0\0\0B\0\0\0Y\0\0\0E"  /* On big endian targets. */
```

7.106 `.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.

7.107 `.subsection name`

This is one of the ELF section stack manipulation directives. The others are `.section` (see [Section 7.96 \[Section\]](#), page 68), `.pushsection` (see [Section 7.90 \[PushSection\]](#), page 66), `.popsection` (see [Section 7.84 \[PopSection\]](#), page 65), and `.previous` (see [Section 7.85 \[Previous\]](#), page 65).

This directive replaces the current subsection with *name*. The current section is not changed. The replaced subsection is put onto the section stack in place of the then current top of stack subsection.

7.108 `.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 can be used like this:

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

If the symbol *name* is defined within the file being assembled, the `.symver` 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.

If the symbol *name* is not defined within the file being assembled, all references to *name* will be changed to *name2@nodename*. If no reference to *name* is made, *name2@nodename* will be removed from the symbol table.

Another usage of the `.symver` directive is:

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

In this case, the symbol *name* must exist and be defined within the file being assembled. It is similar to *name2@nodename*. The difference is *name2@@nodename* will also be used to resolve references to *name2* by the linker.

The third usage of the `.symver` directive is:

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

When *name* is not defined within the file being assembled, it is treated as *name2@nodename*. When *name* is defined within the file being assembled, the symbol name, *name*, will be changed to *name2@@nodename*.

7.109 `.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.

7.110 `.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.

7.111 `.title "heading"`

Use *heading* as the title (second line, immediately after the source file name and pagenum-ber) 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.

7.112 `.type`

This directive is used to set the type of a symbol.

COFF Version

For COFF targets, this directive is permitted only within `.def/.endef` pairs. It is used like this:

```
.type int
```

This records the integer *int* as the type attribute of a symbol table entry.

ELF Version

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

```
.type name , type description
```

This sets the type of symbol *name* to be either a function symbol or an object symbol. There are five different syntaxes supported for the *type description* field, in order to provide compatibility with various other assemblers.

Because some of the characters used in these syntaxes (such as ‘@’ and ‘#’) are comment characters for some architectures, some of the syntaxes below do not work on all architectures. The first variant will be accepted by the GNU assembler on all architectures so that variant should be used for maximum portability, if you do not need to assemble your code with other assemblers.

The syntaxes supported are:

```
.type <name> STT_<TYPE_IN_UPPER_CASE>
.type <name>,#<type>
.type <name>,@<type>
.type <name>,%<type>
.type <name>,"<type>"
```

The types supported are:

STT_FUNC

function Mark the symbol as being a function name.

STT_GNU_IFUNC

gnu_indirect_function

Mark the symbol as an indirect function when evaluated during reloc processing.
(This is only supported on Linux targeted assemblers).

STT_OBJECT

object Mark the symbol as being a data object.

STT_TLS

tls_object

Mark the symbol as being a thread-local data object.

STT_COMMON

common Mark the symbol as being a common data object.

STT_NOTYPE

notype Does not mark the symbol in any way. It is supported just for completeness.

gnu_unique_object

Marks the symbol as being a globally unique data object. The dynamic linker will make sure that in the entire process there is just one symbol with this name and type in use. (This is only supported on Linux targeted assemblers).

Note: Some targets support extra types in addition to those listed above.

7.113 .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 7.102](#) [[.sleb128](#)], [page 71](#).

7.114 `.val addr`

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

7.115 `.version "string"`

This directive creates a `.note` section and places into it an ELF formatted note of type `NT_VERSION`. The note's name is set to *string*.

7.116 `.vtable_entry table, offset`

This directive finds or creates a symbol *table* and creates a `VTABLE_ENTRY` relocation for it with an addend of *offset*.

7.117 `.vtable_inherit child, parent`

This directive finds the symbol *child* and finds or creates the symbol *parent* and then creates a `VTABLE_INHERIT` relocation for the parent whose addend is the value of the child symbol. As a special case the parent name of 0 is treated as referring to the `*ABS*` section.

7.118 `.warning "string"`

Similar to the directive `.error` (see [Section 7.47 \[".error "string"\]](#), page 54), but just emits a warning.

7.119 `.weak names`

This directive sets the weak attribute on the comma separated list of symbol *names*. If the symbols do not already exist, they will be created.

On COFF targets other than PE, weak symbols are a GNU extension. This directive sets the weak attribute on the comma separated list of symbol *names*. If the symbols do not already exist, they will be created.

On the PE target, weak symbols are supported natively as weak aliases. When a weak symbol is created that is not an alias, GAS creates an alternate symbol to hold the default value.

7.120 `.weakref alias, target`

This directive creates an alias to the target symbol that enables the symbol to be referenced with weak-symbol semantics, but without actually making it weak. If direct references or definitions of the symbol are present, then the symbol will not be weak, but if all references to it are through weak references, the symbol will be marked as weak in the symbol table.

The effect is equivalent to moving all references to the alias to a separate assembly source file, renaming the alias to the symbol in it, declaring the symbol as weak there, and running a reloadable link to merge the object files resulting from the assembly of the new source file and the old source file that had the references to the alias removed.

The alias itself never makes to the symbol table, and is entirely handled within the assembler.

7.121 `.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 9 \[Machine Dependencies\]](#), page 81), 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.

7.122 Deprecated Directives

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

`.abort`

`.line`

8 Object Attributes

`as` assembles source files written for a specific architecture into object files for that architecture. But not all object files are alike. Many architectures support incompatible variations. For instance, floating point arguments might be passed in floating point registers if the object file requires hardware floating point support—or floating point arguments might be passed in integer registers if the object file supports processors with no hardware floating point unit. Or, if two objects are built for different generations of the same architecture, the combination may require the newer generation at run-time.

This information is useful during and after linking. At link time, `ld` can warn about incompatible object files. After link time, tools like `gdb` can use it to process the linked file correctly.

Compatibility information is recorded as a series of object attributes. Each attribute has a *vendor*, *tag*, and *value*. The vendor is a string, and indicates who sets the meaning of the tag. The tag is an integer, and indicates what property the attribute describes. The value may be a string or an integer, and indicates how the property affects this object. Missing attributes are the same as attributes with a zero value or empty string value.

Object attributes were developed as part of the ABI for the ARM Architecture. The file format is documented in *ELF for the ARM Architecture*.

8.1 GNU Object Attributes

The `.gnu_attribute` directive records an object attribute with vendor ‘`gnu`’.

Except for ‘`Tag_compatibility`’, which has both an integer and a string for its value, GNU attributes have a string value if the tag number is odd and an integer value if the tag number is even. The second bit (`tag & 2` is set for architecture-independent attributes and clear for architecture-dependent ones.

8.1.1 Common GNU attributes

These attributes are valid on all architectures.

`Tag_compatibility` (32)

The compatibility attribute takes an integer flag value and a vendor name. If the flag value is 0, the file is compatible with other toolchains. If it is 1, then the file is only compatible with the named toolchain. If it is greater than 1, the file can only be processed by other toolchains under some private arrangement indicated by the flag value and the vendor name.

8.1.2 MIPS Attributes

`Tag_GNU_MIPS_ABI_FP` (4)

The floating-point ABI used by this object file. The value will be:

- 0 for files not affected by the floating-point ABI.
- 1 for files using the hardware floating-point with a standard double-precision FPU.
- 2 for files using the hardware floating-point ABI with a single-precision FPU.

- 3 for files using the software floating-point ABI.
- 4 for files using the hardware floating-point ABI with 64-bit wide double-precision floating-point registers and 32-bit wide general purpose registers.

8.1.3 PowerPC Attributes

Tag_GNU_Power_ABI_FP (4)

The floating-point ABI used by this object file. The value will be:

- 0 for files not affected by the floating-point ABI.
- 1 for files using double-precision hardware floating-point ABI.
- 2 for files using the software floating-point ABI.
- 3 for files using single-precision hardware floating-point ABI.

Tag_GNU_Power_ABI_Vector (8)

The vector ABI used by this object file. The value will be:

- 0 for files not affected by the vector ABI.
- 1 for files using general purpose registers to pass vectors.
- 2 for files using AltiVec registers to pass vectors.
- 3 for files using SPE registers to pass vectors.

8.2 Defining New Object Attributes

If you want to define a new GNU object attribute, here are the places you will need to modify. New attributes should be discussed on the ‘`binutils`’ mailing list.

- This manual, which is the official register of attributes.
- The header for your architecture ‘`include/elf`’, to define the tag.
- The ‘`bfd`’ support file for your architecture, to merge the attribute and issue any appropriate link warnings.
- Test cases in ‘`ld/testsuite`’ for merging and link warnings.
- ‘`binutils/readelf.c`’ to display your attribute.
- GCC, if you want the compiler to mark the attribute automatically.

9 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.

9.1 Alpha Dependent Features

9.1.1 Notes

The documentation here is primarily for the ELF object format. `as` also supports the ECOFF and EVAX formats, but features specific to these formats are not yet documented.

9.1.2 Options

`‘-mcpu’` This option specifies the target processor. If an attempt is made to assemble an instruction which will not execute on the target processor, the assembler may either expand the instruction as a macro or issue an error message. This option is equivalent to the `.arch` directive.

The following processor names are recognized: 21064, 21064a, 21066, 21068, 21164, 21164a, 21164pc, 21264, 21264a, 21264b, ev4, ev5, lca45, ev5, ev56, pca56, ev6, ev67, ev68. The special name `all` may be used to allow the assembler to accept instructions valid for any Alpha processor.

In order to support existing practice in OSF/1 with respect to `.arch`, and existing practice within M10 (the Linux ARC bootloader), the numbered processor names (e.g. 21064) enable the processor-specific PALcode instructions, while the “electro-vlasic” names (e.g. ev4) do not.

`‘-mdebug’`

`‘-no-mdebug’`

Enables or disables the generation of `.mdebug` encapsulation for stabs directives and procedure descriptors. The default is to automatically enable `.mdebug` when the first stabs directive is seen.

`‘-relax’` This option forces all relocations to be put into the object file, instead of saving space and resolving some relocations at assembly time. Note that this option does not propagate all symbol arithmetic into the object file, because not all symbol arithmetic can be represented. However, the option can still be useful in specific applications.

`‘-replace’`

`‘-noreplace’`

Enables or disables the optimization of procedure calls, both at assemblage and at link time. These options are only available for VMS targets and `-replace` is the default. See section 1.4.1 of the OpenVMS Linker Utility Manual.

`‘-g’` This option is used when the compiler generates debug information. When `gcc` is using `mips-tfile` to generate debug information for ECOFF, local labels must be passed through to the object file. Otherwise this option has no effect.

`‘-Gsize’` A local common symbol larger than *size* is placed in `.bss`, while smaller symbols are placed in `.sbss`.

`‘-F’`

`‘-32addr’` These options are ignored for backward compatibility.

9.1.3 Syntax

The assembler syntax closely follow the Alpha Reference Manual; assembler directives and general syntax closely follow the OSF/1 and OpenVMS syntax, with a few differences for ELF.

9.1.3.1 Special Characters

‘#’ is the line comment character.

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

9.1.3.2 Register Names

The 32 integer registers are referred to as ‘\$n’ or ‘\$rn’. In addition, registers 15, 28, 29, and 30 may be referred to by the symbols ‘\$fp’, ‘\$at’, ‘\$gp’, and ‘\$sp’ respectively.

The 32 floating-point registers are referred to as ‘\$fn’.

9.1.3.3 Relocations

Some of these relocations are available for ECOFF, but mostly only for ELF. They are modeled after the relocation format introduced in Digital Unix 4.0, but there are additions.

The format is ‘!tag’ or ‘!tag!number’ where *tag* is the name of the relocation. In some cases *number* is used to relate specific instructions.

The relocation is placed at the end of the instruction like so:

```
ldah    $0,a($29)    !gprelhigh
lda     $0,a($0)     !gprellow
ldq     $1,b($29)    !literal!100
ldl     $2,0($1)     !lituse_base!100

!literal
!literal!N
```

Used with an `ldq` instruction to load the address of a symbol from the GOT.

A sequence number *N* is optional, and if present is used to pair `lituse` relocations with this `literal` relocation. The `lituse` relocations are used by the linker to optimize the code based on the final location of the symbol.

Note that these optimizations are dependent on the data flow of the program. Therefore, if *any* `lituse` is paired with a `literal` relocation, then *all* uses of the register set by the `literal` instruction must also be marked with `lituse` relocations. This is because the original `literal` instruction may be deleted or transformed into another instruction.

Also note that there may be a one-to-many relationship between `literal` and `lituse`, but not a many-to-one. That is, if there are two code paths that load up the same address and feed the value to a single use, then the use may not use a `lituse` relocation.

```
!lituse_base!N
```

Used with any memory format instruction (e.g. `ldl`) to indicate that the literal is used for an address load. The offset field of the instruction must be zero. During relaxation, the code may be altered to use a gp-relative load.

!lituse_jsr!N

Used with a register branch format instruction (e.g. `jsr`) to indicate that the literal is used for a call. During relaxation, the code may be altered to use a direct branch (e.g. `bsr`).

!lituse_jsrdirect!N

Similar to `lituse_jsr`, but also that this call cannot be vectored through a PLT entry. This is useful for functions with special calling conventions which do not allow the normal call-clobbered registers to be clobbered.

!lituse_bytoff!N

Used with a byte mask instruction (e.g. `extbl`) to indicate that only the low 3 bits of the address are relevant. During relaxation, the code may be altered to use an immediate instead of a register shift.

!lituse_addr!N

Used with any other instruction to indicate that the original address is in fact used, and the original `ldq` instruction may not be altered or deleted. This is useful in conjunction with `lituse_jsr` to test whether a weak symbol is defined.

```
ldq  $27,foo($29)    !literal!1
beq  $27,is_undef    !lituse_addr!1
jsr  $26,($27),foo   !lituse_jsr!1
```

!lituse_tlsd!N

Used with a register branch format instruction to indicate that the literal is the call to `__tls_get_addr` used to compute the address of the thread-local storage variable whose descriptor was loaded with `!tlsd!N`.

!lituse_tlsldm!N

Used with a register branch format instruction to indicate that the literal is the call to `__tls_get_addr` used to compute the address of the base of the thread-local storage block for the current module. The descriptor for the module must have been loaded with `!tlsldm!N`.

!gpdisp!N

Used with `ldah` and `lda` to load the GP from the current address, a-la the `ldgp` macro. The source register for the `ldah` instruction must contain the address of the `ldah` instruction. There must be exactly one `lda` instruction paired with the `ldah` instruction, though it may appear anywhere in the instruction stream. The immediate operands must be zero.

```
bsr  $26,foo
ldah $29,0($26)    !gpdisp!1
lda  $29,0($29)    !gpdisp!1
```

!gprelhigh

Used with an `ldah` instruction to add the high 16 bits of a 32-bit displacement from the GP.

!gprello

Used with any memory format instruction to add the low 16 bits of a 32-bit displacement from the GP.

| | |
|--|--|
| !gp_{rel} | Used with any memory format instruction to add a 16-bit displacement from the GP. |
| !same_{gp} | Used with any branch format instruction to skip the GP load at the target address. The referenced symbol must have the same GP as the source object file, and it must be declared to either not use \$27 or perform a standard GP load in the first two instructions via the <code>.prologue</code> directive. |
| !t_{lsgd} | |
| !t_{lsgd}!N | Used with an <code>lda</code> instruction to load the address of a TLS descriptor for a symbol in the GOT. The sequence number <i>N</i> is optional, and if present it used to pair the descriptor load with both the <code>literal</code> loading the address of the <code>__tls_get_addr</code> function and the <code>lituse_t_{lsgd}</code> marking the call to that function. For proper relaxation, both the <code>t_{lsgd}</code> , <code>literal</code> and <code>lituse</code> relocations must be in the same extended basic block. That is, the relocation with the lowest address must be executed first at runtime. |
| !t_{lsl_{dm}} | |
| !t_{lsl_{dm}}!N | Used with an <code>lda</code> instruction to load the address of a TLS descriptor for the current module in the GOT. Similar in other respects to <code>t_{lsgd}</code> . |
| !gotd_{tp_{rel}} | Used with an <code>ldq</code> instruction to load the offset of the TLS symbol within its module's thread-local storage block. Also known as the dynamic thread pointer offset or dtp-relative offset. |
| !d_{tp_{rel}}hi | |
| !d_{tp_{rel}}ello | |
| !d_{tp_{rel}} | Like <code>gp_{rel}</code> relocations except they compute dtp-relative offsets. |
| !got_{tp_{rel}} | |
| | Used with an <code>ldq</code> instruction to load the offset of the TLS symbol from the thread pointer. Also known as the tp-relative offset. |
| !t_{p_{rel}}hi | |
| !t_{p_{rel}}ello | |
| !t_{p_{rel}} | Like <code>gp_{rel}</code> relocations except they compute tp-relative offsets. |

9.1.4 Floating Point

The Alpha family uses both IEEE and VAX floating-point numbers.

9.1.5 Alpha Assembler Directives

`as` for the Alpha supports many additional directives for compatibility with the native assembler. This section describes them only briefly.

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

.arch *cpu*

Specifies the target processor. This is equivalent to the ‘-mcpu’ command-line option. See [Section 9.1.2 \[Alpha Options\]](#), page 82, for a list of values for *cpu*.

.ent *function* [, *n*]

Mark the beginning of *function*. An optional number may follow for compatibility with the OSF/1 assembler, but is ignored. When generating *.mdebug* information, this will create a procedure descriptor for the function. In ELF, it will mark the symbol as a function a-la the generic *.type* directive.

.end *function*

Mark the end of *function*. In ELF, it will set the size of the symbol a-la the generic *.size* directive.

.mask *mask*, *offset*

Indicate which of the integer registers are saved in the current function’s stack frame. *mask* is interpreted a bit mask in which bit *n* set indicates that register *n* is saved. The registers are saved in a block located *offset* bytes from the *canonical frame address* (CFA) which is the value of the stack pointer on entry to the function. The registers are saved sequentially, except that the return address register (normally \$26) is saved first.

This and the other directives that describe the stack frame are currently only used when generating *.mdebug* information. They may in the future be used to generate DWARF2 *.debug_frame* unwind information for hand written assembly.

.fmask *mask*, *offset*

Indicate which of the floating-point registers are saved in the current stack frame. The *mask* and *offset* parameters are interpreted as with *.mask*.

.frame *framereg*, *frameoffset*, *retreg* [, *argoffset*]

Describes the shape of the stack frame. The frame pointer in use is *framereg*; normally this is either \$fp or \$sp. The frame pointer is *frameoffset* bytes below the CFA. The return address is initially located in *retreg* until it is saved as indicated in *.mask*. For compatibility with OSF/1 an optional *argoffset* parameter is accepted and ignored. It is believed to indicate the offset from the CFA to the saved argument registers.

.prologue *n*

Indicate that the stack frame is set up and all registers have been spilled. The argument *n* indicates whether and how the function uses the incoming *procedure vector* (the address of the called function) in \$27. 0 indicates that \$27 is not used; 1 indicates that the first two instructions of the function use \$27 to perform a load of the GP register; 2 indicates that \$27 is used in some non-standard way and so the linker cannot elide the load of the procedure vector during relaxation.

.usepv *function*, *which*

Used to indicate the use of the \$27 register, similar to *.prologue*, but without the other semantics of needing to be inside an open *.ent/.end* block.

The *which* argument should be either **no**, indicating that **\$27** is not used, or **std**, indicating that the first two instructions of the function perform a GP load.

One might use this directive instead of **.prologue** if you are also using dwarf2 CFI directives.

.gprel32 *expression*

Computes the difference between the address in *expression* and the GP for the current object file, and stores it in 4 bytes. In addition to being smaller than a full 8 byte address, this also does not require a dynamic relocation when used in a shared library.

.t_floating *expression*

Stores *expression* as an IEEE double precision value.

.s_floating *expression*

Stores *expression* as an IEEE single precision value.

.f_floating *expression*

Stores *expression* as a VAX F format value.

.g_floating *expression*

Stores *expression* as a VAX G format value.

.d_floating *expression*

Stores *expression* as a VAX D format value.

.set *feature*

Enables or disables various assembler features. Using the positive name of the feature enables while using '**no*feature***' disables.

at Indicates that macro expansions may clobber the *assembler temporary* (**\$at** or **\$28**) register. Some macros may not be expanded without this and will generate an error message if **noat** is in effect. When **at** is in effect, a warning will be generated if **\$at** is used by the programmer.

macro Enables the expansion of macro instructions. Note that variants of real instructions, such as **br label** vs **br \$31,label** are considered alternate forms and not macros.

move

reorder

volatile These control whether and how the assembler may re-order instructions. Accepted for compatibility with the OSF/1 assembler, but **as** does not do instruction scheduling, so these features are ignored.

The following directives are recognized for compatibility with the OSF/1 assembler but are ignored.

| | |
|----------------|-----------------|
| .proc | .aproc |
| .reguse | .livereg |
| .option | .aent |
| .ugen | .eflag |
| .alias | .noalias |

9.1.6 Opcodes

For detailed information on the Alpha machine instruction set, see the Alpha Architecture Handbook located at

`ftp://ftp.digital.com/pub/Digital/info/semiconductor/literature/alphaahb.pdf`

9.2 ARC Dependent Features

9.2.1 Options

`-marc[5|6|7|8]`

This option selects the core processor variant. Using `-marc` is the same as `-marc6`, which is also the default.

`arc5` Base instruction set.

`arc6` Jump-and-link (jl) instruction. No requirement of an instruction between setting flags and conditional jump. For example:

```
mov.f r0,r1
beq   foo
```

`arc7` Break (brk) and sleep (sleep) instructions.

`arc8` Software interrupt (swi) instruction.

Note: the `.option` directive can to be used to select a core variant from within assembly 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 - this is the default.

9.2.2 Syntax

9.2.2.1 Special Characters

`*TODO*`

9.2.2.2 Register Names

`*TODO*`

9.2.3 Floating Point

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

9.2.4 ARC Machine Directives

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

`.2byte expressions`
`*TODO*`

`.3byte expressions`
`*TODO*`

`.4byte expressions`
`*TODO*`

.extAuxRegister *name*,*address*,*mode*

The ARCTangent A4 has extensible auxiliary register space. The auxiliary registers can be defined in the assembler source code by using this directive. The first parameter is the *name* of the new auxiliary register. The second parameter is the *address* of the register in the auxiliary register memory map for the variant of the ARC. The third parameter specifies the *mode* in which the register can be operated is and it can be one of:

```
r (readonly)
w (write only)
r|w (read or write)
```

For example:

```
.extAuxRegister mulhi,0x12,w
```

This specifies an extension auxiliary register called *mulhi* which is at address 0x12 in the memory space and which is only writable.

.extCondCode *suffix*,*value*

The condition codes on the ARCTangent A4 are extensible and can be specified by means of this assembler directive. They are specified by the suffix and the value for the condition code. They can be used to specify extra condition codes with any values. For example:

```
.extCondCode is_busy,0x14
```

```
add.is_busy r1,r2,r3
bis_busy    _main
```

.extCoreRegister *name*,*regnum*,*mode*,*shortcut*

Specifies an extension core register *name* for the application. This allows a register *name* with a valid *regnum* between 0 and 60, with the following as valid values for *mode*

```
'r (readonly)'
'w (write only)'
'r|w (read or write)'
```

The other parameter gives a description of the register having a *shortcut* in the pipeline. The valid values are:

```
can_shortcut
cannot_shortcut
```

For example:

```
.extCoreRegister mlo,57,r,can_shortcut
```

This defines an extension core register mlo with the value 57 which can shortcut the pipeline.

.extInstruction *name*,*opcode*,*subopcode*,*suffixclass*,*syntaxclass*

The ARCTangent A4 allows the user to specify extension instructions. The extension instructions are not macros. The assembler creates encodings for use of these instructions according to the specification by the user. The parameters are:

- *name* Name of the extension instruction
- *opcode* Opcode to be used. (Bits 27:31 in the encoding). Valid values 0x10-0x1f or 0x03
- *subopcode* Subopcode to be used. Valid values are from 0x09-0x3f. However the correct value also depends on *syntaxclass*
- *suffixclass* Determines the kinds of suffixes to be allowed. Valid values are SUFFIX_NONE, SUFFIX_COND, SUFFIX_FLAG which indicates the absence or presence of conditional suffixes and flag setting by the extension instruction. It is also possible to specify that an instruction sets the flags and is conditional by using SUFFIX_CODE | SUFFIX_FLAG.
- *syntaxclass* Determines the syntax class for the instruction. It can have the following values:
 SYNTAX_2OP:
 2 Operand Instruction
 SYNTAX_3OP:
 3 Operand Instruction

In addition there could be modifiers for the syntax class as described below:

Syntax Class Modifiers are:

- OP1_MUST_BE_IMM: Modifies syntax class SYNTAX_3OP, specifying that the first operand of a three-operand instruction must be an immediate (i.e., the result is discarded). OP1_MUST_BE_IMM is used by bitwise ORing it with SYNTAX_3OP as given in the example below. This could usually be used to set the flags using specific instructions and not retain results.
- OP1_IMM IMPLIED: Modifies syntax class SYNTAX_2OP, it specifies that there is an implied immediate destination operand which does not appear in the syntax. For example, if the source code contains an instruction like:

```
inst r1,r2
```

it really means that the first argument is an implied immediate (that is, the result is discarded). This is the same as though the source code were: inst 0,r1,r2. You use OP1_IMM IMPLIED by bitwise ORing it with SYNTAX_2OP.

For example, defining 64-bit multiplier with immediate operands:

```
.extInstruction mp64,0x14,0x0,SUFFIX_COND | SUFFIX_FLAG ,  
                 SYNTAX_3OP|OP1_MUST_BE_IMM
```

The above specifies an extension instruction called `mp64` which has 3 operands, sets the flags, can be used with a condition code, for which the first operand is an immediate. (Equivalent to discarding the result of the operation).

```
.extInstruction mul64,0x14,0x00,SUFFIX_COND, SYNTAX_2OP|OP1_IMM IMPLIED
```

This describes a 2 operand instruction with an implicit first immediate operand. The result of this operation would be discarded.

```
.half expressions
    *TODO*
```

```
.long expressions
    *TODO*
```

```
.option arc|arc5|arc6|arc7|arc8
```

The `.option` directive must be followed by the desired core version. Again `arc` is an alias for `arc6`.

Note: the `.option` directive overrides the command line option `-marc`; a warning is emitted when the version is not consistent between the two - even for the implicit default core version (`arc6`).

```
.short expressions
    *TODO*
```

```
.word expressions
    *TODO*
```

9.2.5 Opcodes

For information on the ARC instruction set, see *ARC Programmers Reference Manual*, ARC International (www.arc.com)

9.3 ARM Dependent Features

9.3.1 Options

`-mcpu=processor[+extension...]`

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. The following processor names are recognized: `arm1`, `arm2`, `arm250`, `arm3`, `arm6`, `arm60`, `arm600`, `arm610`, `arm620`, `arm7`, `arm7m`, `arm7d`, `arm7dm`, `arm7di`, `arm7dmi`, `arm70`, `arm700`, `arm700i`, `arm710`, `arm710t`, `arm720`, `arm720t`, `arm740t`, `arm710c`, `arm7100`, `arm7500`, `arm7500fe`, `arm7t`, `arm7tdmi`, `arm7tdmi-s`, `arm8`, `arm810`, `strongarm`, `strongarm1`, `strongarm110`, `strongarm1100`, `strongarm1110`, `arm9`, `arm920`, `arm920t`, `arm922t`, `arm940t`, `arm9tdmi`, `fa526` (Faraday FA526 processor), `fa626` (Faraday FA626 processor), `arm9e`, `arm926e`, `arm926ej-s`, `arm946e-r0`, `arm946e`, `arm946e-s`, `arm966e-r0`, `arm966e`, `arm966e-s`, `arm968e-s`, `arm10t`, `arm10tdmi`, `arm10e`, `arm1020`, `arm1020t`, `arm1020e`, `arm1022e`, `arm1026ej-s`, `fa626te` (Faraday FA626TE processor), `fa726te` (Faraday FA726TE processor), `arm1136j-s`, `arm1136jf-s`, `arm1156t2-s`, `arm1156t2f-s`, `arm1176jz-s`, `arm1176jzf-s`, `mpcore`, `mpcorenovfp`, `cortex-a5`, `cortex-a8`, `cortex-a9`, `cortex-a15`, `cortex-r4`, `cortex-r4f`, `cortex-m4`, `cortex-m3`, `cortex-m1`, `cortex-m0`, `ep9312` (ARM920 with Cirrus Maverick coprocessor), `i80200` (Intel XScale processor) `iwmmxt` (Intel(r) XScale processor with Wireless MMX(tm) technology coprocessor) and `xscale`. The special name `all` may be used to allow the assembler to accept instructions valid for any ARM processor.

In addition to the basic instruction set, the assembler can be told to accept various extension mnemonics that extend the processor using the co-processor instruction space. For example, `-mcpu=arm920+maverick` is equivalent to specifying `-mcpu=ep9312`.

Multiple extensions may be specified, separated by a `+`. The extensions should be specified in ascending alphabetical order.

Some extensions may be restricted to particular architectures; this is documented in the list of extensions below.

Extension mnemonics may also be removed from those the assembler accepts. This is done by prepending `no` to the option that adds the extension. Extensions that are removed should be listed after all extensions which have been added, again in ascending alphabetical order. For example, `-mcpu=ep9312+nomaverick` is equivalent to specifying `-mcpu=arm920`.

The following extensions are currently supported: `idiv`, (Integer Divide Extensions for v7-A architecture), `iwmmxt`, `iwmmxt2`, `maverick`, `mp` (Multiprocessing Extensions for v7-A and v7-R architectures), `os` (Operating System for v6M architecture), `sec` (Security Extensions for v6K and v7-A architectures), `virt` (Virtualization Extensions for v7-A architecture, implies `idiv`), and `xscale`.

-march=architecture[+extension...]

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. The following architecture names are recognized: `armv1`, `armv2`, `armv2a`, `armv2s`, `armv3`, `armv3m`, `armv4`, `armv4xm`, `armv4t`, `armv4txm`, `armv5`, `armv5t`, `armv5txm`, `armv5te`, `armv5teexp`, `armv6`, `armv6j`, `armv6k`, `armv6z`, `armv6zk`, `armv6-m`, `armv6s-m`, `armv7`, `armv7-a`, `armv7-r`, `armv7-m`, `armv7e-m`, `iwmmxt` and `xscale`. If both `-mcpu` and `-march` are specified, the assembler will use the setting for `-mcpu`.

The architecture option can be extended with the same instruction set extension options as the `-mcpu` option.

-mfpu=floating-point-format

This option specifies the floating point format to assemble for. The assembler will issue an error message if an attempt is made to assemble an instruction which will not execute on the target floating point unit. The following format options are recognized: `softfpa`, `fpe`, `fpe2`, `fpe3`, `fpa`, `fpa10`, `fpa11`, `arm7500fe`, `softvfp`, `softvfp+vfp`, `vfp`, `vfp10`, `vfp10-r0`, `vfp9`, `vfpvd`, `vfpv2`, `vfpv3`, `vfpv3-fp16`, `vfpv3-d16`, `vfpv3-d16-fp16`, `vfpv3xd`, `vfpv3xd-d16`, `vfpv4`, `vfpv4-d16`, `fpv4-sp-d16`, `arm1020t`, `arm1020e`, `arm1136jf-s`, `maverick`, `neon`, and `neon-vfpv4`.

In addition to determining which instructions are assembled, this option also affects the way in which the `.double` assembler directive behaves when assembling little-endian code.

The default is dependent on the processor selected. For Architecture 5 or later, the default is to assemble for VFP instructions; for earlier architectures the default is to assemble for FPA instructions.

-mthumb This option specifies that the assembler should start assembling Thumb instructions; that is, it should behave as though the file starts with a `.code 16` directive.

-mthumb-interwork

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

-mimplicit-it=never

-mimplicit-it=always

-mimplicit-it=arm

-mimplicit-it=thumb

The `-mimplicit-it` option controls the behavior of the assembler when conditional instructions are not enclosed in IT blocks. There are four possible behaviors. If `never` is specified, such constructs cause a warning in ARM code and an error in Thumb-2 code. If `always` is specified, such constructs are accepted in both ARM and Thumb-2 code, where the IT instruction is added implicitly. If `arm` is specified, such constructs are accepted in ARM code and cause an error in Thumb-2 code. If `thumb` is specified, such constructs cause a warning in ARM code and are accepted in Thumb-2 code. If you omit this option, the behavior is equivalent to `-mimplicit-it=arm`.

`-mapcs-26`

`-mapcs-32`

These options specify that the output generated by the assembler should be marked as supporting the indicated version of the Arm Procedure Calling Standard.

`-matpcs` This option specifies that the output generated by the assembler should be marked as supporting the Arm/Thumb Procedure Calling Standard. If enabled this option will cause the assembler to create an empty debugging section in the object file called `.arm.atpcs`. Debuggers can use this to determine the ABI being used by.

`-mapcs-float`

This indicates 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.

`-mfloat-abi=abi`

This option specifies that the output generated by the assembler should be marked as using specified floating point ABI. The following values are recognized: `soft`, `softfp` and `hard`.

`-meabi=ver`

This option specifies which EABI version the produced object files should conform to. The following values are recognized: `gnu`, 4 and 5.

`-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 specifies that the output of the assembler should be marked as position-independent code (PIC).

`--fix-v4bx`

Allow `BX` instructions in ARMv4 code. This is intended for use with the linker option of the same name.

`-mwarn-deprecated`

`-mno-warn-deprecated`

Enable or disable warnings about using deprecated options or features. The default is to warn.

9.3.2 Syntax

9.3.2.1 Instruction Set Syntax

Two slightly different syntaxes are support for ARM and THUMB instructions. The default, `divided`, uses the old style where ARM and THUMB instructions had their own, separate

syntaxes. The new, **unified** syntax, which can be selected via the `.syntax` directive, and has the following main features:

- Immediate operands do not require a `#` prefix.
- The `IT` instruction may appear, and if it does it is validated against subsequent conditional affixes. In ARM mode it does not generate machine code, in THUMB mode it does.
- For ARM instructions the conditional affixes always appear at the end of the instruction. For THUMB instructions conditional affixes can be used, but only inside the scope of an `IT` instruction.
- All of the instructions new to the V6T2 architecture (and later) are available. (Only a few such instructions can be written in the **divided** syntax).
- The `.N` and `.W` suffixes are recognized and honored.
- All instructions set the flags if and only if they have an `s` affix.

9.3.2.2 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.

The `';` character 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.

9.3.2.3 Register Names

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

9.3.2.4 NEON Alignment Specifiers

Some NEON load/store instructions allow an optional address alignment qualifier. The ARM documentation specifies that this is indicated by `'@ align'`. However GAS already interprets the `'@'` character as a "line comment" start, so `': align'` is used instead. For example:

```
vld1.8 {q0}, [r0, :128]
```

9.3.3 Floating Point

The ARM family uses IEEE floating-point numbers.

9.3.3.1 ARM relocation generation

Specific data relocations can be generated by putting the relocation name in parentheses after the symbol name. For example:

```
.word foo(TARGET1)
```

This will generate an `'R_ARM_TARGET1'` relocation against the symbol `foo`. The following relocations are supported: `GOT`, `GOTOFF`, `TARGET1`, `TARGET2`, `SBREL`, `TLSGD`, `TLSLDM`, `TLSLDO`, `GOTTPOFF`, `GOT_PREL` and `TPOFF`.

For compatibility with older toolchains the assembler also accepts (PLT) after branch targets. This will generate the deprecated ‘R_ARM_PLT32’ relocation.

Relocations for ‘MOVW’ and ‘MOVT’ instructions can be generated by prefixing the value with ‘#:lower16:’ and ‘#:upper16:’ respectively. For example to load the 32-bit address of foo into r0:

```
MOVW r0, #:lower16:foo
MOVT r0, #:upper16:foo
```

9.3.4 ARM Machine Directives

`.2byte expression [, expression]*`

`.4byte expression [, expression]*`

`.8byte expression [, expression]*`

These directives write 2, 4 or 8 byte values to the output section.

`.align expression [, expression]`

This is the generic `.align` directive. For the ARM however if the first argument is zero (ie no alignment is needed) the assembler will behave as if the argument had been 2 (ie pad to the next four byte boundary). This is for compatibility with ARM’s own assembler.

`.arch name`

Select the target architecture. Valid values for *name* are the same as for the ‘-march’ commandline option.

Specifying `.arch` clears any previously selected architecture extensions.

`.arch_extension name`

Add or remove an architecture extension to the target architecture. Valid values for *name* are the same as those accepted as architectural extensions by the ‘-mcpu’ commandline option.

`.arch_extension` may be used multiple times to add or remove extensions incrementally to the architecture being compiled for.

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

`.pad #count`

Generate unwinder annotations for a stack adjustment of *count* bytes. A positive value indicates the function prologue allocated stack space by decrementing the stack pointer.

`.bss` This directive switches to the `.bss` section.

`.cantunwind`

Prevents unwinding through the current function. No personality routine or exception table data is required or permitted.

`.code [16|32]`

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

`.cpu name`

Select the target processor. Valid values for *name* are the same as for the ‘-mcpu’ commandline option.

Specifying `.cpu` clears any previously selected architecture extensions.

`name .dn register name [.type] [[index]]`

`name .qn register name [.type] [[index]]`

The `dn` and `qn` directives are used to create typed and/or indexed register aliases for use in Advanced SIMD Extension (Neon) instructions. The former should be used to create aliases of double-precision registers, and the latter to create aliases of quad-precision registers.

If these directives are used to create typed aliases, those aliases can be used in Neon instructions instead of writing types after the mnemonic or after each operand. For example:

```
x .dn d2.f32
y .dn d3.f32
z .dn d4.f32[1]
vmul x,y,z
```

This is equivalent to writing the following:

```
vmul.f32 d2,d3,d4[1]
```

Aliases created using `dn` or `qn` can be destroyed using `unreq`.

`.eabi_attribute tag, value`

Set the EABI object attribute `tag` to `value`.

The `tag` is either an attribute number, or one of the following: `Tag_CPU_raw_name`, `Tag_CPU_name`, `Tag_CPU_arch`, `Tag_CPU_arch_profile`, `Tag_ARM_ISA_use`, `Tag_THUMB_ISA_use`, `Tag_FP_arch`, `Tag_WMMX_arch`, `Tag_Advanced_SIMD_arch`, `Tag_PCS_config`, `Tag_ABI_PCS_R9_use`, `Tag_ABI_PCS_RW_data`, `Tag_ABI_PCS_RO_data`, `Tag_ABI_PCS_GOT_use`, `Tag_ABI_PCS_wchar_t`, `Tag_ABI_FP_rounding`, `Tag_ABI_FP_denormal`, `Tag_ABI_FP_exceptions`, `Tag_ABI_FP_user_exceptions`, `Tag_ABI_FP_number_model`, `Tag_ABI_align_needed`, `Tag_ABI_align_preserved`, `Tag_ABI_enum_size`, `Tag_ABI_HardFP_use`, `Tag_ABI_VFP_args`, `Tag_ABI_WMMX_args`, `Tag_ABI_optimization_goals`, `Tag_ABI_FP_optimization_goals`, `Tag_compatibility`, `Tag_CPU_unaligned_access`, `Tag_FP_HP_extension`, `Tag_ABI_FP_16bit_format`, `Tag_MPextension_use`, `Tag_DIV_use`, `Tag_nofaults`, `Tag_also_compatible_with`, `Tag_conformance`, `Tag_T2EE_use`, `Tag_Virtualization_use`

The `value` is either a number, "string", or number, "string" depending on the tag.

Note - the following legacy values are also accepted by `tag`: `Tag_VFP_arch`, `Tag_ABI_align8_needed`, `Tag_ABI_align8_preserved`, `Tag_VFP_HP_extension`,

`.even` This directive aligns to an even-numbered address.

`.extend expression [, expression]*`

`.ldouble expression [, expression]*`

These directives write 12byte long double floating-point values to the output section. These are not compatible with current ARM processors or ABIs.

`.fnend` Marks the end of a function with an unwind table entry. The unwind index table entry is created when this directive is processed.

If no personality routine has been specified then standard personality routine 0 or 1 will be used, depending on the number of unwind opcodes required.

.fnstart Marks the start of a function with an unwind table entry.

.force_thumb

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

.fpu *name*

Select the floating-point unit to assemble for. Valid values for *name* are the same as for the ‘-mfpu’ commandline option.

.handlerdata

Marks the end of the current function, and the start of the exception table entry for that function. Anything between this directive and the **.fnend** directive will be added to the exception table entry.

Must be preceded by a **.personality** or **.personalityindex** directive.

.inst opcode [, ...]

.inst.n opcode [, ...]

.inst.w opcode [, ...]

Generates the instruction corresponding to the numerical value *opcode*. **.inst.n** and **.inst.w** allow the Thumb instruction size to be specified explicitly, overriding the normal encoding rules.

.ldouble *expression* [, *expression*]*

See **.extend**.

.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). **GAS** maintains a separate literal pool for each section and each sub-section. The **.ltorg** directive will only affect the literal pool of the current section and sub-section. At the end of assembly all remaining, un-empty literal pools will automatically be dumped.

Note - older versions of **GAS** would dump the current literal pool any time a section change occurred. This is no longer done, since it prevents accurate control of the placement of literal pools.

.movsp *reg* [, #*offset*]

Tell the unwinder that *reg* contains an offset from the current stack pointer. If *offset* is not specified then it is assumed to be zero.

.object_arch *name*

Override the architecture recorded in the EABI object attribute section. Valid values for *name* are the same as for the **.arch** directive. Typically this is useful when code uses runtime detection of CPU features.

.packed *expression* [, *expression*]*

This directive writes 12-byte packed floating-point values to the output section. These are not compatible with current ARM processors or ABIs.

.pad #count

Generate unwinder annotations for a stack adjustment of *count* bytes. A positive value indicates the function prologue allocated stack space by decrementing the stack pointer.

.personality name

Sets the personality routine for the current function to *name*.

.personalityindex index

Sets the personality routine for the current function to the EABI standard routine number *index*

.pool This is a synonym for **.ltorg**.**name .req register name**

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

```
foo .req r0
```

.save reglist

Generate unwinder annotations to restore the registers in *reglist*. The format of *reglist* is the same as the corresponding store-multiple instruction.

core registers

```
.save {r4, r5, r6, lr}
stmfd sp!, {r4, r5, r6, lr}
```

FPA registers

```
.save f4, 2
sfmfd f4, 2, [sp]!
```

VFP registers

```
.save {d8, d9, d10}
fstmdx sp!, {d8, d9, d10}
```

iWMMXt registers

```
.save {wr10, wr11}
wstrd wr11, [sp, #-8]!
wstrd wr10, [sp, #-8]!
```

or

```
.save wr11
wstrd wr11, [sp, #-8]!
.save wr10
wstrd wr10, [sp, #-8]!
```

.setfp fpreg, spreg [, #offset]

Make all unwinder annotations relative to a frame pointer. Without this the unwinder will use offsets from the stack pointer.

The syntax of this directive is the same as the **add** or **mov** instruction used to set the frame pointer. *spreg* must be either **sp** or mentioned in a previous **.movsp** directive.

```
.movsp ip
mov ip, sp
...
.setfp fp, ip, #4
add fp, ip, #4
```

.secrel32 expression [, expression]*

This directive emits relocations that evaluate to the section-relative offset of each expression's symbol. This directive is only supported for PE targets.

.syntax [unified | divided]

This directive sets the Instruction Set Syntax as described in the [Section 9.3.2.1 \[ARM-Instruction-Set\]](#), page 95 section.

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

.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. The presence of this directive also implies *.thumb*

This directive is not necessary when generating EABI objects. On these targets the encoding is implicit when generating Thumb code.

.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.

.unreq *alias-name*

This undefines a register alias which was previously defined using the *req*, *dn* or *qn* directives. For example:

```
foo .req r0
.unreq foo
```

An error occurs if the name is undefined. Note - this pseudo op can be used to delete builtin in register name aliases (eg 'r0'). This should only be done if it is really necessary.

.unwind_raw *offset, byte1, ...*

Insert one or more arbitrary unwind opcode bytes, which are known to adjust the stack pointer by *offset* bytes.

For example *.unwind_raw 4, 0xb1, 0x01* is equivalent to *.save {r0}*

.vsave *vfp-reglist*

Generate unwinder annotations to restore the VFP registers in *vfp-reglist* using FLDMD. Also works for VFPv3 registers that are to be restored using VLDM. The format of *vfp-reglist* is the same as the corresponding store-multiple instruction.

VFP registers

```
.vsave {d8, d9, d10}
fstmdd sp!, {d8, d9, d10}
```

VFPv3 registers

```
.vsave {d15, d16, d17}
vstm sp!, {d15, d16, d17}
```

Since FLDMD and FSTM are now deprecated, this directive should be used in favour of *.save* for saving VFP registers for ARMv6 and above.

9.3.5 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 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.

9.3.6 Mapping Symbols

The ARM ELF specification requires that special symbols be inserted into object files to mark certain features:

- \$a At the start of a region of code containing ARM instructions.
- \$t At the start of a region of code containing THUMB instructions.
- \$d At the start of a region of data.

The assembler will automatically insert these symbols for you - there is no need to code them yourself. Support for tagging symbols (\$b, \$f, \$p and \$m) which is also mentioned in the current ARM ELF specification is not implemented. This is because they have been dropped from the new EABI and so tools cannot rely upon their presence.

9.3.7 Unwinding

The ABI for the ARM Architecture specifies a standard format for exception unwind information. This information is used when an exception is thrown to determine where control should be transferred. In particular, the unwind information is used to determine which function called the function that threw the exception, and which function called that one, and so forth. This information is also used to restore the values of callee-saved registers in the function catching the exception.

If you are writing functions in assembly code, and those functions call other functions that throw exceptions, you must use assembly pseudo ops to ensure that appropriate exception unwind information is generated. Otherwise, if one of the functions called by your assembly code throws an exception, the run-time library will be unable to unwind the stack through your assembly code and your program will not behave correctly.

To illustrate the use of these pseudo ops, we will examine the code that G++ generates for the following C++ input:

```
void callee (int *);
```

```
int
caller ()
{
    int i;
    callee (&i);
    return i;
}
```

This example does not show how to throw or catch an exception from assembly code. That is a much more complex operation and should always be done in a high-level language, such as C++, that directly supports exceptions.

The code generated by one particular version of G++ when compiling the example above is:

```
_Z6callerv:
    .fnstart
.LFB2:
    @ Function supports interworking.
    @ args = 0, pretend = 0, frame = 8
    @ frame_needed = 1, uses_anonymous_args = 0
    stmfd    sp!, {fp, lr}
    .save {fp, lr}
.LCFI0:
    .setfp fp, sp, #4
    add      fp, sp, #4
.LCFI1:
```

```

        .pad #8
        sub    sp, sp, #8
.LCFI2:
        sub    r3, fp, #8
        mov    r0, r3
        bl     _Z6calleePi
        ldr    r3, [fp, #-8]
        mov    r0, r3
        sub    sp, fp, #4
        ldmdf  sp!, {fp, lr}
        bx     lr
.LFE2:
        .fnend

```

Of course, the sequence of instructions varies based on the options you pass to GCC and on the version of GCC in use. The exact instructions are not important since we are focusing on the pseudo ops that are used to generate unwind information.

An important assumption made by the unwinder is that the stack frame does not change during the body of the function. In particular, since we assume that the assembly code does not itself throw an exception, the only point where an exception can be thrown is from a call, such as the `bl` instruction above. At each call site, the same saved registers (including `lr`, which indicates the return address) must be located in the same locations relative to the frame pointer.

The `.fnstart` (see [\[.fnstart pseudo op\]](#), page 99) pseudo op appears immediately before the first instruction of the function while the `.fnend` (see [\[.fnend pseudo op\]](#), page 98) pseudo op appears immediately after the last instruction of the function. These pseudo ops specify the range of the function.

Only the order of the other pseudos ops (e.g., `.setfp` or `.pad`) matters; their exact locations are irrelevant. In the example above, the compiler emits the pseudo ops with particular instructions. That makes it easier to understand the code, but it is not required for correctness. It would work just as well to emit all of the pseudo ops other than `.fnend` in the same order, but immediately after `.fnstart`.

The `.save` (see [\[.save pseudo op\]](#), page 100) pseudo op indicates registers that have been saved to the stack so that they can be restored before the function returns. The argument to the `.save` pseudo op is a list of registers to save. If a register is “callee-saved” (as specified by the ABI) and is modified by the function you are writing, then your code must save the value before it is modified and restore the original value before the function returns. If an exception is thrown, the run-time library restores the values of these registers from their locations on the stack before returning control to the exception handler. (Of course, if an exception is not thrown, the function that contains the `.save` pseudo op restores these registers in the function epilogue, as is done with the `ldmdf` instruction above.)

You do not have to save callee-saved registers at the very beginning of the function and you do not need to use the `.save` pseudo op immediately following the point at which the registers are saved. However, if you modify a callee-saved register, you must save it on the stack before modifying it and before calling any functions which might throw an exception. And, you must use the `.save` pseudo op to indicate that you have done so.

The `.pad` (see [\[.pad\]](#), page 97) pseudo op indicates a modification of the stack pointer that does not save any registers. The argument is the number of bytes (in decimal) that are subtracted from the stack pointer. (On ARM CPUs, the stack grows downwards, so subtracting from the stack pointer increases the size of the stack.)

The `.setfp` (see [\[.setfp pseudo op\]](#), page 100) pseudo op indicates the register that contains the frame pointer. The first argument is the register that is set, which is typically `fp`. The second argument indicates the register from which the frame pointer takes its value. The third argument, if present, is the value (in decimal) added to the register specified by the second argument to compute the value of the frame pointer. You should not modify the frame pointer in the body of the function.

If you do not use a frame pointer, then you should not use the `.setfp` pseudo op. If you do not use a frame pointer, then you should avoid modifying the stack pointer outside of the function prologue. Otherwise, the run-time library will be unable to find saved registers when it is unwinding the stack.

The pseudo ops described above are sufficient for writing assembly code that calls functions which may throw exceptions. If you need to know more about the object-file format used to represent unwind information, you may consult the *Exception Handling ABI for the ARM Architecture* available from <http://infocenter.arm.com>.

9.4 AVR Dependent Features

9.4.1 Options

`-mmcu=mcu`

Specify ATMEL AVR instruction set or MCU type.

Instruction set `avr1` is for the minimal AVR core, not supported by the C compiler, only for assembler programs (MCU types: `at90s1200`, `attiny11`, `attiny12`, `attiny15`, `attiny28`).

Instruction set `avr2` (default) is for the classic AVR core with up to 8K program memory space (MCU types: `at90s2313`, `at90s2323`, `at90s2333`, `at90s2343`, `attiny22`, `attiny26`, `at90s4414`, `at90s4433`, `at90s4434`, `at90s8515`, `at90c8534`, `at90s8535`).

Instruction set `avr25` is for the classic AVR core with up to 8K program memory space plus the MOVW instruction (MCU types: `attiny13`, `attiny13a`, `attiny2313`, `attiny2313a`, `attiny24`, `attiny24a`, `attiny4313`, `attiny44`, `attiny44a`, `attiny84`, `attiny84a`, `attiny25`, `attiny45`, `attiny85`, `attiny261`, `attiny261a`, `attiny461`, `attiny461a`, `attiny861`, `attiny861a`, `attiny87`, `attiny43u`, `attiny48`, `attiny88`, `at86rf401`, `ata6289`).

Instruction set `avr3` is for the classic AVR core with up to 128K program memory space (MCU types: `at43usb355`, `at76c711`).

Instruction set `avr31` is for the classic AVR core with exactly 128K program memory space (MCU types: `atmega103`, `at43usb320`).

Instruction set `avr35` is for classic AVR core plus MOVW, CALL, and JMP instructions (MCU types: `attiny167`, `at90usb82`, `at90usb162`, `atmega8u2`, `atmega16u2`, `atmega32u2`).

Instruction set `avr4` is for the enhanced AVR core with up to 8K program memory space (MCU types: `atmega48`, `atmega48a`, `atmega48p`, `atmega8`, `atmega88`, `atmega88a`, `atmega88p`, `atmega88pa`, `atmega8515`, `atmega8535`, `atmega8hva`, `at90pwm1`, `at90pwm2`, `at90pwm2b`, `at90pwm3`, `at90pwm3b`, `at90pwm81`).

Instruction set `avr5` is for the enhanced AVR core with up to 128K program memory space (MCU types: `atmega16`, `atmega16a`, `atmega161`, `atmega162`, `atmega163`, `atmega164a`, `atmega164p`, `atmega165`, `atmega165a`, `atmega165p`, `atmega168`, `atmega168a`, `atmega168p`, `atmega169`, `atmega169a`, `atmega169p`, `atmega169pa`, `atmega32`, `atmega323`, `atmega324a`, `atmega324p`, `atmega325`, `atmega325a`, `atmega325p`, `atmega3250`, `atmega3250a`, `atmega3250p`, `atmega328`, `atmega328p`, `atmega329`, `atmega329a`, `atmega329p`, `atmega329pa`, `atmega3290`, `atmega3290a`, `atmega3290p`, `atmega406`, `atmega64`, `atmega640`, `atmega644`, `atmega644a`, `atmega644p`, `atmega644pa`, `atmega645`, `atmega645a`, `atmega645p`, `atmega6450`, `atmega6450a`, `atmega6450p`, `atmega649`, `atmega649a`, `atmega649p`, `atmega6490`, `atmega6490a`, `atmega6490p`, `atmega16hva`, `atmega16hva2`, `atmega16hvb`, `atmega32hvb`, `atmega64hve`, `at90can32`, `at90can64`, `at90pwm216`, `at90pwm316`, `atmega32c1`, `atmega64c1`, `atmega16m1`, `atmega32m1`, `atmega64m1`, `atmega16u4`, `atmega32u4`, `atmega32u6`, `at90usb646`, `at90usb647`, `at94k`, `at90scr100`).

Instruction set `avr51` is for the enhanced AVR core with exactly 128K program memory space (MCU types: `atmega128`, `atmega1280`, `atmega1281`, `atmega1284p`, `atmega128rfa1`, `at90can128`, `at90usb1286`, `at90usb1287`, `m3000`).

Instruction set `avr6` is for the enhanced AVR core with a 3-byte PC (MCU types: `atmega2560`, `atmega2561`).

`-mall-opcodes`

Accept all AVR opcodes, even if not supported by `-mmcu`.

`-mno-skip-bug`

This option disable warnings for skipping two-word instructions.

`-mno-wrap`

This option reject `rjmp/rcall` instructions with 8K wrap-around.

9.4.2 Syntax

9.4.2.1 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.

The ‘\$’ character can be used instead of a newline to separate statements.

9.4.2.2 Register Names

The AVR has 32 x 8-bit general purpose working registers ‘`r0`’, ‘`r1`’, ... ‘`r31`’. Six of the 32 registers can be used as three 16-bit indirect address register pointers for Data Space addressing. One of the these address pointers can also be used as an address pointer for look up tables in Flash program memory. These added function registers are the 16-bit ‘`X`’, ‘`Y`’ and ‘`Z`’ - registers.

`X` = `r26:r27`

`Y` = `r28:r29`

`Z` = `r30:r31`

9.4.2.3 Relocatable Expression Modifiers

The assembler supports several modifiers when using relocatable addresses in AVR instruction operands. The general syntax is the following:

`modifier(relocatable-expression)`

`lo8`

This modifier allows you to use bits 0 through 7 of an address expression as 8 bit relocatable expression.

`hi8`

This modifier allows you to use bits 7 through 15 of an address expression as 8 bit relocatable expression. This is useful with, for example, the AVR ‘`ldi`’ instruction and ‘`lo8`’ modifier.

For example

```
ldi r26, lo8(sym+10)
```

```
ldi r27, hi8(sym+10)
```

hh8

This modifier allows you to use bits 16 through 23 of an address expression as 8 bit relocatable expression. Also, can be useful for loading 32 bit constants.

hlo8

Synonym of ‘**hh8**’.

hhi8

This modifier allows you to use bits 24 through 31 of an expression as 8 bit expression. This is useful with, for example, the AVR ‘**ldi**’ instruction and ‘**lo8**’, ‘**hi8**’, ‘**hlo8**’, ‘**hhi8**’, modifier.

For example

```
ldi r26, lo8(285774925)
ldi r27, hi8(285774925)
ldi r28, hlo8(285774925)
ldi r29, hhi8(285774925)
; r29,r28,r27,r26 = 285774925
```

pm_lo8

This modifier allows you to use bits 0 through 7 of an address expression as 8 bit relocatable expression. This modifier useful for addressing data or code from Flash/Program memory. The using of ‘**pm_lo8**’ similar to ‘**lo8**’.

pm_hi8

This modifier allows you to use bits 8 through 15 of an address expression as 8 bit relocatable expression. This modifier useful for addressing data or code from Flash/Program memory.

pm_hh8

This modifier allows you to use bits 15 through 23 of an address expression as 8 bit relocatable expression. This modifier useful for addressing data or code from Flash/Program memory.

9.4.3 Opcodes

For detailed information on the AVR machine instruction set, see www.atmel.com/products/AVR.

as implements all the standard AVR opcodes. The following table summarizes the AVR opcodes, and their arguments.

Legend:

| | |
|----------|---|
| r | any register |
| d | ‘ldi’ register (r16-r31) |
| v | ‘movw’ even register (r0, r2, ..., r28, r30) |
| a | ‘fmul’ register (r16-r23) |
| w | ‘adiw’ register (r24,r26,r28,r30) |
| e | pointer registers (X,Y,Z) |
| b | base pointer register and displacement ([YZ]+disp) |
| z | Z pointer register (for [e]lpm Rd,Z[+]) |
| M | immediate value from 0 to 255 |
| n | immediate value from 0 to 255 (n = ~M). Relocation impossible |
| s | immediate value from 0 to 7 |
| P | Port address value from 0 to 63. (in, out) |
| p | Port address value from 0 to 31. (cbi, sbi, sbic, sbis) |

K immediate value from 0 to 63 (used in 'adiw', 'sbiw')
 i immediate value
 l signed pc relative offset from -64 to 63
 L signed pc relative offset from -2048 to 2047
 h absolute code address (call, jmp)
 S immediate value from 0 to 7 ($S = s \ll 4$)
 ? use this opcode entry if no parameters, else use next opcode entry

```

1001010010001000    clc
1001010011011000    clh
1001010011111000    cli
1001010010101000    cln
1001010011001000    cls
1001010011101000    clt
1001010010111000    clv
1001010010011000    clz
1001010000001000    sec
1001010001011000    seh
1001010001111000    sei
1001010000101000    sen
1001010001001000    ses
1001010001101000    set
1001010000111000    sev
1001010000011000    sez
100101001SSS1000    bclr    S
100101000SSS1000    bset    S
1001010100001001    icall
1001010000001001    ijmp
1001010111001000    lpm      ?
1001000dddd010+    lpm      r,z
1001010111011000    elpm     ?
1001000dddd011+    elpm     r,z
0000000000000000    nop
1001010100001000    ret
1001010100011000    reti
1001010110001000    sleep
1001010110011000    break
1001010110101000    wdr
1001010111101000    spm
000111rddddrrrr    adc      r,r
000011rddddrrrr    add      r,r
001000rddddrrrr    and      r,r
000101rddddrrrr    cp       r,r
000001rddddrrrr    cpc      r,r
000100rddddrrrr    cpse     r,r
001001rddddrrrr    eor      r,r
001011rddddrrrr    mov      r,r
100111rddddrrrr    mul      r,r
001010rddddrrrr    or       r,r
000010rddddrrrr    sbc      r,r
000110rddddrrrr    sub      r,r
001001rddddrrrr    clr      r
000011rddddrrrr    lsl      r
000111rddddrrrr    rol      r
001000rddddrrrr    tst      r
0111KKKKdddKKKK    andi     d,M
0111KKKKdddKKKK    cbr      d,n
1110KKKKdddKKKK    ldi      d,M
  
```

| | | |
|------------------|--------|-----|
| 11101111dddd1111 | ser | d |
| 0110KKKKddddKKKK | ori | d,M |
| 0110KKKKddddKKKK | sbr | d,M |
| 0011KKKKddddKKKK | cpi | d,M |
| 0100KKKKddddKKKK | sbc | d,M |
| 0101KKKKddddKKKK | subi | d,M |
| 1111110rrrrr0sss | sbr | r,s |
| 1111111rrrrr0sss | sbrs | r,s |
| 1111100dddd0sss | bld | r,s |
| 1111101dddd0sss | bst | r,s |
| 10110PPddddPPPP | in | r,P |
| 10111PPrrrrPPPP | out | P,r |
| 10010110KKddKKKK | adiw | w,K |
| 10010111KKddKKKK | sbiw | w,K |
| 10011000pppppsss | cbi | p,s |
| 10011010pppppsss | sbi | p,s |
| 10011001pppppsss | sbic | p,s |
| 10011011pppppsss | sbis | p,s |
| 111101111111000 | brcc | l |
| 1111001111111000 | brcs | l |
| 1111001111111001 | bre | l |
| 1111011111111100 | brge | l |
| 1111011111111101 | brhc | l |
| 1111001111111101 | brhs | l |
| 1111011111111111 | brid | l |
| 1111001111111111 | brie | l |
| 1111001111111000 | brlo | l |
| 1111001111111100 | brlt | l |
| 1111001111111010 | brmi | l |
| 1111011111111001 | brne | l |
| 1111011111111010 | brpl | l |
| 1111011111111000 | brsh | l |
| 1111011111111110 | brtc | l |
| 1111001111111110 | brts | l |
| 1111011111111011 | brvc | l |
| 1111001111111011 | brvs | l |
| 1111011111111sss | brbc | s,l |
| 1111001111111sss | brbs | s,l |
| 1101LLLLLLLLLLLL | rcall | L |
| 1100LLLLLLLLLLLL | rjmp | L |
| 1001010hhhhh111h | call | h |
| 1001010hhhhh110h | jmp | h |
| 1001010rrrrr0101 | asr | r |
| 1001010rrrrr0000 | com | r |
| 1001010rrrrr1010 | dec | r |
| 1001010rrrrr0011 | inc | r |
| 1001010rrrrr0110 | lsr | r |
| 1001010rrrrr0001 | neg | r |
| 1001000rrrrr1111 | pop | r |
| 1001001rrrrr1111 | push | r |
| 1001010rrrrr0111 | ror | r |
| 1001010rrrrr0010 | swap | r |
| 00000001ddddrrrr | movw | v,v |
| 00000010ddddrrrr | mul | d,d |
| 000000110ddd0rrr | mulsu | a,a |
| 000000110ddd1rrr | fmul | a,a |
| 000000111ddd0rrr | fmuls | a,a |
| 000000111ddd1rrr | fmulsu | a,a |

| | | |
|------------------|--------|-----|
| 1001001dddd0000 | sts | i,r |
| 1001000dddd0000 | lds | r,i |
| 10o0oo0ddddbooo | ldd | r,b |
| 100!000ddddee-+ | ld | r,e |
| 10o0oo1rrrrrbooo | std | b,r |
| 100!001rrrrree-+ | st | e,r |
| 1001010100011001 | eicall | |
| 1001010000011001 | eijmp | |

9.5 Blackfin Dependent Features

9.5.1 Options

`-mcpu=processor``[-sirevision]`

This option specifies the target processor. The optional *sirevision* is not used in assembler. It's here such that GCC can easily pass down its `-mcpu=` option. The assembler will issue an error message if an attempt is made to assemble an instruction which will not execute on the target processor. The following processor names are recognized: bf504, bf506, bf512, bf514, bf516, bf518, bf522, bf523, bf524, bf525, bf526, bf527, bf531, bf532, bf533, bf534, bf535 (not implemented yet), bf536, bf537, bf538, bf539, bf542, bf542m, bf544, bf544m, bf547, bf547m, bf548, bf548m, bf549, bf549m, bf561, and bf592.

`-mfdpic` Assemble for the FDPIC ABI.

`-mno-fdpic`

`-mnopic` Disable `-mfdpic`.

9.5.2 Syntax

Special Characters

Assembler input is free format and may appear anywhere on the line. One instruction may extend across multiple lines or more than one instruction may appear on the same line. White space (space, tab, comments or newline) may appear anywhere between tokens. A token must not have embedded spaces. Tokens include numbers, register names, keywords, user identifiers, and also some multicharacter special symbols like `"+="`, `"/**"` or `"||"`.

Instruction Delimiting

A semicolon must terminate every instruction. Sometimes a complete instruction will consist of more than one operation. There are two cases where this occurs. The first is when two general operations are combined. Normally a comma separates the different parts, as in

```
a0= r3.h * r2.l, a1 = r3.l * r2.h ;
```

The second case occurs when a general instruction is combined with one or two memory references for joint issue. The latter portions are set off by a `"||"` token.

```
a0 = r3.h * r2.l || r1 = [p3++] || r4 = [i2++];
```

Register Names

The assembler treats register names and instruction keywords in a case insensitive manner. User identifiers are case sensitive. Thus, R3.l, R3.L, r3.l and r3.L are all equivalent input to the assembler.

Register names are reserved and may not be used as program identifiers.

Some operations (such as "Move Register") require a register pair. Register pairs are always data registers and are denoted using a colon, eg., R3:2. The larger number must be written first. Note that the hardware only supports odd-even pairs, eg., R7:6, R5:4, R3:2, and R1:0.

Some instructions (such as `-SP` (Push Multiple)) require a group of adjacent registers. Adjacent registers are denoted in the syntax by the range enclosed in parentheses and separated by a colon, eg., `(R7:3)`. Again, the larger number appears first.

Portions of a particular register may be individually specified. This is written with a dot (".") following the register name and then a letter denoting the desired portion. For 32-bit registers, ".H" denotes the most significant ("High") portion. ".L" denotes the least-significant portion. The subdivisions of the 40-bit registers are described later.

Accumulators

The set of 40-bit registers A1 and A0 that normally contain data that is being manipulated. Each accumulator can be accessed in four ways.

one 40-bit register

The register will be referred to as A1 or A0.

one 32-bit register

The registers are designated as A1.W or A0.W.

two 16-bit registers

The registers are designated as A1.H, A1.L, A0.H or A0.L.

one 8-bit register

The registers are designated as A1.X or A0.X for the bits that extend beyond bit 31.

Data Registers

The set of 32-bit registers (R0, R1, R2, R3, R4, R5, R6 and R7) that normally contain data for manipulation. These are abbreviated as D-register or Dreg. Data registers can be accessed as 32-bit registers or as two independent 16-bit registers. The least significant 16 bits of each register is called the "low" half and is designated with ".L" following the register name. The most significant 16 bits are called the "high" half and is designated with ".H" following the name.

R7.L, r2.h, r4.L, R0.H

Pointer Registers

The set of 32-bit registers (P0, P1, P2, P3, P4, P5, SP and FP) that normally contain byte addresses of data structures. These are abbreviated as P-register or Preg.

p2, p5, fp, sp

Stack Pointer SP

The stack pointer contains the 32-bit address of the last occupied byte location in the stack. The stack grows by decrementing the stack pointer.

Frame Pointer FP

The frame pointer contains the 32-bit address of the previous frame pointer in the stack. It is located at the top of a frame.

Loop Top LT0 and LT1. These registers contain the 32-bit address of the top of a zero overhead loop.

Loop Count

LC0 and LC1. These registers contain the 32-bit counter of the zero overhead loop executions.

Loop Bottom

LB0 and LB1. These registers contain the 32-bit address of the bottom of a zero overhead loop.

Index Registers

The set of 32-bit registers (I0, I1, I2, I3) that normally contain byte addresses of data structures. Abbreviated I-register or Ireg.

Modify Registers

The set of 32-bit registers (M0, M1, M2, M3) that normally contain offset values that are added and subtracted to one of the index registers. Abbreviated as Mreg.

Length Registers

The set of 32-bit registers (L0, L1, L2, L3) that normally contain the length in bytes of the circular buffer. Abbreviated as Lreg. Clear the Lreg to disable circular addressing for the corresponding Ireg.

Base Registers

The set of 32-bit registers (B0, B1, B2, B3) that normally contain the base address in bytes of the circular buffer. Abbreviated as Breg.

Floating Point

The Blackfin family has no hardware floating point but the `.float` directive generates ieee floating point numbers for use with software floating point libraries.

Blackfin Opcodes

For detailed information on the Blackfin machine instruction set, see the Blackfin(r) Processor Instruction Set Reference.

9.5.3 Directives

The following directives are provided for compatibility with the VDSP assembler.

| | |
|---------------------|---|
| <code>.byte2</code> | Initializes a four byte data object. |
| <code>.byte4</code> | Initializes a two byte data object. |
| <code>.db</code> | TBD |
| <code>.dd</code> | TBD |
| <code>.dw</code> | TBD |
| <code>.var</code> | Define and initialize a 32 bit data object. |

9.6 CR16 Dependent Features

9.6.1 CR16 Operand Qualifiers

The National Semiconductor CR16 target of `as` has a few machine dependent operand qualifiers.

Operand expression type qualifier is an optional field in the instruction operand, to determines the type of the expression field of an operand. The `@` is required. CR16 architecture uses one of the following expression qualifiers:

- `s` - Specifies expression operand type as small
- `m` - Specifies expression operand type as medium
- `l` - Specifies expression operand type as large
- `c` - Specifies the CR16 Assembler generates a relocation entry for the operand, where `pc` has implied bit, the expression is adjusted accordingly. The linker uses the relocation entry to update the operand address at link time.
- `got/GOT` - Specifies the CR16 Assembler generates a relocation entry for the operand, offset from Global Offset Table. The linker uses this relocation entry to update the operand address at link time
- `cgot/cGOT` - Specifies the CompactRISC Assembler generates a relocation entry for the operand, where `pc` has implied bit, the expression is adjusted accordingly. The linker uses the relocation entry to update the operand address at link time.

CR16 target operand qualifiers and its size (in bits):

- 'Immediate Operand'
 - `s` — 4 bits
 - ' - `m` — 16 bits, for `movb` and `movw` instructions.
 - ' - `m` — 20 bits, `movd` instructions.
 - ' - `l` — 32 bits
- 'Absolute Operand'
 - `s` — Illegal specifier for this operand.
 - ' - `m` — 20 bits, `movd` instructions.
- 'Displacement Operand'
 - `s` — 8 bits
 - ' - `m` — 16 bits
 - ' - `l` — 24 bits

For example:

```
1  movw $_myfun@c,r1
```

This loads the address of `_myfun`, shifted right by 1, into `r1`.

```
2  movd $_myfun@c,(r2,r1)
```

This loads the address of `_myfun`, shifted right by 1, into register-pair `r2-r1`.

```
3  _myfun_ptr:
    .long _myfun@c
    loadadd _myfun_ptr, (r1,r0)
    jal (r1,r0)
```

This `.long` directive, the address of `_myfunc`, shifted right by 1 at link time.

```
4  loadadd _data1@GOT(r12), (r1,r0)
```

This loads the address of `_data1`, into global offset table (ie GOT) and its offset value from GOT loads into register-pair `r2-r1`.

```
5  loadadd _myfunc@cGOT(r12), (r1,r0)
```

This loads the address of `_myfun`, shifted right by 1, into global offset table (ie GOT) and its offset value from GOT loads into register-pair `r1-r0`.

9.7 CRIS Dependent Features

9.7.1 Command-line Options

The CRIS version of `as` has these machine-dependent command-line options.

The format of the generated object files can be either ELF or a.out, specified by the command-line options `--emulation=crisaout` and `--emulation=criself`. The default is ELF (criself), unless `as` has been configured specifically for a.out by using the configuration name `cris-axis-aout`.

There are two different link-incompatible ELF object file variants for CRIS, for use in environments where symbols are expected to be prefixed by a leading `'_'` character and for environments without such a symbol prefix. The variant used for GNU/Linux port has no symbol prefix. Which variant to produce is specified by either of the options `--underscore` and `--no-underscore`. The default is `--underscore`. Since symbols in CRIS a.out objects are expected to have a `'_'` prefix, specifying `--no-underscore` when generating a.out objects is an error. Besides the object format difference, the effect of this option is to parse register names differently (see [\[crisnous\]](#), page 120). The `--no-underscore` option makes a `'$'` register prefix mandatory.

The option `--pic` must be passed to `as` in order to recognize the symbol syntax used for ELF (SVR4 PIC) position-independent-code (see [\[crispic\]](#), page 119). This will also affect expansion of instructions. The expansion with `--pic` will use PC-relative rather than (slightly faster) absolute addresses in those expansions.

The option `--march=architecture` specifies the recognized instruction set and recognized register names. It also controls the architecture type of the object file. Valid values for *architecture* are:

| | |
|-----------------------------|---|
| <code>v0_v10</code> | All instructions and register names for any architecture variant in the set <code>v0...v10</code> are recognized. This is the default if the target is configured as <code>cris-*</code> . |
| <code>v10</code> | Only instructions and register names for CRIS v10 (as found in ETRAX 100 LX) are recognized. This is the default if the target is configured as <code>crisv10-*</code> . |
| <code>v32</code> | Only instructions and register names for CRIS v32 (code name Guinness) are recognized. This is the default if the target is configured as <code>crisv32-*</code> . This value implies <code>--no-mul-bug-abort</code> . (A subsequent <code>--mul-bug-abort</code> will turn it back on.) |
| <code>common_v10_v32</code> | Only instructions with register names and addressing modes with opcodes common to the v10 and v32 are recognized. |

When `-N` is specified, `as` will emit a warning when a 16-bit branch instruction is expanded into a 32-bit multiple-instruction construct (see [Section 9.7.2 \[CRIS-Expand\]](#), page 118).

Some versions of the CRIS v10, for example in the Etrax 100 LX, contain a bug that causes destabilizing memory accesses when a multiply instruction is executed with certain values in the first operand just before a cache-miss. When the `--mul-bug-abort` command line option is active (the default value), `as` will refuse to assemble a file containing a multiply instruction at a dangerous offset, one that could be the last on a cache-line, or is in a

section with insufficient alignment. This placement checking does not catch any case where the multiply instruction is dangerously placed because it is located in a delay-slot. The ‘`--mul-bug-abort`’ command line option turns off the checking.

9.7.2 Instruction expansion

`as` will silently choose an instruction that fits the operand size for ‘`[register+constant]`’ operands. For example, the offset 127 in `move.d [r3+127],r4` fits in an instruction using a signed-byte offset. Similarly, `move.d [r2+32767],r1` will generate an instruction using a 16-bit offset. For symbolic expressions and constants that do not fit in 16 bits including the sign bit, a 32-bit offset is generated.

For branches, `as` will expand from a 16-bit branch instruction into a sequence of instructions that can reach a full 32-bit address. Since this does not correspond to a single instruction, such expansions can optionally be warned about. See [Section 9.7.1 \[CRIS-Opts\]](#), [page 117](#).

If the operand is found to fit the range, a `lapc` mnemonic will translate to a `lapcq` instruction. Use `lapc.d` to force the 32-bit `lapc` instruction.

Similarly, the `addo` mnemonic will translate to the shortest fitting instruction of `addoq`, `addo.w` and `addo.d`, when used with a operand that is a constant known at assembly time.

9.7.3 Symbols

Some symbols are defined by the assembler. They’re intended to be used in conditional assembly, for example:

```
.if ..asm.arch.cris.v32
code for CRIS v32
.elseif ..asm.arch.cris.common_v10_v32
code common to CRIS v32 and CRIS v10
.elseif ..asm.arch.cris.v10 | ..asm.arch.cris.any_v0_v10
code for v10
.else
.error "Code needs to be added here."
.endif
```

These symbols are defined in the assembler, reflecting command-line options, either when specified or the default. They are always defined, to 0 or 1.

```
..asm.arch.cris.any_v0_v10
```

This symbol is non-zero when ‘`--march=v0_v10`’ is specified or the default.

```
..asm.arch.cris.common_v10_v32
```

Set according to the option ‘`--march=common_v10_v32`’.

```
..asm.arch.cris.v10
```

Reflects the option ‘`--march=v10`’.

```
..asm.arch.cris.v32
```

Corresponds to ‘`--march=v10`’.

Speaking of symbols, when a symbol is used in code, it can have a suffix modifying its value for use in position-independent code. See [Section 9.7.4.2 \[CRIS-Pic\]](#), [page 119](#).

9.7.4 Syntax

There are different aspects of the CRIS assembly syntax.

9.7.4.1 Special Characters

The character ‘#’ is a line comment character. It starts a comment if and only if it is placed at the beginning of a line.

A ‘;’ character starts a comment anywhere on the line, causing all characters up to the end of the line to be ignored.

A ‘@’ character is handled as a line separator equivalent to a logical new-line character (except in a comment), so separate instructions can be specified on a single line.

9.7.4.2 Symbols in position-independent code

When generating position-independent code (SVR4 PIC) for use in `cris-axis-linux-gnu` or `crisv32-axis-linux-gnu` shared libraries, symbol suffixes are used to specify what kind of run-time symbol lookup will be used, expressed in the object as different *relocation types*. Usually, all absolute symbol values must be located in a table, the *global offset table*, leaving the code position-independent; independent of values of global symbols and independent of the address of the code. The suffix modifies the value of the symbol, into for example an index into the global offset table where the real symbol value is entered, or a PC-relative value, or a value relative to the start of the global offset table. All symbol suffixes start with the character ‘:’ (omitted in the list below). Every symbol use in code or a read-only section must therefore have a PIC suffix to enable a useful shared library to be created. Usually, these constructs must not be used with an additive constant offset as is usually allowed, i.e. no `symbol + 4` is allowed. This restriction is checked at link-time, not at assembly-time.

GOT

Attaching this suffix to a symbol in an instruction causes the symbol to be entered into the global offset table. The value is a 32-bit index for that symbol into the global offset table. The name of the corresponding relocation is ‘`R_CRIS_32_GOT`’. Example: `move.d [$r0+extsym:GOT], $r9`

GOT16

Same as for ‘GOT’, but the value is a 16-bit index into the global offset table. The corresponding relocation is ‘`R_CRIS_16_GOT`’. Example: `move.d [$r0+asymbol:GOT16], $r10`

PLT

This suffix is used for function symbols. It causes a *procedure linkage table*, an array of code stubs, to be created at the time the shared object is created or linked against, together with a global offset table entry. The value is a pc-relative offset to the corresponding stub code in the procedure linkage table. This arrangement causes the run-time symbol resolver to be called to look up and set the value of the symbol the first time the function is called (at latest; depending environment variables). It is only safe to leave the symbol unresolved this way if all references are function calls. The name of the relocation is ‘`R_CRIS_32_PLT_PCREL`’. Example: `add.d ffname:PLT, $pc`

PLTG

Like PLT, but the value is relative to the beginning of the global offset table. The relocation is ‘R_CRIS_32_PLT_GOTREL’. Example: `move.d fnname:PLTG,$r3`

GOTPLT

Similar to ‘PLT’, but the value of the symbol is a 32-bit index into the global offset table. This is somewhat of a mix between the effect of the ‘GOT’ and the ‘PLT’ suffix; the difference to ‘GOT’ is that there will be a procedure linkage table entry created, and that the symbol is assumed to be a function entry and will be resolved by the run-time resolver as with ‘PLT’. The relocation is ‘R_CRIS_32_GOTPLT’. Example: `jsr [$r0+fnname:GOTPLT]`

GOTPLT16

A variant of ‘GOTPLT’ giving a 16-bit value. Its relocation name is ‘R_CRIS_16_GOTPLT’. Example: `jsr [$r0+fnname:GOTPLT16]`

GOTOFF

This suffix must only be attached to a local symbol, but may be used in an expression adding an offset. The value is the address of the symbol relative to the start of the global offset table. The relocation name is ‘R_CRIS_32_GOTREL’. Example: `move.d [$r0+localsym:GOTOFF],r3`

9.7.4.3 Register names

A ‘\$’ character may always prefix a general or special register name in an instruction operand but is mandatory when the option ‘--no-underscore’ is specified or when the `.syntax register_prefix` directive is in effect (see [crisnous], page 120). Register names are case-insensitive.

9.7.4.4 Assembler Directives

There are a few CRIS-specific pseudo-directives in addition to the generic ones. See [Chapter 7 \[Pseudo Ops\], page 47](#). Constants emitted by pseudo-directives are in little-endian order for CRIS. There is no support for floating-point-specific directives for CRIS.

`.dword` EXPRESSIONS

The `.dword` directive is a synonym for `.int`, expecting zero or more EXPRESSIONS, separated by commas. For each expression, a 32-bit little-endian constant is emitted.

`.syntax` ARGUMENT

The `.syntax` directive takes as *ARGUMENT* one of the following case-sensitive choices.

`no_register_prefix`

The `.syntax no_register_prefix` directive makes a ‘\$’ character prefix on all registers optional. It overrides a previous setting, including the corresponding effect of the option ‘--no-underscore’. If this directive is used when ordinary symbols do not have a ‘_’ character prefix, care must be taken to avoid ambiguities whether an operand is a register or a symbol; using symbols with names the same as general or special registers then invoke undefined behavior.

register_prefix

This directive makes a '\$' character prefix on all registers mandatory. It overrides a previous setting, including the corresponding effect of the option '--underscore'.

leading_underscore

This is an assertion directive, emitting an error if the '--no-underscore' option is in effect.

no_leading_underscore

This is the opposite of the `.syntax leading_underscore` directive and emits an error if the option '--underscore' is in effect.

.arch ARGUMENT

This is an assertion directive, giving an error if the specified *ARGUMENT* is not the same as the specified or default value for the '--march=*architecture*' option (see [\[march-option\]](#), page 117).

9.8 D10V Dependent Features

9.8.1 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.

`‘--gstabs-packing’`
`‘--no-gstabs-packing’`
 `as` packs adjacent short instructions into a single packed instruction. `‘--no-gstabs-packing’` turns instruction packing off if `‘--gstabs’` is specified as well; `‘--gstabs-packing’` (the default) turns instruction packing on even when `‘--gstabs’` is specified.

9.8.2 Syntax

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

9.8.2.1 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.

9.8.2.2 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.

9.8.2.3 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.

9.8.2.4 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 |

9.8.2.5 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.

| | |
|--------------------------|--|
| <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) |

9.8.2.6 @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
```

9.8.3 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.

9.8.4 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

9.9 D30V Dependent Features

9.9.1 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.

9.9.2 Syntax

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

9.9.2.1 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.

9.9.2.2 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.

9.9.2.3 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 ‘`-O`’ 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.

9.9.2.4 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:

- ‘/tx’ Execute the instruction if flag f0 is true.

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

9.9.2.5 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) |

9.9.2.6 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.

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

9.9.3 Floating Point

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

9.9.4 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

9.10 H8/300 Dependent Features

9.10.1 Options

The Renesas H8/300 version of **as** has one machine-dependent option:

-h-tick-hex

Support H'00 style hex constants in addition to 0x00 style.

9.10.2 Syntax

9.10.2.1 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.

9.10.2.2 Register Names

You can use predefined symbols of the form '**rn***h*' and '**rn***l*' 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**.

9.10.2.3 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.

9.10.3 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.

9.10.4 H8/300 Machine Directives

`as` has the following machine-dependent directives 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.
- `.h8300s` Recognize and emit additional instructions for the H8S variant, and also make `.int` emit 32-bit numbers rather than the usual (16-bit) for the H8/300 family.
- `.h8300hn` Recognize and emit additional instructions for the H8/300H variant in normal mode, and also make `.int` emit 32-bit numbers rather than the usual (16-bit) for the H8/300 family.
- `.h8300sn` Recognize and emit additional instructions for the H8S variant in normal mode, 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.

9.10.5 Opcodes

For detailed information on the H8/300 machine instruction set, see *H8/300 Series Programming Manual*. For information specific to the H8/300H, see *H8/300H Series Programming Manual* (Renesas).

`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.

9.11 HPPA Dependent Features

9.11.1 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.

9.11.2 Options

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

9.11.3 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.

9.11.4 Floating Point

The HPPA family uses IEEE floating-point numbers.

9.11.5 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 **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 29](#), 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), `'comdat'` (subsection is comdat), `'common'` (subsection is common block), `'dup_comm'` (subsection 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.

`'comdat'`, `'common'` and `'dup_comm'` can be used to implement various flavors of one-only support when using the SOM linker. The SOM linker only supports specific combinations of these flags. The details are not documented. A brief description is provided here.

`'comdat'` provides a form of linkonce support. It is useful for both code and data subspaces. A `'comdat'` subspace has a key symbol marked by the `'is_comdat'` flag or `'ST_COMDAT'`. Only the first subspace for any given key is selected. The key symbol becomes universal in shared links. This is similar to the behavior of `'secondary_def'` symbols.

`'common'` provides Fortran named common support. It is only useful for data subspaces. Symbols with the flag `'is_common'` retain this flag in shared links. Referencing a `'is_common'` symbol in a shared library from outside the library doesn't work. Thus, `'is_common'` symbols must be output whenever they are needed.

`'common'` and `'dup_comm'` together provide Cobol common support. The subspaces in this case must all be the same length. Otherwise, this support is similar to the Fortran common support.

`'dup_comm'` by itself provides a type of one-only support for code. Only the first `'dup_comm'` subspace is selected. There is a rather complex algorithm to compare subspaces. Code symbols marked with the `'dup_common'` flag are hidden. This support was intended for "C++ duplicate inlines".

A simplified technique is used to mark the flags of symbols based on the flags of their subspace. A symbol with the scope `SS_UNIVERSAL` and type `ST_ENTRY`, `ST_CODE` or `ST_DATA` is marked with the corresponding settings of `'comdat'`, `'common'` and `'dup_comm'` from the subspace, respectively. This avoids having to introduce additional directives to mark these symbols. The HP assembler sets `'is_common'` from `'common'`. However, it doesn't set the `'dup_common'` from `'dup_comm'`. It doesn't have `'comdat'` support.

`.version "str"`

Write *str* as version identifier in object code.

9.11.6 Opcodes

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

9.12 ESA/390 Dependent Features

9.12.1 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.

9.12.2 Options

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

9.12.3 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` or `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'
```



```
0 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).

9.12.4 Floating Point

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

9.12.5 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.

9.12.6 Opcodes

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

9.13 80386 Dependent Features

The i386 version of **as** supports both the original Intel 386 architecture in both 16 and 32-bit mode as well as AMD x86-64 architecture extending the Intel architecture to 64-bits.

9.13.1 Options

The i386 version of **as** has a few machine dependent options:

--32 | --64

Select the word size, either 32 bits or 64 bits. Selecting 32-bit implies Intel i386 architecture, while 64-bit implies AMD x86-64 architecture.

These options are only available with the ELF object file format, and require that the necessary BFD support has been included (on a 32-bit platform you have to add `-enable-64-bit-bfd` to configure enable 64-bit usage and use x86-64 as target platform).

-n By default, x86 GAS replaces multiple `nop` instructions used for alignment within code sections with multi-byte `nop` instructions such as `leal 0(%esi,1),%esi`. This switch disables the optimization.

--divide On SVR4-derived platforms, the character `'/'` is treated as a comment character, which means that it cannot be used in expressions. The `'--divide'` option turns `'/'` into a normal character. This does not disable `'/'` at the beginning of a line starting a comment, or affect using `'#'` for starting a comment.

-march=CPU[+EXTENSION...]

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. The following processor names are recognized: `i8086`, `i186`, `i286`, `i386`, `i486`, `i586`, `i686`, `pentium`, `pentiumpro`, `pentiumii`, `pentiumiii`, `pentium4`, `prescott`, `nocona`, `core`, `core2`, `corei7`, `l1om`, `k6`, `k6_2`, `athlon`, `opteron`, `k8`, `amdfam10`, `bdver1`, `generic32` and `generic64`.

In addition to the basic instruction set, the assembler can be told to accept various extension mnemonics. For example, `-march=i686+sse4+vmx` extends `i686` with `sse4` and `vmx`. The following extensions are currently supported: `8087`, `287`, `387`, `no87`, `mmx`, `nommx`, `sse`, `sse2`, `sse3`, `ssse3`, `sse4.1`, `sse4.2`, `sse4`, `nosse`, `avx`, `noavx`, `vmx`, `smx`, `xsave`, `xsaveopt`, `aes`, `pclmul`, `fsgsbase`, `rdrnd`, `f16c`, `fma`, `movbe`, `ept`, `clflush`, `lwp`, `fma4`, `xop`, `syscall`, `rdtscp`, `3dnow`, `3dnowa`, `sse4a`, `sse5`, `svme`, `abm` and `padlock`. Note that rather than extending a basic instruction set, the extension mnemonics starting with `no` revoke the respective functionality.

When the `.arch` directive is used with `'-march'`, the `.arch` directive will take precedent.

-mtune=CPU

This option specifies a processor to optimize for. When used in conjunction with the `'-march'` option, only instructions of the processor specified by the `'-march'` option will be generated.

Valid `CPU` values are identical to the processor list of `'-march=CPU'`.

`-msse2avx`

This option specifies that the assembler should encode SSE instructions with VEX prefix.

`-msse-check=none`

`-msse-check=warning`

`-msse-check=error`

These options control if the assembler should check SSE instructions. ‘`-msse-check=none`’ will make the assembler not to check SSE instructions, which is the default. ‘`-msse-check=warning`’ will make the assembler issue a warning for any SSE instruction. ‘`-msse-check=error`’ will make the assembler issue an error for any SSE instruction.

`-mavxscalar=128`

`-mavxscalar=256`

This options control how the assembler should encode scalar AVX instructions. ‘`-mavxscalar=128`’ will encode scalar AVX instructions with 128bit vector length, which is the default. ‘`-mavxscalar=256`’ will encode scalar AVX instructions with 256bit vector length.

`-mmnemonic=att`

`-mmnemonic=intel`

This option specifies instruction mnemonic for matching instructions. The `.att_mnemonic` and `.intel_mnemonic` directives will take precedent.

`-msyntax=att`

`-msyntax=intel`

This option specifies instruction syntax when processing instructions. The `.att_syntax` and `.intel_syntax` directives will take precedent.

`-mnaked-reg`

This option specifies that registers don’t require a ‘%’ prefix. The `.att_syntax` and `.intel_syntax` directives will take precedent.

9.13.2 x86 specific Directives

`.lcomm symbol , length[, alignment]`

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. Since *symbol* is not declared global, it is normally not visible to `ld`. The optional third parameter, *alignment*, specifies the desired alignment of the symbol in the bss section.

This directive is only available for COFF based x86 targets.

9.13.3 AT&T Syntax versus Intel Syntax

`as` now supports assembly using Intel assembler syntax. `.intel_syntax` selects Intel mode, and `.att_syntax` switches back to the usual AT&T mode for compatibility with the output of `gcc`. Either of these directives may have an optional argument, `prefix`, or `noprefix` specifying whether registers require a ‘%’ prefix. AT&T System V/386 assembler syntax 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 'bound', 'invlpga', and instructions with 2 immediate operands, such as the 'enter' instruction, do *not* have reversed order. [Section 9.13.14 \[i386-Bugs\], page 149](#).
- In AT&T syntax the size of memory operands is determined from the last character of the instruction mnemonic. Mnemonic suffixes of 'b', 'w', 'l' and 'q' specify byte (8-bit), word (16-bit), long (32-bit) and quadruple word (64-bit) memory references. Intel syntax accomplishes this by prefixing memory operands (*not* the instruction mnemonics) with 'byte ptr', 'word ptr', 'dword ptr' and 'qword ptr'. Thus, Intel 'mov al, byte ptr foo' is 'movb foo, %al' in AT&T syntax.
In 64-bit code, 'movabs' can be used to encode the 'mov' instruction with the 64-bit displacement or immediate operand.
- 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.

9.13.4 Instruction Naming

Instruction mnemonics are suffixed with one character modifiers which specify the size of operands. The letters 'b', 'w', 'l' and 'q' specify byte, word, long and quadruple word 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), 'wl' (from word to long), 'bq' (from byte to quadruple word), 'wq' (from word to quadruple word), and 'lq' (from long to quadruple word).

Different encoding options can be specified via optional mnemonic suffix. ‘.s’ suffix swaps 2 register operands in encoding when moving from one register to another. ‘.d32’ suffix forces 32bit displacement in encoding.

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’,
- ‘cqd’ — sign-extend word in ‘%ax’ to long in ‘%dx:%ax’,
- ‘cdq’ — sign-extend dword in ‘%eax’ to quad in ‘%edx:%eax’,
- ‘cdqe’ — sign-extend dword in ‘%eax’ to quad in ‘%rax’ (x86-64 only),
- ‘cqo’ — sign-extend quad in ‘%rax’ to octuple in ‘%rdx:%rax’ (x86-64 only),

are called ‘cbtw’, ‘cwtl’, ‘cwtd’, ‘cltd’, ‘cltq’, and ‘cqto’ 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.

9.13.5 AT&T Mnemonic versus Intel Mnemonic

`as` supports assembly using Intel mnemonic. `.intel_mnemonic` selects Intel mnemonic with Intel syntax, and `.att_mnemonic` switches back to the usual AT&T mnemonic with AT&T syntax for compatibility with the output of `gcc`. Several x87 instructions, ‘fadd’, ‘fdiv’, ‘fdivp’, ‘fdivr’, ‘fdivrp’, ‘fmul’, ‘fsub’, ‘fsubp’, ‘fsubr’ and ‘fsubrp’, are implemented in AT&T System V/386 assembler with different mnemonics from those in Intel IA32 specification. `gcc` generates those instructions with AT&T mnemonic.

9.13.6 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)’. These registers are overloaded by 8 MMX registers ‘%mm0’, ‘%mm1’, ‘%mm2’, ‘%mm3’, ‘%mm4’, ‘%mm5’, ‘%mm6’ and ‘%mm7’.
- the 8 SSE registers registers ‘%xmm0’, ‘%xmm1’, ‘%xmm2’, ‘%xmm3’, ‘%xmm4’, ‘%xmm5’, ‘%xmm6’ and ‘%xmm7’.

The AMD x86-64 architecture extends the register set by:

- enhancing the 8 32-bit registers to 64-bit: `%rax` (the accumulator), `%rbx`, `%rcx`, `%rdx`, `%rdi`, `%rsi`, `%rbp` (the frame pointer), `%rsp` (the stack pointer)
- the 8 extended registers `%r8`–`%r15`.
- the 8 32-bit low ends of the extended registers: `%r8d`–`%r15d`
- the 8 16-bit low ends of the extended registers: `%r8w`–`%r15w`
- the 8 8-bit low ends of the extended registers: `%r8b`–`%r15b`
- the 4 8-bit registers: `%sil`, `%dil`, `%bpl`, `%spl`.
- the 8 debug registers: `%db8`–`%db15`.
- the 8 SSE registers: `%xmm8`–`%xmm15`.

9.13.7 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 jmp1 *(%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).
- The `'rex'` family of prefixes is used by x86-64 to encode extensions to i386 instruction set. The `'rex'` prefix has four bits — an operand size overwrite (64) used to change operand size from 32-bit to 64-bit and X, Y and Z extensions bits used to extend the register set.

You may write the `'rex'` prefixes directly. The `'rex64xyz'` instruction emits `'rex'` prefix with all the bits set. By omitting the 64, x, y or z you may write other prefixes as well. Normally, there is no need to write the prefixes explicitly, since `gas` will automatically generate them based on the instruction operands.

9.13.8 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, long, or quadruple) with an instruction mnemonic suffix (‘b’, ‘w’, ‘l’ or ‘q’, respectively).

The x86-64 architecture adds an RIP (instruction pointer relative) addressing. This addressing mode is specified by using ‘rip’ as a base register. Only constant offsets are valid. For example:

AT&T: ‘1234(%rip)’, Intel: ‘[rip + 1234]’

Points to the address 1234 bytes past the end of the current instruction.

AT&T: ‘symbol(%rip)’, Intel: ‘[rip + symbol]’

Points to the **symbol** in RIP relative way, this is shorter than the default absolute addressing.

Other addressing modes remain unchanged in x86-64 architecture, except registers used are 64-bit instead of 32-bit.

9.13.9 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 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. (See also see [Section 9.13.15 \[i386-Arch\], page 149](#))

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:
```

9.13.10 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 ‘fildq’ (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’)

9.13.11 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.

9.13.12 AMD’s Lightweight Profiling Instructions

`as` supports AMD’s Lightweight Profiling (LWP) instruction set, available on AMD’s Family 15h (Orochi) processors.

LWP enables applications to collect and manage performance data, and react to performance events. The collection of performance data requires no context switches. LWP runs in the context of a thread and so several counters can be used independently across multiple threads. LWP can be used in both 64-bit and legacy 32-bit modes.

For detailed information on the LWP instruction set, see the *AMD Lightweight Profiling Specification* available at [Lightweight Profiling Specification](#).

9.13.13 Writing 16-bit Code

While `as` normally writes only “pure” 32-bit i386 code or 64-bit x86-64 code depending on the default configuration, 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` to writing 32-bit code with the ‘`.code32`’ directive or 64-bit code with the ‘`.code64`’ 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`’ (i.e., 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.

9.13.14 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)'.

9.13.15 Specifying CPU Architecture

as may be told to assemble for a particular CPU (sub-)architecture with the `.arch cpu_type` directive. This directive enables a warning when gas detects an instruction that is not supported on the CPU specified. The choices for `cpu_type` are:

| | | | |
|--------------|-------------|--------------|------------|
| 'i8086' | 'i186' | 'i286' | 'i386' |
| 'i486' | 'i586' | 'i686' | 'pentium' |
| 'pentiumpro' | 'pentiumii' | 'pentiumiii' | 'pentium4' |
| 'prescott' | 'nocona' | 'core' | 'core2' |
| 'corei7' | 'l1om' | | |
| 'k6' | 'k6_2' | 'athlon' | 'k8' |
| 'amdfam10' | 'bdver1' | | |
| 'generic32' | 'generic64' | | |
| 'mmx' | 'sse' | 'sse2' | 'sse3' |
| 'ssse3' | 'sse4.1' | 'sse4.2' | 'sse4' |
| 'avx' | 'vmx' | 'smx' | 'ept' |
| 'clflush' | 'movbe' | 'xsave' | 'xsaveopt' |
| 'aes' | 'pclmul' | 'fma' | 'fsgsbase' |
| 'rdrnd' | 'f16c' | | |
| '3dnow' | '3dnowa' | 'sse4a' | 'sse5' |
| 'syscall' | 'rdtscp' | 'svme' | 'abm' |
| 'lwp' | 'fma4' | 'xop' | |
| 'padlock' | | | |

Apart from the warning, there are only two other effects on `as` operation; Firstly, if you specify a CPU other than 'i486', then shift by one instructions such as '`sarl $1, %eax`' will automatically use a two byte opcode sequence. The larger three byte opcode sequence is used on the 486 (and when no architecture is specified) because it executes faster on the 486. Note that you can explicitly request the two byte opcode by writing '`sarl %eax`'. Secondly, if you specify 'i8086', 'i186', or 'i286', and '.code16' or '.code16gcc' then byte offset conditional jumps will be promoted when necessary to a two instruction sequence consisting of a conditional jump of the opposite sense around an unconditional jump to the target.

Following the CPU architecture (but not a sub-architecture, which are those starting with a dot), you may specify 'jumps' or 'nojumps' to control automatic promotion of conditional

jumps. ‘jumps’ is the default, and enables jump promotion; All external jumps will be of the long variety, and file-local jumps will be promoted as necessary. (see [Section 9.13.9 \[i386-Jumps\]](#), page 147) ‘nojumps’ leaves external conditional jumps as byte offset jumps, and warns about file-local conditional jumps that **as** promotes. Unconditional jumps are treated as for ‘jumps’.

For example

```
.arch i8086,nojumps
```

9.13.16 Notes

There is some trickery concerning the ‘mul’ and ‘imul’ instructions that deserves mention. The 16-, 32-, 64- and 128-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’.

9.14 Intel i860 Dependent Features

9.14.1 i860 Notes

This is a fairly complete i860 assembler which is compatible with the UNIX System V/860 Release 4 assembler. However, it does not currently support SVR4 PIC (i.e., `@GOT`, `@GOTOFF`, `@PLT`).

Like the SVR4/860 assembler, the output object format is ELF32. Currently, this is the only supported object format. If there is sufficient interest, other formats such as COFF may be implemented.

Both the Intel and AT&T/SVR4 syntaxes are supported, with the latter being the default. One difference is that AT&T syntax requires the '%' prefix on register names while Intel syntax does not. Another difference is in the specification of relocatable expressions. The Intel syntax is `ha%expression` whereas the SVR4 syntax is `[expression]@ha` (and similarly for the "l" and "h" selectors).

9.14.2 i860 Command-line Options

9.14.2.1 SVR4 compatibility options

- V Print assembler version.
- Qy Ignored.
- Qn Ignored.

9.14.2.2 Other options

- EL Select little endian output (this is the default).
- EB Select big endian output. Note that the i860 always reads instructions as little endian data, so this option only effects data and not instructions.
- mwarn-expand
 Emit a warning message if any pseudo-instruction expansions occurred. For example, a `or` instruction with an immediate larger than 16-bits will be expanded into two instructions. This is a very undesirable feature to rely on, so this flag can help detect any code where it happens. One use of it, for instance, has been to find and eliminate any place where `gcc` may emit these pseudo-instructions.
- mxp Enable support for the i860XP instructions and control registers. By default, this option is disabled so that only the base instruction set (i.e., i860XR) is supported.
- mintel-syntax
 The i860 assembler defaults to AT&T/SVR4 syntax. This option enables the Intel syntax.

9.14.3 i860 Machine Directives

- .dual Enter dual instruction mode. While this directive is supported, the preferred way to use dual instruction mode is to explicitly code the dual bit with the `d.` prefix.

- .enddual** Exit dual instruction mode. While this directive is supported, the preferred way to use dual instruction mode is to explicitly code the dual bit with the **d.** prefix.
- .atmp** Change the temporary register used when expanding pseudo operations. The default register is **r31**.

The **.dual**, **.enddual**, and **.atmp** directives are available only in the Intel syntax mode.

Both syntaxes allow for the standard **.align** directive. However, the Intel syntax additionally allows keywords for the alignment parameter: **".align type"**, where 'type' is one of **.short**, **.long**, **.quad**, **.single**, **.double** representing alignments of 2, 4, 16, 4, and 8, respectively.

9.14.4 i860 Opcodes

All of the Intel i860XR and i860XP machine instructions are supported. Please see either *i860 Microprocessor Programmer's Reference Manual* or *i860 Microprocessor Architecture* for more information.

9.14.4.1 Other instruction support (pseudo-instructions)

For compatibility with some other i860 assemblers, a number of pseudo-instructions are supported. While these are supported, they are a very undesirable feature that should be avoided – in particular, when they result in an expansion to multiple actual i860 instructions. Below are the pseudo-instructions that result in expansions.

- Load large immediate into general register:

The pseudo-instruction **mov imm,%rn** (where the immediate does not fit within a signed 16-bit field) will be expanded into:

```
orh large_imm@h,%r0,%rn
or large_imm@l,%rn,%rn
```

- Load/store with relocatable address expression:

For example, the pseudo-instruction **ld.b addr_exp(%rx),%rn** will be expanded into:

```
orh addr_exp@ha,%rx,%r31
ld.l addr_exp@l(%r31),%rn
```

The analogous expansions apply to **ld.x**, **st.x**, **fld.x**, **pfld.x**, **fst.x**, and **pst.x** as well.

- Signed large immediate with add/subtract:

If any of the arithmetic operations **adds**, **addu**, **subs**, **subu** are used with an immediate larger than 16-bits (signed), then they will be expanded. For instance, the pseudo-instruction **adds large_imm,%rx,%rn** expands to:

```
orh large_imm@h,%r0,%r31
or large_imm@l,%r31,%r31
adds %r31,%rx,%rn
```

- Unsigned large immediate with logical operations:

Logical operations (**or**, **andnot**, **or**, **xor**) also result in expansions. The pseudo-instruction **or large_imm,%rx,%rn** results in:

```
orh large_imm@h,%rx,%r31
or large_imm@l,%r31,%rn
```

Similarly for the others, except for **and** which expands to:

```
andnot (-1 - large_imm)@h,%rx,%r31  
andnot (-1 - large_imm)@l,%r31,%rn
```

9.15 Intel 80960 Dependent Features

9.15.1 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 labeled ‘__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.



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'`.

9.15.2 Floating Point

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

9.15.3 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 7.67 \[.lcomm\]](#), page 59.

.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).

9.15.4 i960 Opcodes

All Intel 960 machine instructions are supported; see [Section 9.15.1 \[i960 Command-line Options\]](#), [page 154](#) 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.

9.15.4.1 `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.

9.15.4.2 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</i> | | |
|----------------------|----------------------|--------------------------|
| <i>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> |
| <code>cmpibge</code> | <code>cmpijge</code> | <code>cmpi; bge</code> |
| <code>cmpibl</code> | <code>cmpijl</code> | <code>cmpi; bl</code> |
| <code>cmpible</code> | <code>cmpijle</code> | <code>cmpi; ble</code> |
| <code>cmpibno</code> | <code>cmpijno</code> | <code>cmpi; bno</code> |
| <code>cmpibne</code> | <code>cmpijne</code> | <code>cmpi; bne</code> |
| <code>cmpibo</code> | <code>cmpijo</code> | <code>cmpi; bo</code> |
| <code>cmpobe</code> | <code>cmpoje</code> | <code>cmpo; be</code> |
| <code>cmpobg</code> | <code>cmpojg</code> | <code>cmpo; bg</code> |
| <code>cmpobge</code> | <code>cmpojge</code> | <code>cmpo; bge</code> |
| <code>cmpobl</code> | <code>cmpojl</code> | <code>cmpo; bl</code> |
| <code>cmpoble</code> | <code>cmpojle</code> | <code>cmpo; ble</code> |

`cmpobne` `cmpojne` `cmpo; bne`

9.16 IA-64 Dependent Features

9.16.1 Options

`‘-mconstant-gp’`

This option instructs the assembler to mark the resulting object file as using the “constant GP” model. With this model, it is assumed that the entire program uses a single global pointer (GP) value. Note that this option does not in any fashion affect the machine code emitted by the assembler. All it does is turn on the `EF_IA_64_CONS_GP` flag in the ELF file header.

`‘-mauto-pic’`

This option instructs the assembler to mark the resulting object file as using the “constant GP without function descriptor” data model. This model is like the “constant GP” model, except that it additionally does away with function descriptors. What this means is that the address of a function refers directly to the function’s code entry-point. Normally, such an address would refer to a function descriptor, which contains both the code entry-point and the GP-value needed by the function. Note that this option does not in any fashion affect the machine code emitted by the assembler. All it does is turn on the `EF_IA_64_NOFUNCDESC_CONS_GP` flag in the ELF file header.

`‘-milp32’`

`‘-milp64’`

`‘-mlp64’`

`‘-mp64’`

These options select the data model. The assembler defaults to `-mlp64` (LP64 data model).

`‘-mle’`

`‘-mbe’`

These options select the byte order. The `-mle` option selects little-endian byte order (default) and `-mbe` selects big-endian byte order. Note that IA-64 machine code always uses little-endian byte order.

`‘-mtune=itanium1’`

`‘-mtune=itanium2’`

Tune for a particular IA-64 CPU, *itanium1* or *itanium2*. The default is *itanium2*.

`‘-munwind-check=warning’`

`‘-munwind-check=error’`

These options control what the assembler will do when performing consistency checks on unwind directives. `-munwind-check=warning` will make the assembler issue a warning when an unwind directive check fails. This is the default. `-munwind-check=error` will make the assembler issue an error when an unwind directive check fails.

`‘-mhint.b=ok’`

`‘-mhint.b=warning’`

`‘-mhint.b=error’`

These options control what the assembler will do when the `‘hint.b’` instruction is used. `-mhint.b=ok` will make the assembler accept `‘hint.b’`.

`-mint.b=warning` will make the assembler issue a warning when `'hint.b'` is used. `-mhint.b=error` will make the assembler treat `'hint.b'` as an error, which is the default.

`'-x'`

`'-xexplicit'`

These options turn on dependency violation checking.

`'-xauto'` This option instructs the assembler to automatically insert stop bits where necessary to remove dependency violations. This is the default mode.

`'-xnone'` This option turns off dependency violation checking.

`'-xdebug'` This turns on debug output intended to help tracking down bugs in the dependency violation checker.

`'-xdebugn'`

This is a shortcut for `-xnone -xdebug`.

`'-xdebugx'`

This is a shortcut for `-xexplicit -xdebug`.

9.16.2 Syntax

The assembler syntax closely follows the IA-64 Assembly Language Reference Guide.

9.16.2.1 Special Characters

`'//'` is the line comment token.

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

9.16.2.2 Register Names

The 128 integer registers are referred to as `'rn'`. The 128 floating-point registers are referred to as `'fn'`. The 128 application registers are referred to as `'arn'`. The 128 control registers are referred to as `'crn'`. The 64 one-bit predicate registers are referred to as `'pn'`. The 8 branch registers are referred to as `'bn'`. In addition, the assembler defines a number of aliases: `'gp'` (`'r1'`), `'sp'` (`'r12'`), `'rp'` (`'b0'`), `'ret0'` (`'r8'`), `'ret1'` (`'r9'`), `'ret2'` (`'r10'`), `'ret3'` (`'r9'`), `'fargn'` (`'f8+n'`), and `'fretn'` (`'f8+n'`).

For convenience, the assembler also defines aliases for all named application and control registers. For example, `'ar.bsp'` refers to the register backing store pointer (`'ar17'`). Similarly, `'cr.eoi'` refers to the end-of-interrupt register (`'cr67'`).

9.16.2.3 IA-64 Processor-Status-Register (PSR) Bit Names

The assembler defines bit masks for each of the bits in the IA-64 processor status register. For example, `'psr.ic'` corresponds to a value of 0x2000. These masks are primarily intended for use with the `'ssm'`/`'sum'` and `'rsm'`/`'rum'` instructions, but they can be used anywhere else where an integer constant is expected.

9.16.2.4 Relocations

In addition to the standard IA-64 relocations, the following relocations are implemented by as:

@slotcount(*V*)

Convert the address offset *V* into a slot count. This pseudo function is available only on VMS. The expression *V* must be known at assembly time: it can't reference undefined symbols or symbols in different sections.

9.16.3 Opcodes

For detailed information on the IA-64 machine instruction set, see the IA-64 Assembly Language Reference Guide available at

http://developer.intel.com/design/itanium/arch_spec.htm

9.17 IP2K Dependent Features

9.17.1 IP2K Options

The Uvicom IP2K version of `as` has a few machine dependent options:

`-mip2022ext`

`as` can assemble the extended IP2022 instructions, but it will only do so if this is specifically allowed via this command line option.

`-mip2022` This option restores the assembler's default behaviour of not permitting the extended IP2022 instructions to be assembled.

9.18 LM32 Dependent Features

9.18.1 Options

- `-mmultiply-enabled`
Enable multiply instructions.
- `-mdivide-enabled`
Enable divide instructions.
- `-mbarrel-shift-enabled`
Enable barrel-shift instructions.
- `-msign-extend-enabled`
Enable sign extend instructions.
- `-muser-enabled`
Enable user defined instructions.
- `-micache-enabled`
Enable instruction cache related CSRs.
- `-mdcache-enabled`
Enable data cache related CSRs.
- `-mbreak-enabled`
Enable break instructions.
- `-mall-enabled`
Enable all instructions and CSRs.

9.18.2 Syntax

9.18.2.1 Register Names

LM32 has 32 x 32-bit general purpose registers ‘r0’, ‘r1’, ... ‘r31’.

The following aliases are defined: ‘gp’ - ‘r26’, ‘fp’ - ‘r27’, ‘sp’ - ‘r28’, ‘ra’ - ‘r29’, ‘ea’ - ‘r30’, ‘ba’ - ‘r31’.

LM32 has the following Control and Status Registers (CSRs).

| | |
|-----|----------------------------|
| IE | Interrupt enable. |
| IM | Interrupt mask. |
| IP | Interrupt pending. |
| ICC | Instruction cache control. |
| DCC | Data cache control. |
| CC | Cycle counter. |
| CFG | Configuration. |
| EBA | Exception base address. |
| DC | Debug control. |

| | |
|------|-------------------------------|
| DEBA | Debug exception base address. |
| JTX | JTAG transmit. |
| JRX | JTAG receive. |
| BP0 | Breakpoint 0. |
| BP1 | Breakpoint 1. |
| BP2 | Breakpoint 2. |
| BP3 | Breakpoint 3. |
| WP0 | Watchpoint 0. |
| WP1 | Watchpoint 1. |
| WP2 | Watchpoint 2. |
| WP3 | Watchpoint 3. |

9.18.2.2 Relocatable Expression Modifiers

The assembler supports several modifiers when using relocatable addresses in LM32 instruction operands. The general syntax is the following:

`modifier(relocatable-expression)`

`lo`

This modifier allows you to use bits 0 through 15 of an address expression as 16 bit relocatable expression.

`hi`

This modifier allows you to use bits 16 through 23 of an address expression as 16 bit relocatable expression.

For example

```
ori r4, r4, lo(sym+10)
orhi r4, r4, hi(sym+10)
```

`gp`

This modifier creates a 16-bit relocatable expression that is the offset of the symbol from the global pointer.

```
mva r4, gp(sym)
```

`got`

This modifier places a symbol in the GOT and creates a 16-bit relocatable expression that is the offset into the GOT of this symbol.

```
lw r4, (gp+got(sym))
```

`gotofflo16`

This modifier allows you to use the bits 0 through 15 of an address which is an offset from the GOT.

`gotoffhi16`

This modifier allows you to use the bits 16 through 31 of an address which is an offset from the GOT.

```
orhi r4, r4, gotoffhi16(1sym)
addi r4, r4, gotofflo16(1sym)
```

9.18.3 Opcodes

For detailed information on the LM32 machine instruction set, see <http://www.latticesemi.com/products/i> as implements all the standard LM32 opcodes.

9.19 M32C Dependent Features

as can assemble code for several different members of the Renesas M32C family. Normally the default is to assemble code for the M16C microprocessor. The `-m32c` option may be used to change the default to the M32C microprocessor.

9.19.1 M32C Options

The Renesas M32C version of **as** has these machine-dependent options:

- `-m32c` Assemble M32C instructions.
- `-m16c` Assemble M16C instructions (default).
- `-relax` Enable support for link-time relaxations.
- `-h-tick-hex` Support H'00 style hex constants in addition to 0x00 style.

9.19.2 Symbolic Operand Modifiers

The assembler supports several modifiers when using symbol addresses in M32C instruction operands. The general syntax is the following:

`%modifier(symbol)`

`%dsp8`

`%dsp16`

These modifiers override the assembler's assumptions about how big a symbol's address is. Normally, when it sees an operand like `'sym[a0]'` it assumes `'sym'` may require the widest displacement field (16 bits for `'-m16c'`, 24 bits for `'-m32c'`). These modifiers tell it to assume the address will fit in an 8 or 16 bit (respectively) unsigned displacement. Note that, of course, if it doesn't actually fit you will get linker errors. Example:

```
mov.w %dsp8(sym)[a0],r1
mov.b #0,%dsp8(sym)[a0]
```

`%hi8`

This modifier allows you to load bits 16 through 23 of a 24 bit address into an 8 bit register. This is useful with, for example, the M16C `'smovf'` instruction, which expects a 20 bit address in `'r1h'` and `'a0'`. Example:

```
mov.b #%hi8(sym),r1h
mov.w #%lo16(sym),a0
smovf.b
```

`%lo16`

Likewise, this modifier allows you to load bits 0 through 15 of a 24 bit address into a 16 bit register.

`%hi16`

This modifier allows you to load bits 16 through 31 of a 32 bit address into a 16 bit register. While the M32C family only has 24 bits of address space, it does support addresses in pairs of 16 bit registers (like `'a1a0'` for the `'lde'` instruction). This modifier is for loading the upper half in such cases. Example:

```
mov.w  #hi16(sym),a1
mov.w  #lo16(sym),a0
...
lde.w  [a1a0],r1
```

9.20 M32R Dependent Features

9.20.1 M32R Options

The Renesas M32R version of `as` has a few machine dependent options:

- `-m32rx` `as` can assemble code for several different members of the Renesas M32R family. Normally the default is to assemble code for the M32R microprocessor. This option may be used to change the default to the M32RX microprocessor, which adds some more instructions to the basic M32R instruction set, and some additional parameters to some of the original instructions.
- `-m32r2` This option changes the target processor to the the M32R2 microprocessor.
- `-m32r` This option can be used to restore the assembler's default behaviour of assembling for the M32R microprocessor. This can be useful if the default has been changed by a previous command line option.
- `-little` This option tells the assembler to produce little-endian code and data. The default is dependent upon how the toolchain was configured.
- `-EL` This is a synonym for *-little*.
- `-big` This option tells the assembler to produce big-endian code and data.
- `-EB` This is a synonym for *-big*.
- `-KPIC` This option specifies that the output of the assembler should be marked as position-independent code (PIC).
- `-parallel` This option tells the assembler to attempts to combine two sequential instructions into a single, parallel instruction, where it is legal to do so.
- `-no-parallel` This option disables a previously enabled *-parallel* option.
- `-no-bitinst` This option disables the support for the extended bit-field instructions provided by the M32R2. If this support needs to be re-enabled the *-bitinst* switch can be used to restore it.
- `-O` This option tells the assembler to attempt to optimize the instructions that it produces. This includes filling delay slots and converting sequential instructions into parallel ones. This option implies *-parallel*.
- `-warn-explicit-parallel-conflicts` Instructs `as` to produce warning messages when questionable parallel instructions are encountered. This option is enabled by default, but `gcc` disables it when it invokes `as` directly. Questionable instructions are those whose behaviour would be different if they were executed sequentially. For example the code fragment `'mv r1, r2 || mv r3, r1'` produces a different result from `'mv r1, r2 \n mv r3, r1'` since the former moves r1 into r3 and then r2 into r1, whereas the later moves r2 into r1 and r3.
- `-Wp` This is a shorter synonym for the *-warn-explicit-parallel-conflicts* option.

- no-warn-explicit-parallel-conflicts**
Instructs **as** not to produce warning messages when questionable parallel instructions are encountered.
- Wnp** This is a shorter synonym for the *-no-warn-explicit-parallel-conflicts* option.
- ignore-parallel-conflicts**
This option tells the assembler's to stop checking parallel instructions for constraint violations. This ability is provided for hardware vendors testing chip designs and should not be used under normal circumstances.
- no-ignore-parallel-conflicts**
This option restores the assembler's default behaviour of checking parallel instructions to detect constraint violations.
- Ip** This is a shorter synonym for the *-ignore-parallel-conflicts* option.
- nIp** This is a shorter synonym for the *-no-ignore-parallel-conflicts* option.
- warn-unmatched-high**
This option tells the assembler to produce a warning message if a **.high** pseudo op is encountered without a matching **.low** pseudo op. The presence of such an unmatched pseudo op usually indicates a programming error.
- no-warn-unmatched-high**
Disables a previously enabled *-warn-unmatched-high* option.
- Wuh** This is a shorter synonym for the *-warn-unmatched-high* option.
- Wnuh** This is a shorter synonym for the *-no-warn-unmatched-high* option.

9.20.2 M32R Directives

The Renease M32R version of **as** has a few architecture specific directives:

low expression

The **low** directive computes the value of its expression and places the lower 16-bits of the result into the immediate-field of the instruction. For example:

```
or3   r0, r0, #low(0x12345678) ; compute r0 = r0 | 0x5678
add3, r0, r0, #low(fred)      ; compute r0 = r0 + low 16-bits of address of fred
```

high expression

The **high** directive computes the value of its expression and places the upper 16-bits of the result into the immediate-field of the instruction. For example:

```
seth  r0, #high(0x12345678) ; compute r0 = 0x12340000
seth, r0, #high(fred)       ; compute r0 = upper 16-bits of address of fred
```

shigh expression

The **shigh** directive is very similar to the **high** directive. It also computes the value of its expression and places the upper 16-bits of the result into the immediate-field of the instruction. The difference is that **shigh** also checks to see if the lower 16-bits could be interpreted as a signed number, and if so it assumes that a borrow will occur from the upper-16 bits. To compensate for this the **shigh** directive pre-biases the upper 16 bit value by adding one to it. For example:

For example:

```
seth  r0, #shigh(0x12345678) ; compute r0 = 0x12340000
seth  r0, #shigh(0x00008000) ; compute r0 = 0x00010000
```

In the second example the lower 16-bits are 0x8000. If these are treated as a signed value and sign extended to 32-bits then the value becomes 0xffff8000. If this value is then added to 0x00010000 then the result is 0x00008000.

This behaviour is to allow for the different semantics of the `or3` and `add3` instructions. The `or3` instruction treats its 16-bit immediate argument as unsigned whereas the `add3` treats its 16-bit immediate as a signed value. So for example:

```
seth  r0, #shigh(0x00008000)
add3  r0, r0, #low(0x00008000)
```

Produces the correct result in r0, whereas:

```
seth  r0, #shigh(0x00008000)
or3   r0, r0, #low(0x00008000)
```

Stores 0xffff8000 into r0.

Note - the `shigh` directive does not know where in the assembly source code the lower 16-bits of the value are going set, so it cannot check to make sure that an `or3` instruction is being used rather than an `add3` instruction. It is up to the programmer to make sure that correct directives are used.

- `.m32r` The directive performs a similar thing as the `-m32r` command line option. It tells the assembler to only accept M32R instructions from now on. An instructions from later M32R architectures are refused.
- `.m32rx` The directive performs a similar thing as the `-m32rx` command line option. It tells the assembler to start accepting the extra instructions in the M32RX ISA as well as the ordinary M32R ISA.
- `.m32r2` The directive performs a similar thing as the `-m32r2` command line option. It tells the assembler to start accepting the extra instructions in the M32R2 ISA as well as the ordinary M32R ISA.
- `.little` The directive performs a similar thing as the `-little` command line option. It tells the assembler to start producing little-endian code and data. This option should be used with care as producing mixed-endian binary files is fraught with danger.
- `.big` The directive performs a similar thing as the `-big` command line option. It tells the assembler to start producing big-endian code and data. This option should be used with care as producing mixed-endian binary files is fraught with danger.

9.20.3 M32R Warnings

There are several warning and error messages that can be produced by `as` which are specific to the M32R:

output of 1st instruction is the same as an input to 2nd instruction - is this intentional ?

This message is only produced if warnings for explicit parallel conflicts have been enabled. It indicates that the assembler has encountered a parallel instruction in which the destination register of the left hand instruction is used

as an input register in the right hand instruction. For example in this code fragment `'mv r1, r2 || neg r3, r1'` register r1 is the destination of the move instruction and the input to the neg instruction.

output of 2nd instruction is the same as an input to 1st instruction - is this intentional ?

This message is only produced if warnings for explicit parallel conflicts have been enabled. It indicates that the assembler has encountered a parallel instruction in which the destination register of the right hand instruction is used as an input register in the left hand instruction. For example in this code fragment `'mv r1, r2 || neg r2, r3'` register r2 is the destination of the neg instruction and the input to the move instruction.

instruction `'...'` is for the M32RX only

This message is produced when the assembler encounters an instruction which is only supported by the M32Rx processor, and the `'-m32rx'` command line flag has not been specified to allow assembly of such instructions.

unknown instruction `'...'`

This message is produced when the assembler encounters an instruction which it does not recognize.

only the NOP instruction can be issued in parallel on the m32r

This message is produced when the assembler encounters a parallel instruction which does not involve a NOP instruction and the `'-m32rx'` command line flag has not been specified. Only the M32Rx processor is able to execute two instructions in parallel.

instruction `'...'` cannot be executed in parallel.

This message is produced when the assembler encounters a parallel instruction which is made up of one or two instructions which cannot be executed in parallel.

Instructions share the same execution pipeline

This message is produced when the assembler encounters a parallel instruction whose components both use the same execution pipeline.

Instructions write to the same destination register.

This message is produced when the assembler encounters a parallel instruction where both components attempt to modify the same register. For example these code fragments will produce this message: `'mv r1, r2 || neg r1, r3'` `'jl r0 || mv r14, r1'` `'st r2, @-r1 || mv r1, r3'` `'mv r1, r2 || ld r0, @r1+' 'cmp r1, r2 || addx r3, r4' (Both write to the condition bit)`

9.21 M680x0 Dependent Features

9.21.1 M680x0 Options

The Motorola 680x0 version of **as** has a few machine dependent options:

‘-march=architecture’

This option specifies a target architecture. The following architectures are recognized: 68000, 68010, 68020, 68030, 68040, 68060, **cpu32**, **isaa**, **isaaplus**, **isab**, **isac** and **cfv4e**.

‘-mcpu=cpu’

This option specifies a target cpu. When used in conjunction with the **‘-march’** option, the cpu must be within the specified architecture. Also, the generic features of the architecture are used for instruction generation, rather than those of the specific chip.

‘-m[no-]68851’

‘-m[no-]68881’

‘-m[no-]div’

‘-m[no-]usp’

‘-m[no-]float’

‘-m[no-]mac’

‘-m[no-]emac’

Enable or disable various architecture specific features. If a chip or architecture by default supports an option (for instance **‘-march=isaaplus’** includes the **‘-mdiv’** option), explicitly disabling the option will override the default.

‘-l’

You can use the **‘-l’** option to shorten the size of references to undefined symbols. If you do not use the **‘-l’** 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.

‘--register-prefix-optional’

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.

‘--bitwise-or’

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.

`--base-size-default-16 --base-size-default-32`

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.

`--disp-size-default-16 --disp-size-default-32`

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.

`--pcrel` Always keep branches PC-relative. In the M680x0 architecture all branches are defined as PC-relative. However, on some processors they are limited to word displacements maximum. When `as` needs a long branch that is not available, it normally emits an absolute jump instead. This option disables this substitution. When this option is given and no long branches are available, only word branches will be emitted. An error message will be generated if a word branch cannot reach its target. This option has no effect on 68020 and other processors that have long branches. see [Section 9.21.6.1 \[Branch Improvement\]](#), page 177.

`-m68000` `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'`
`'-m5202'`
`'-m5204'`
`'-m5206'`
`'-m5206e'`
`'-m521x'`
`'-m5249'`
`'-m528x'`
`'-m5307'`
`'-m5407'`
`'-m547x'`
`'-m548x'`
`'-mcfv4'`
`'-mcfv4e'` 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.

9.21.2 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, ‘movl’ is equivalent to ‘mov.l’.

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)@(number,register:size:scale)’

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

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

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

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

9.21.3 Motorola Syntax

The standard Motorola syntax for this chip differs from the syntax already discussed (see [Section 9.21.2 \[Syntax\], page 174](#)). *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*, *number*)’

The *number*, 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.

9.21.4 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.

9.21.5 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.

.arch *name*
 Select the target architecture and extension features. Valid values for *name* are the same as for the ‘**-march**’ command line option. This directive cannot be specified after any instructions have been assembled. If it is given multiple times, or in conjunction with the ‘**-march**’ option, all uses must be for the same architecture and extension set.

.cpu *name*
 Select the target cpu. Valid valuse for *name* are the same as for the ‘**-mcpu**’ command line option. This directive cannot be specified after any instructions have been assembled. If it is given multiple times, or in conjunction with the ‘**-mopt**’ option, all uses must be for the same cpu.

9.21.6 Opcodes

9.21.6.1 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 | 68020 LONG | 68000/10, not PC-relative ABSOLUTE LONG JUMP | OK ** |
| jbsr | bsrs | bsrw | bsrl | jsr | |
| jra | bras | braw | bral | jmp | |
| * jXX | bXXs | bXXw | bXXl | bNXs;jmp | |
| * dbXX | N/A | dbXXw | dbXX;bras;bral | dbXX;bras;jmp | |
| fjXX | N/A | fbXXw | fbXXl | N/A | |

XX: condition

NX: negative of condition XX

*—see full description below

**—this expansion mode is disallowed by ‘--pcrel’

jbsr

jra

These are the simplest jump pseudo-operations; they always map to one particular machine instruction, depending on the displacement to the branch target. This instruction will be a byte or word branch if that is sufficient. Otherwise, a long branch will be emitted if available. If no long branches are available and the ‘--pcrel’ option is not given, an absolute long jump will be emitted instead. If no long branches are available, the ‘--pcrel’ option is given, and a word branch cannot reach the target, an error message is generated.

In addition to standard branch operands, **as** allows these pseudo-operations to have all operands that are allowed for jsr and jmp, substituting these instructions if the operand given is not valid for a branch instruction.

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 |

Usually, each of these pseudo-operations expands to a single branch instruction. However, if a word branch is not sufficient, no long branches are available, and the ‘--pcrel’ option is not given, **as** issues a longer code fragment in terms of NX, the opposite condition to XX. For example, under these conditions:

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

```

Motorola ‘dbXX’ instructions allow word displacements only. When a word displacement is sufficient, each of these pseudo-operations expands to the corresponding Motorola instruction. When a word displacement is not sufficient and long branches are available, when the source reads ‘dbXX foo’, **as** emits

```

    dbXX oo1
    bras oo2
oo1:bral foo
oo2:

```

If, however, long branches are not available and the ‘--pcrel’ option is not given, **as** emits

```

    dbXX oo1
    bras oo2
oo1:jmp 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  fjsef  fjsef  fjst   fjueq  fjuge
    fjugt  fjule  fjult  fjun

```

Each of these pseudo-operations always expands to a single Motorola coprocessor branch instruction, word or long. All Motorola coprocessor branch instructions allow both word and long displacements.

9.21.6.2 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.

9.22 M68HC11 and M68HC12 Dependent Features

9.22.1 M68HC11 and M68HC12 Options

The Motorola 68HC11 and 68HC12 version of `as` have a few machine dependent options.

-m68hc11 This option switches the assembler in the M68HC11 mode. In this mode, the assembler only accepts 68HC11 operands and mnemonics. It produces code for the 68HC11.

-m68hc12 This option switches the assembler in the M68HC12 mode. In this mode, the assembler also accepts 68HC12 operands and mnemonics. It produces code for the 68HC12. A few 68HC11 instructions are replaced by some 68HC12 instructions as recommended by Motorola specifications.

-m68hcs12 This option switches the assembler in the M68HCS12 mode. This mode is similar to `-m68hc12` but specifies to assemble for the 68HCS12 series. The only difference is on the assembling of the `'movb'` and `'movw'` instruction when a PC-relative operand is used.

-mshort This option controls the ABI and indicates to use a 16-bit integer ABI. It has no effect on the assembled instructions. This is the default.

-mlong This option controls the ABI and indicates to use a 32-bit integer ABI.

-mshort-double This option controls the ABI and indicates to use a 32-bit float ABI. This is the default.

-mlong-double This option controls the ABI and indicates to use a 64-bit float ABI.

--strict-direct-mode You can use the `'--strict-direct-mode'` option to disable the automatic translation of direct page mode addressing into extended mode when the instruction does not support direct mode. For example, the `'clr'` instruction does not support direct page mode addressing. When it is used with the direct page mode, `as` will ignore it and generate an absolute addressing. This option prevents `as` from doing this, and the wrong usage of the direct page mode will raise an error.

--short-branches The `'--short-branches'` option turns off the translation of relative branches into absolute branches when the branch offset is out of range. By default `as` transforms the relative branch (`'bsr'`, `'bgt'`, `'bge'`, `'beq'`, `'bne'`, `'ble'`, `'blt'`, `'bhi'`, `'bcc'`, `'bls'`, `'bcs'`, `'bmi'`, `'bvs'`, `'bvs'`, `'bra'`) into an absolute branch when the offset is out of the -128 .. 127 range. In that case, the `'bsr'` instruction is translated into a `'jsr'`, the `'bra'` instruction is translated into a `'jmp'` and the conditional branches instructions are inverted and followed by a `'jmp'`. This option disables these translations and `as` will generate an error if a relative branch is out of range. This option does not affect the optimization associated to the `'jbra'`, `'jbsr'` and `'jbXX'` pseudo opcodes.

--force-long-branches

The ‘--force-long-branches’ option forces the translation of relative branches into absolute branches. This option does not affect the optimization associated to the ‘jbra’, ‘jbsr’ and ‘jbXX’ pseudo opcodes.

--print-insn-syntax

You can use the ‘--print-insn-syntax’ option to obtain the syntax description of the instruction when an error is detected.

--print-opcodes

The ‘--print-opcodes’ option prints the list of all the instructions with their syntax. The first item of each line represents the instruction name and the rest of the line indicates the possible operands for that instruction. The list is printed in alphabetical order. Once the list is printed **as** exits.

--generate-example

The ‘--generate-example’ option is similar to ‘--print-opcodes’ but it generates an example for each instruction instead.

9.22.2 Syntax

In the M68HC11 syntax, the instruction name comes first and it may be followed by one or several operands (up to three). Operands are separated by comma (‘,’). In the normal mode, **as** will complain if too many operands are specified for a given instruction. In the MRI mode (turned on with ‘-M’ option), it will treat them as comments. Example:

```
inx
lda #23
bset 2,x #4
brclr *bot #8 foo
```

The following addressing modes are understood for 68HC11 and 68HC12:

Immediate

‘*number*’

Address Register

‘*number,X*’, ‘*number,Y*’

The *number* may be omitted in which case 0 is assumed.

Direct Addressing mode

‘**symbol*’, or ‘**digits*’

Absolute ‘*symbol*’, or ‘*digits*’

The M68HC12 has other more complex addressing modes. All of them are supported and they are represented below:

Constant Offset Indexed Addressing Mode

‘*number,reg*’

The *number* may be omitted in which case 0 is assumed. The register can be either ‘X’, ‘Y’, ‘SP’ or ‘PC’. The assembler will use the smaller post-byte definition according to the constant value (5-bit constant offset, 9-bit constant offset or 16-bit constant offset). If the constant is not known by the assembler it will use the 16-bit constant offset post-byte and the value will be resolved at link time.

Offset Indexed Indirect

`'[number,reg]'`

The register can be either 'X', 'Y', 'SP' or 'PC'.

Auto Pre-Increment/Pre-Decrement/Post-Increment/Post-Decrement

`'number,-reg'` `'number,+reg'` `'number,reg--'` `'number,reg++'`

The number must be in the range '-8'..' +8' and must not be 0. The register can be either 'X', 'Y', 'SP' or 'PC'.

Accumulator Offset

`'acc,reg'`

The accumulator register can be either 'A', 'B' or 'D'. The register can be either 'X', 'Y', 'SP' or 'PC'.

Accumulator D offset indexed-indirect

`'[D,reg]'`

The register can be either 'X', 'Y', 'SP' or 'PC'.

For example:

```
ldab 1024,sp
ldd [10,x]
orab 3,+x
stab -2,y-
ldx a,pc
sty [d,sp]
```

9.22.3 Symbolic Operand Modifiers

The assembler supports several modifiers when using symbol addresses in 68HC11 and 68HC12 instruction operands. The general syntax is the following:

`%modifier(symbol)`

- | | |
|--------------|--|
| %addr | This modifier indicates to the assembler and linker to use the 16-bit physical address corresponding to the symbol. This is intended to be used on memory window systems to map a symbol in the memory bank window. If the symbol is in a memory expansion part, the physical address corresponds to the symbol address within the memory bank window. If the symbol is not in a memory expansion part, this is the symbol address (using or not using the %addr modifier has no effect in that case). |
| %page | This modifier indicates to use the memory page number corresponding to the symbol. If the symbol is in a memory expansion part, its page number is computed by the linker as a number used to map the page containing the symbol in the memory bank window. If the symbol is not in a memory expansion part, the page number is 0. |
| %hi | This modifier indicates to use the 8-bit high part of the physical address of the symbol. |
| %lo | This modifier indicates to use the 8-bit low part of the physical address of the symbol. |

For example a 68HC12 call to a function 'foo_example' stored in memory expansion part could be written as follows:

```
call %addr(foo_example),%page(foo_example)
```

and this is equivalent to

```
call foo_example
```

And for 68HC11 it could be written as follows:

```
ldab #%page(foo_example)
stab _page_switch
jsr  %addr(foo_example)
```

9.22.4 Assembler Directives

The 68HC11 and 68HC12 version of **as** have the following specific assembler directives:

.relax The **relax** directive is used by the ‘GNU Compiler’ to emit a specific relocation to mark a group of instructions for linker relaxation. The sequence of instructions within the group must be known to the linker so that relaxation can be performed.

.mode [mshort|mlong|mshort-double|mlong-double]
This directive specifies the ABI. It overrides the ‘-mshort’, ‘-mlong’, ‘-mshort-double’ and ‘-mlong-double’ options.

.far *symbol*
This directive marks the symbol as a ‘far’ symbol meaning that it uses a ‘call/rtc’ calling convention as opposed to ‘jsr/rts’. During a final link, the linker will identify references to the ‘far’ symbol and will verify the proper calling convention.

.interrupt *symbol*
This directive marks the symbol as an interrupt entry point. This information is then used by the debugger to correctly unwind the frame across interrupts.

.xrefb *symbol*
This directive is defined for compatibility with the ‘Specification for Motorola 8 and 16-Bit Assembly Language Input Standard’ and is ignored.

9.22.5 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.

9.22.6 Opcodes

9.22.6.1 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 prepending ‘j’ to the start of Motorola mnemonic. These pseudo opcodes are not affected by the ‘--short-branches’ or ‘--force-long-branches’ options.

The following table summarizes the pseudo-operations.

| Displacement Width | | | | |
|--------------------|--------------|-----------------------|--------------|-------------------|
| Options | | | | |
| --short-branches | | --force-long-branches | | |
| Op | BYTE | WORD | BYTE | WORD |
| bsr | bsr <pc-rel> | <error> | | jsr <abs> |
| bra | bra <pc-rel> | <error> | | jmp <abs> |
| jbsr | bsr <pc-rel> | jsr <abs> | bsr <pc-rel> | jsr <abs> |
| jbra | bra <pc-rel> | jmp <abs> | bra <pc-rel> | jmp <abs> |
| bXX | bXX <pc-rel> | <error> | | bNX +3; jmp <abs> |
| jbXX | bXX <pc-rel> | bNX +3; jmp <abs> | bXX <pc-rel> | bNX +3; jmp <abs> |

XX: condition

NX: negative of condition XX

jbsr

jbra

These are the simplest jump pseudo-operations; they always map to one particular machine instruction, depending on the displacement to the branch target.

jbXX

Here, ‘jbXX’ 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:

```

        jbcc  jbeq  jbge  jbgt  jbhi  jbvs  jbpl  jblo
        jbcs  jbne  jblt  jble  jbls  jbvc  jbmi

```

For the cases of non-PC relative displacements and long displacements, as issues a longer code fragment in terms of NX, the opposite condition to XX. For example, for the non-PC relative case:

```
jbXX foo
```

gives

```

        bNXs oof
        jmp foo
oof:

```

9.23 MicroBlaze Dependent Features

The Xilinx MicroBlaze processor family includes several variants, all using the same core instruction set. This chapter covers features of the GNU assembler that are specific to the MicroBlaze architecture. For details about the MicroBlaze instruction set, please see the *MicroBlaze Processor Reference Guide (UG081)* available at www.xilinx.com.

9.23.1 Directives

A number of assembler directives are available for MicroBlaze.

`.data8 expression,...`

This directive is an alias for `.byte`. Each expression is assembled into an eight-bit value.

`.data16 expression,...`

This directive is an alias for `.hword`. Each expression is assembled into an 16-bit value.

`.data32 expression,...`

This directive is an alias for `.word`. Each expression is assembled into an 32-bit value.

`.ent name[,label]`

This directive is an alias for `.func` denoting the start of function *name* at (optional) *label*.

`.end name[,label]`

This directive is an alias for `.endfunc` denoting the end of function *name*.

`.gpword label,...`

This directive is an alias for `.rva`. The resolved address of *label* is stored in the data section.

`.weakext label`

Declare that *label* is a weak external symbol.

`.rodata` Switch to `.rodata` section. Equivalent to `.section .rodata`

`.sdata2` Switch to `.sdata2` section. Equivalent to `.section .sdata2`

`.sdata` Switch to `.sdata` section. Equivalent to `.section .sdata`

`.bss` Switch to `.bss` section. Equivalent to `.section .bss`

`.sbss` Switch to `.sbss` section. Equivalent to `.section .sbss`

9.24 MIPS Dependent Features

GNU `as` for MIPS architectures supports several different MIPS processors, and MIPS ISA levels I through V, MIPS32, and MIPS64. 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.

9.24.1 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.

-KPIC Generate SVR4-style PIC. This option tells the assembler to generate SVR4-style position-independent macro expansions. It also tells the assembler to mark the output file as PIC.

-mvxworks-pic

Generate VxWorks PIC. This option tells the assembler to generate VxWorks-style position-independent macro expansions.

-mips1

-mips2

-mips3

-mips4

-mips5xo

-mips32

-mips32r2

-mips64

-mips64r2

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. ‘-mips5’, ‘-mips32’, ‘-mips32r2’, ‘-mips64’, and ‘-mips64r2’ correspond to generic MIPS V, MIPS32, MIPS32 RELEASE 2, MIPS64, and MIPS64 RELEASE 2 ISA processors, respectively. You can also switch instruction sets during the assembly; see [Section 9.24.5 \[MIPS ISA\]](#), [page 191](#).

-mfp32

-mfp32 Some macros have different expansions for 32-bit and 64-bit registers. The register sizes are normally inferred from the ISA and ABI, but these flags force a certain group of registers to be treated as 32 bits wide at all times. ‘-mfp32’

controls the size of general-purpose registers and ‘-mfp32’ controls the size of floating-point registers.

The `.set gp=32` and `.set fp=32` directives allow the size of registers to be changed for parts of an object. The default value is restored by `.set gp=default` and `.set fp=default`.

On some MIPS variants there is a 32-bit mode flag; when this flag is set, 64-bit instructions generate a trap. Also, some 32-bit OSes only save the 32-bit registers on a context switch, so it is essential never to use the 64-bit registers.

`-mfp64`

`-mfp64` Assume that 64-bit registers are available. This is provided in the interests of symmetry with ‘-mfp32’ and ‘-mfp32’.

The `.set gp=64` and `.set fp=64` directives allow the size of registers to be changed for parts of an object. The default value is restored by `.set gp=default` and `.set fp=default`.

`-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.

`-msmartmips`

`-mno-smartmips`

Enables the SmartMIPS extensions to the MIPS32 instruction set, which provides a number of new instructions which target smartcard and cryptographic applications. This is equivalent to putting `.set smartmips` at the start of the assembly file. ‘-mno-smartmips’ turns off this option.

`-mips3d`

`-no-mips3d`

Generate code for the MIPS-3D Application Specific Extension. This tells the assembler to accept MIPS-3D instructions. ‘-no-mips3d’ turns off this option.

`-mdmx`

`-no-mdmx` Generate code for the MDMX Application Specific Extension. This tells the assembler to accept MDMX instructions. ‘-no-mdmx’ turns off this option.

`-mdsp`

`-mno-dsp` Generate code for the DSP Release 1 Application Specific Extension. This tells the assembler to accept DSP Release 1 instructions. ‘-mno-dsp’ turns off this option.

`-mdspr2`

`-mno-dspr2`

Generate code for the DSP Release 2 Application Specific Extension. This option implies -mdsp. This tells the assembler to accept DSP Release 2 instructions. ‘-mno-dspr2’ turns off this option.

`-mmt`

`-mno-mt` Generate code for the MT Application Specific Extension. This tells the assembler to accept MT instructions. ‘-mno-mt’ turns off this option.

`-mfix7000`

`-mno-fix7000`

Cause nops to be inserted if the read of the destination register of an `mfhi` or `mflo` instruction occurs in the following two instructions.

`-mfix-loongson2f-jump`

`-mno-fix-loongson2f-jump`

Eliminate instruction fetch from outside 256M region to work around the Loongson2F ‘`jump`’ instructions. Without it, under extreme cases, the kernel may crash. The issue has been solved in latest processor batches, but this fix has no side effect to them.

`-mfix-loongson2f-nop`

`-mno-fix-loongson2f-nop`

Replace nops by `or at,at,zero` to work around the Loongson2F ‘`nop`’ errata. Without it, under extreme cases, cpu might deadlock. The issue has been solved in latest loongson2f batches, but this fix has no side effect to them.

`-mfix-vr4120`

`-mno-fix-vr4120`

Insert nops to work around certain VR4120 errata. This option is intended to be used on GCC-generated code: it is not designed to catch all problems in hand-written assembler code.

`-mfix-vr4130`

`-mno-fix-vr4130`

Insert nops to work around the VR4130 ‘`mflo`’/‘`mfhi`’ errata.

`-mfix-24k`

`-no-mfix-24k`

Insert nops to work around the 24K ‘`eret`’/‘`deret`’ errata.

`-mfix-cn63xxp1`

`-mno-fix-cn63xxp1`

Replace `pref` hints 0 - 4 and 6 - 24 with hint 28 to work around certain CN63XXP1 errata.

`-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.

`-march=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, vr4120, vr4130, vr4181,
4300, 4400, 4600, 4650, 5000, rm5200, rm5230, rm5231, rm5261,
rm5721, vr5400, vr5500, 6000, rm7000, 8000, rm9000, 10000,
12000, 14000, 16000, 4kc, 4km, 4kp, 4ksc, 4kec, 4kem, 4kep,
4ksd, m4k, m4kp, 24kc, 24kf2_1, 24kf, 24kf1_1, 24kec, 24kef2_1,
24kef, 24kef1_1, 34kc, 34kf2_1, 34kf, 34kf1_1, 74kc, 74kf2_1, 74kf,
74kf1_1, 74kf3_2, 1004kc, 1004kf2_1, 1004kf, 1004kf1_1, 5kc, 5kf,
20kc, 25kf, sb1, sb1a, loongson2e, loongson2f, octeon, xlr

For compatibility reasons, ‘`nx`’ and ‘`bfx`’ are accepted as synonyms for ‘`nf1_1`’. These values are deprecated.

`-mtune=cpu`

Schedule and tune for a particular MIPS *cpu*. Valid *cpu* values are identical to ‘`-march=cpu`’.

`-mabi=abi`

Record which ABI the source code uses. The recognized arguments are: ‘`32`’, ‘`32`’, ‘`o64`’, ‘`64`’ and ‘`eabi`’.

`-msym32`

`-mno-sym32`

Equivalent to adding `.set sym32` or `.set nosym32` to the beginning of the assembler input. See [Section 9.24.4 \[MIPS symbol sizes\]](#), page 190.

`-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.

`-msoft-float`

`-mhard-float`

Disable or enable floating-point instructions. Note that by default floating-point instructions are always allowed even with CPU targets that don’t have support for these instructions.

`-msingle-float`

`-mdouble-float`

Disable or enable double-precision floating-point operations. Note that by default double-precision floating-point operations are always allowed even with CPU targets that don't have support for these operations.

`--construct-floats`

`--no-construct-floats`

The `--no-construct-floats` option disables the construction of double width floating point constants by loading the two halves of the value into the two single width floating point registers that make up the double width register. This feature is useful if the processor support the FR bit in its status register, and this bit is known (by the programmer) to be set. This bit prevents the aliasing of the double width register by the single width registers.

By default `--construct-floats` is selected, allowing construction of these floating point constants.

`--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.

`-mpdr`

`-mno-pdr` Control generation of `.pdr` sections. Off by default on IRIX, on elsewhere.

`-mshared`

`-mno-shared`

When generating code using the Unix calling conventions (selected by '`-KPIC`' or '`-mcall_shared`'), `gas` will normally generate code which can go into a shared library. The '`-mno-shared`' option tells `gas` to generate code which uses the calling convention, but can not go into a shared library. The resulting code is slightly more efficient. This option only affects the handling of the '`.cpload`' and '`.cpsetup`' pseudo-ops.

9.24.2 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.

9.24.3 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`, `.endef`, `.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!

9.24.4 Directives to override the size of symbols

The n64 ABI allows symbols to have any 64-bit value. Although this provides a great deal of flexibility, it means that some macros have much longer expansions than their 32-bit counterparts. For example, the non-PIC expansion of ‘`dla $4,sym`’ is usually:

```
lui    $4,%highest(sym)
lui    $1,%hi(sym)
daddiu $4,$4,%higher(sym)
daddiu $1,$1,%lo(sym)
dsll32 $4,$4,0
daddu  $4,$4,$1
```

whereas the 32-bit expansion is simply:

```
lui    $4,%hi(sym)
daddiu $4,$4,%lo(sym)
```

n64 code is sometimes constructed in such a way that all symbolic constants are known to have 32-bit values, and in such cases, it’s preferable to use the 32-bit expansion instead of the 64-bit expansion.

You can use the `.set sym32` directive to tell the assembler that, from this point on, all expressions of the form ‘`symbol`’ or ‘`symbol + offset`’ have 32-bit values. For example:

```
.set sym32
dla    $4,sym
lw     $4,sym+16
sw     $4,sym+0x8000($4)
```

will cause the assembler to treat ‘`sym`’, `sym+16` and `sym+0x8000` as 32-bit values. The handling of non-symbolic addresses is not affected.

The directive `.set nosym32` ends a `.set sym32` block and reverts to the normal behavior. It is also possible to change the symbol size using the command-line options ‘`-msym32`’ and ‘`-mno-sym32`’.

These options and directives are always accepted, but at present, they have no effect for anything other than n64.

9.24.5 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 5, or 32, 32r2, 64 or 64r2. The values other than 0 make 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 MIPS3 instructions while assembling in 32 bit mode. Use this directive with care!

The `.set arch=cpu` directive provides even finer control. It changes the effective CPU target and allows the assembler to use instructions specific to a particular CPU. All CPUs supported by the `‘-march’` command line option are also selectable by this directive. The original value is restored by `.set arch=default`.

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.

9.24.6 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.

9.24.7 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.

The `.global` and `.globl` directives supported by `as` will by default mark the symbol as pointing to a region of data not code. This means that, for example, any instructions following such a symbol will not be disassembled by `objdump` as it will regard them as data. To change this behaviour an optional section name can be placed after the symbol name in the `.global` directive. If this section exists and is known to be a code section, then the symbol will be marked as pointing at code not data. The syntax for the directive is:

```
.global symbol[ section][, symbol[ section]] ...
```

Here is a short example:

```
.global foo .text, bar, baz .data
foo:
```

```

        nop
bar:
        .word 0x0
baz:
        .word 0x1

```

9.24.8 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.

9.24.9 Directives to control generation of MIPS ASE instructions

The directive `.set mips3d` makes the assembler accept instructions from the MIPS-3D Application Specific Extension from that point on in the assembly. The `.set nomips3d` directive prevents MIPS-3D instructions from being accepted.

The directive `.set smartmips` makes the assembler accept instructions from the SmartMIPS Application Specific Extension to the MIPS32 ISA from that point on in the assembly. The `.set nosmartmips` directive prevents SmartMIPS instructions from being accepted.

The directive `.set mdmx` makes the assembler accept instructions from the MDMX Application Specific Extension from that point on in the assembly. The `.set nomdmx` directive prevents MDMX instructions from being accepted.

The directive `.set dsp` makes the assembler accept instructions from the DSP Release 1 Application Specific Extension from that point on in the assembly. The `.set nodsp` directive prevents DSP Release 1 instructions from being accepted.

The directive `.set dspr2` makes the assembler accept instructions from the DSP Release 2 Application Specific Extension from that point on in the assembly. This directive implies `.set dsp`. The `.set nodspr2` directive prevents DSP Release 2 instructions from being accepted.

The directive `.set mt` makes the assembler accept instructions from the MT Application Specific Extension from that point on in the assembly. The `.set nomt` directive prevents MT instructions from being accepted.

Traditional MIPS assemblers do not support these directives.

9.24.10 Directives to override floating-point options

The directives `.set softfloat` and `.set hardfloat` provide finer control of disabling and enabling float-point instructions. These directives always override the default (that hard-float instructions are accepted) or the command-line options (`'-msoft-float'` and `'-mhard-float'`).

The directives `.set singlefloat` and `.set doublefloat` provide finer control of disabling and enabling double-precision float-point operations. These directives always override the default (that double-precision operations are accepted) or the command-line options (`'-msingle-float'` and `'mdouble-float'`).

Traditional MIPS assemblers do not support these directives.

9.25 MMIX Dependent Features

9.25.1 Command-line Options

The MMIX version of `as` has some machine-dependent options.

When ‘`--fixed-special-register-names`’ is specified, only the register names specified in [Section 9.25.3.3 \[MMIX-Regs\], page 196](#) are recognized in the instructions `PUT` and `GET`.

You can use the ‘`--globalize-symbols`’ to make all symbols global. This option is useful when splitting up a `mmixal` program into several files.

The ‘`--gnu-syntax`’ turns off most syntax compatibility with `mmixal`. Its usability is currently doubtful.

The ‘`--relax`’ option is not fully supported, but will eventually make the object file prepared for linker relaxation.

If you want to avoid inadvertently calling a predefined symbol and would rather get an error, for example when using `as` with a compiler or other machine-generated code, specify ‘`--no-predefined-syms`’. This turns off built-in predefined definitions of all such symbols, including rounding-mode symbols, segment symbols, ‘`BIT`’ symbols, and `TRAP` symbols used in `mmix` “system calls”. It also turns off predefined special-register names, except when used in `PUT` and `GET` instructions.

By default, some instructions are expanded to fit the size of the operand or an external symbol (see [Section 9.25.2 \[MMIX-Expand\], page 195](#)). By passing ‘`--no-expand`’, no such expansion will be done, instead causing errors at link time if the operand does not fit.

The `mmixal` documentation (see [\[mmixsite\], page 195](#)) specifies that global registers allocated with the ‘`GREG`’ directive (see [\[MMIX-greg\], page 197](#)) and initialized to the same non-zero value, will refer to the same global register. This isn’t strictly enforceable in `as` since the final addresses aren’t known until link-time, but it will do an effort unless the ‘`--no-merge-gregs`’ option is specified. (Register merging isn’t yet implemented in `ld`.)

`as` will warn every time it expands an instruction to fit an operand unless the option ‘`-x`’ is specified. It is believed that this behaviour is more useful than just mimicking `mmixal`’s behaviour, in which instructions are only expanded if the ‘`-x`’ option is specified, and assembly fails otherwise, when an instruction needs to be expanded. It needs to be kept in mind that `mmixal` is both an assembler and linker, while `as` will expand instructions that at link stage can be contracted. (Though linker relaxation isn’t yet implemented in `ld`.) The option ‘`-x`’ also implies ‘`--linker-allocated-gregs`’.

If instruction expansion is enabled, `as` can expand a ‘`PUSHJ`’ instruction into a series of instructions. The shortest expansion is to not expand it, but just mark the call as redirectable to a stub, which `ld` creates at link-time, but only if the original ‘`PUSHJ`’ instruction is found not to reach the target. The stub consists of the necessary instructions to form a jump to the target. This happens if `as` can assert that the ‘`PUSHJ`’ instruction can reach such a stub. The option ‘`--no-pushj-stubs`’ disables this shorter expansion, and the longer series of instructions is then created at assembly-time. The option ‘`--no-stubs`’ is a synonym, intended for compatibility with future releases, where generation of stubs for other instructions may be implemented.

Usually a two-operand-expression (see [\[GREG-base\], page 198](#)) without a matching ‘`GREG`’ directive is treated as an error by `as`. When the option ‘`--linker-allocated-gregs`’

is in effect, they are instead passed through to the linker, which will allocate as many global registers as is needed.

9.25.2 Instruction expansion

When `as` encounters an instruction with an operand that is either not known or does not fit the operand size of the instruction, `as` (and `ld`) will expand the instruction into a sequence of instructions semantically equivalent to the operand fitting the instruction. Expansion will take place for the following instructions:

‘GETA’ Expands to a sequence of four instructions: `SETL`, `INCML`, `INCMH` and `INCH`. The operand must be a multiple of four.

Conditional branches

A branch instruction is turned into a branch with the complemented condition and prediction bit over five instructions; four instructions setting `$255` to the operand value, which like with `GETA` must be a multiple of four, and a final `GO $255,$255,0`.

‘PUSHJ’ Similar to expansion for conditional branches; four instructions set `$255` to the operand value, followed by a `PUSHGO $255,$255,0`.

‘JMP’ Similar to conditional branches and `PUSHJ`. The final instruction is `GO $255,$255,0`.

The linker `ld` is expected to shrink these expansions for code assembled with `‘--relax’` (though not currently implemented).

9.25.3 Syntax

The assembly syntax is supposed to be upward compatible with that described in Sections 1.3 and 1.4 of ‘*The Art of Computer Programming, Volume 1*’. Draft versions of those chapters as well as other MMIX information is located at <http://www-cs-faculty.stanford.edu/~knuth/mmix-news.html>. Most code examples from the `mmixal` package located there should work unmodified when assembled and linked as single files, with a few noteworthy exceptions (see Section 9.25.4 [MMIX-`mmixal`], page 199).

Before an instruction is emitted, the current location is aligned to the next four-byte boundary. If a label is defined at the beginning of the line, its value will be the aligned value.

In addition to the traditional hex-prefix `‘0x’`, a hexadecimal number can also be specified by the prefix character `‘#’`.

After all operands to an MMIX instruction or directive have been specified, the rest of the line is ignored, treated as a comment.

9.25.3.1 Special Characters

The characters `‘*’` and `‘#’` are line comment characters; each start a comment at the beginning of a line, but only at the beginning of a line. A `‘#’` prefixes a hexadecimal number if found elsewhere on a line.

Two other characters, ‘%’ and ‘!’, each start a comment anywhere on the line. Thus you can’t use the ‘modulus’ and ‘not’ operators in expressions normally associated with these two characters.

A ‘;’ is a line separator, treated as a new-line, so separate instructions can be specified on a single line.

9.25.3.2 Symbols

The character ‘:’ is permitted in identifiers. There are two exceptions to it being treated as any other symbol character: if a symbol begins with ‘:’, it means that the symbol is in the global namespace and that the current prefix should not be prepended to that symbol (see [\[MMIX-prefix\], page 199](#)). The ‘:’ is then not considered part of the symbol. For a symbol in the label position (first on a line), a ‘:’ at the end of a symbol is silently stripped off. A label is permitted, but not required, to be followed by a ‘:’, as with many other assembly formats.

The character ‘@’ in an expression, is a synonym for ‘.’, the current location.

In addition to the common forward and backward local symbol formats (see [Section 5.3 \[Symbol Names\], page 39](#)), they can be specified with upper-case ‘B’ and ‘F’, as in ‘8B’ and ‘9F’. A local label defined for the current position is written with a ‘H’ appended to the number:

```
3H LDB $0,$1,2
```

This and traditional local-label formats cannot be mixed: a label must be defined and referred to using the same format.

There’s a minor caveat: just as for the ordinary local symbols, the local symbols are translated into ordinary symbols using control characters are to hide the ordinal number of the symbol. Unfortunately, these symbols are not translated back in error messages. Thus you may see confusing error messages when local symbols are used. Control characters ‘\003’ (control-C) and ‘\004’ (control-D) are used for the MMIX-specific local-symbol syntax.

The symbol ‘Main’ is handled specially; it is always global.

By defining the symbols ‘`__MMIX.start..text`’ and ‘`__MMIX.start..data`’, the address of respectively the ‘`.text`’ and ‘`.data`’ segments of the final program can be defined, though when linking more than one object file, the code or data in the object file containing the symbol is not guaranteed to be start at that position; just the final executable. See [\[MMIX-loc\], page 197](#).

9.25.3.3 Register names

Local and global registers are specified as ‘\$0’ to ‘\$255’. The recognized special register names are ‘rJ’, ‘rA’, ‘rB’, ‘rC’, ‘rD’, ‘rE’, ‘rF’, ‘rG’, ‘rH’, ‘rI’, ‘rK’, ‘rL’, ‘rM’, ‘rN’, ‘rO’, ‘rP’, ‘rQ’, ‘rR’, ‘rS’, ‘rT’, ‘rU’, ‘rV’, ‘rW’, ‘rX’, ‘rY’, ‘rZ’, ‘rBB’, ‘rTT’, ‘rWW’, ‘rXX’, ‘rYY’ and ‘rZZ’. A leading ‘:’ is optional for special register names.

Local and global symbols can be equated to register names and used in place of ordinary registers.

Similarly for special registers, local and global symbols can be used. Also, symbols equated from numbers and constant expressions are allowed in place of a special reg-

ister, except when either of the options `--no-predefined-syms` and `--fixed-special-register-names` are specified. Then only the special register names above are allowed for the instructions having a special register operand; GET and PUT.

9.25.3.4 Assembler Directives

LOC

The LOC directive sets the current location to the value of the operand field, which may include changing sections. If the operand is a constant, the section is set to either `.data` if the value is `0x2000000000000000` or larger, else it is set to `.text`. Within a section, the current location may only be changed to monotonically higher addresses. A LOC expression must be a previously defined symbol or a “pure” constant.

An example, which sets the label *prev* to the current location, and updates the current location to eight bytes forward:

```
prev LOC @+8
```

When a LOC has a constant as its operand, a symbol `__MMIX.start..text` or `__MMIX.start..data` is defined depending on the address as mentioned above. Each such symbol is interpreted as special by the linker, locating the section at that address. Note that if multiple files are linked, the first object file with that section will be mapped to that address (not necessarily the file with the LOC definition).

LOCAL

Example:

```
LOCAL external_symbol
LOCAL 42
.local asymbol
```

This directive-operation generates a link-time assertion that the operand does not correspond to a global register. The operand is an expression that at link-time resolves to a register symbol or a number. A number is treated as the register having that number. There is one restriction on the use of this directive: the pseudo-directive must be placed in a section with contents, code or data.

IS

The IS directive:

```
asymbol IS an_expression
```

sets the symbol ‘asymbol’ to ‘an_expression’. A symbol may not be set more than once using this directive. Local labels may be set using this directive, for example:

```
5H IS @+4
```

GREG

This directive reserves a global register, gives it an initial value and optionally gives it a symbolic name. Some examples:

```
areg GREG
breg GREG data_value
```

```
GREG data_buffer
.greg creg, another_data_value
```

The symbolic register name can be used in place of a (non-special) register. If a value isn't provided, it defaults to zero. Unless the option `'--no-merge-gregs'` is specified, non-zero registers allocated with this directive may be eliminated by `as`; another register with the same value used in its place. Any of the instructions `'CSWAP'`, `'GO'`, `'LDA'`, `'LDBU'`, `'LDB'`, `'LDHT'`, `'LDOU'`, `'LDO'`, `'LDSF'`, `'LDTU'`, `'LDT'`, `'LDUNC'`, `'LDVTS'`, `'LDWU'`, `'LDW'`, `'PREGO'`, `'PRELD'`, `'PREST'`, `'PUSHGO'`, `'STBU'`, `'STB'`, `'STCO'`, `'STHT'`, `'STOU'`, `'STSF'`, `'STTU'`, `'STT'`, `'STUNC'`, `'SYNCD'`, `'SYNCID'`, can have a value nearby an initial value in place of its second and third operands. Here, "nearby" is defined as within the range 0...255 from the initial value of such an allocated register.

```
buffer1 BYTE 0,0,0,0,0
buffer2 BYTE 0,0,0,0,0
...
GREG buffer1
LDOU $42,buffer2
```

In the example above, the 'Y' field of the `LDOUI` instruction (`LDOU` with a constant Z) will be replaced with the global register allocated for `'buffer1'`, and the 'Z' field will have the value 5, the offset from `'buffer1'` to `'buffer2'`. The result is equivalent to this code:

```
buffer1 BYTE 0,0,0,0,0
buffer2 BYTE 0,0,0,0,0
...
tmpreg GREG buffer1
LDOU $42,tmpreg,(buffer2-buffer1)
```

Global registers allocated with this directive are allocated in order higher-to-lower within a file. Other than that, the exact order of register allocation and elimination is undefined. For example, the order is undefined when more than one file with such directives are linked together. With the options `'-x'` and `'--linker-allocated-gregs'`, `'GREG'` directives for two-operand cases like the one mentioned above can be omitted. Sufficient global registers will then be allocated by the linker.

BYTE

The `'BYTE'` directive takes a series of operands separated by a comma. If an operand is a string (see [Section 3.6.1.1 \[Strings\]](#), [page 29](#)), each character of that string is emitted as a byte. Other operands must be constant expressions without forward references, in the range 0...255. If you need operands having expressions with forward references, use `'.byte'` (see [Section 7.8 \[Byte\]](#), [page 49](#)). An operand can be omitted, defaulting to a zero value.

WYDE TETRA OCTA

The directives `'WYDE'`, `'TETRA'` and `'OCTA'` emit constants of two, four and eight bytes size respectively. Before anything else happens for the directive, the current location is aligned to the respective constant-size boundary. If a label is defined at the beginning of the line, its value will be that after the alignment.

A single operand can be omitted, defaulting to a zero value emitted for the directive. Operands can be expressed as strings (see [Section 3.6.1.1 \[Strings\]](#), [page 29](#)), in which case each character in the string is emitted as a separate constant of the size indicated by the directive.

PREFIX

The ‘PREFIX’ directive sets a symbol name prefix to be prepended to all symbols (except local symbols, see [Section 9.25.3.2 \[MMIX-Symbols\]](#), [page 196](#)), that are not prefixed with ‘:’, until the next ‘PREFIX’ directive. Such prefixes accumulate. For example,

```
PREFIX a
PREFIX b
c IS 0
```

defines a symbol ‘abc’ with the value 0.

BSPEC

ESPEC

A pair of ‘BSPEC’ and ‘ESPEC’ directives delimit a section of special contents (without specified semantics). Example:

```
BSPEC 42
TETRA 1,2,3
ESPEC
```

The single operand to ‘BSPEC’ must be number in the range 0...255. The ‘BSPEC’ number 80 is used by the GNU binutils implementation.

9.25.4 Differences to mmixal

The binutils `as` and `ld` combination has a few differences in function compared to `mmixal` (see [\[mmixsite\]](#), [page 195](#)).

The replacement of a symbol with a GREG-allocated register (see [\[GREG-base\]](#), [page 198](#)) is not handled the exactly same way in `as` as in `mmixal`. This is apparent in the `mmixal` example file `inout.mms`, where different registers with different offsets, eventually yielding the same address, are used in the first instruction. This type of difference should however not affect the function of any program unless it has specific assumptions about the allocated register number.

Line numbers (in the ‘mmo’ object format) are currently not supported.

Expression operator precedence is not that of `mmixal`: operator precedence is that of the C programming language. It’s recommended to use parentheses to explicitly specify wanted operator precedence whenever more than one type of operators are used.

The serialize unary operator `&`, the fractional division operator `‘/’`, the logical not operator `!` and the modulus operator `‘%’` are not available.

Symbols are not global by default, unless the option ‘`--globalize-symbols`’ is passed. Use the ‘`.global`’ directive to globalize symbols (see [Section 7.55 \[Global\]](#), [page 55](#)).

Operand syntax is a bit stricter with `as` than `mmixal`. For example, you can’t say `addu 1,2,3`, instead you must write `addu $1,$2,3`.

You can’t LOC to a lower address than those already visited (i.e., “backwards”).

A LOC directive must come before any emitted code.

Predefined symbols are visible as file-local symbols after use. (In the ELF file, that is—the linked mmo file has no notion of a file-local symbol.)

Some mapping of constant expressions to sections in LOC expressions is attempted, but that functionality is easily confused and should be avoided unless compatibility with `mmixal` is required. A LOC expression to `'0x2000000000000000'` or higher, maps to the `'.data'` section and lower addresses map to the `'.text'` section (see [MMIX-loc], page 197).

The code and data areas are each contiguous. Sparse programs with far-away LOC directives will take up the same amount of space as a contiguous program with zeros filled in the gaps between the LOC directives. If you need sparse programs, you might try and get the wanted effect with a linker script and splitting up the code parts into sections (see Section 7.96 [Section], page 68). Assembly code for this, to be compatible with `mmixal`, would look something like:

```
.if 0
LOC away_expression
.else
.section away,"ax"
.fi
```

`as` will not execute the LOC directive and `mmixal` ignores the lines with `..`. This construct can be used generally to help compatibility.

Symbols can't be defined twice—not even to the same value.

Instruction mnemonics are recognized case-insensitive, though the `'IS'` and `'GREG'` pseudo-operations must be specified in upper-case characters.

There's no unicode support.

The following is a list of programs in `'mmix.tar.gz'`, available at <http://www-cs-faculty.stanford.edu/~> last checked with the version dated 2001-08-25 (md5sum c393470cfc86fac040487d22d2bf0172) that assemble with `mmixal` but do not assemble with `as`:

```
silly.mms    LOC to a previous address.
sim.mms      Redefines symbol 'Done'.
test.mms     Uses the serial operator '&'.
```

9.26 MSP 430 Dependent Features

9.26.1 Options

- m select the mpu arch. Currently has no effect.
- mP enables polymorph instructions handler.
- mQ enables relaxation at assembly time. DANGEROUS!

9.26.2 Syntax

9.26.2.1 Macros

The macro syntax used on the MSP 430 is like that described in the MSP 430 Family Assembler Specification. Normal **as** macros should still work.

Additional built-in macros are:

- llo(exp)** Extracts least significant word from 32-bit expression 'exp'.
- lhi(exp)** Extracts most significant word from 32-bit expression 'exp'.
- hlo(exp)** Extracts 3rd word from 64-bit expression 'exp'.
- hhi(exp)** Extracts 4rd word from 64-bit expression 'exp'.

They normally being used as an immediate source operand.

```
mov #llo(1), r10 ; == mov #1, r10
mov #lhi(1), r10 ; == mov #0, r10
```

9.26.2.2 Special Characters

';' is the line comment character.

The character '\$' in jump instructions indicates current location and implemented only for TI syntax compatibility.

9.26.2.3 Register Names

General-purpose registers are represented by predefined symbols of the form '**rN**' (for global registers), where *N* represents a number between 0 and 15. The leading letters may be in either upper or lower case; for example, '**r13**' and '**R7**' are both valid register names.

Register names '**PC**', '**SP**' and '**SR**' cannot be used as register names and will be treated as variables. Use '**r0**', '**r1**', and '**r2**' instead.

9.26.2.4 Assembler Extensions

- @rN** As destination operand being treated as '**0(rn)**'
- 0(rN)** As source operand being treated as '**@rn**'
- jCOND +N** Skips next *N* bytes followed by jump instruction and equivalent to '**jCOND \$+N+2**'

Also, there are some instructions, which cannot be found in other assemblers. These are branch instructions, which has different opcodes upon jump distance. They all got PC relative addressing mode.

| | |
|-------------------------|---|
| <code>beq label</code> | A polymorph instruction which is ‘ <code>jeq label</code> ’ in case if jump distance within allowed range for cpu’s jump instruction. If not, this unrolls into a sequence of <pre> jne \$+6 br label </pre> |
| <code>bne label</code> | A polymorph instruction which is ‘ <code>jne label</code> ’ or ‘ <code>jeq +4; br label</code> ’ |
| <code>blt label</code> | A polymorph instruction which is ‘ <code>jlt label</code> ’ or ‘ <code>jge +4; br label</code> ’ |
| <code>bltn label</code> | A polymorph instruction which is ‘ <code>jn label</code> ’ or ‘ <code>jn +2; jmp +4; br label</code> ’ |
| <code>bltu label</code> | A polymorph instruction which is ‘ <code>jlo label</code> ’ or ‘ <code>jhs +2; br label</code> ’ |
| <code>bge label</code> | A polymorph instruction which is ‘ <code>jge label</code> ’ or ‘ <code>jlt +4; br label</code> ’ |
| <code>bgeu label</code> | A polymorph instruction which is ‘ <code>jhs label</code> ’ or ‘ <code>jlo +4; br label</code> ’ |
| <code>bgt label</code> | A polymorph instruction which is ‘ <code>jeq +2; jge label</code> ’ or ‘ <code>jeq +6; jlt +4; br label</code> ’ |
| <code>bgtu label</code> | A polymorph instruction which is ‘ <code>jeq +2; jhs label</code> ’ or ‘ <code>jeq +6; jlo +4; br label</code> ’ |
| <code>bleu label</code> | A polymorph instruction which is ‘ <code>jeq label; jlo label</code> ’ or ‘ <code>jeq +2; jhs +4; br label</code> ’ |
| <code>ble label</code> | A polymorph instruction which is ‘ <code>jeq label; jlt label</code> ’ or ‘ <code>jeq +2; jge +4; br label</code> ’ |
| <code>jump label</code> | A polymorph instruction which is ‘ <code>jmp label</code> ’ or ‘ <code>br label</code> ’ |

9.26.3 Floating Point

The MSP 430 family uses IEEE 32-bit floating-point numbers.

9.26.4 MSP 430 Machine Directives

| | |
|------------------------|--|
| <code>.file</code> | This directive is ignored; it is accepted for compatibility with other MSP 430 assemblers. <i>Warning:</i> in other versions of the GNU assembler, <code>.file</code> is used for the directive called <code>.app-file</code> in the MSP 430 support. |
| <code>.line</code> | This directive is ignored; it is accepted for compatibility with other MSP 430 assemblers. |
| <code>.arch</code> | Currently this directive is ignored; it is accepted for compatibility with other MSP 430 assemblers. |
| <code>.profiler</code> | This directive instructs assembler to add new profile entry to the object file. |

9.26.5 Opcodes

as implements all the standard MSP 430 opcodes. No additional pseudo-instructions are needed on this family.

For information on the 430 machine instruction set, see *MSP430 User's Manual, document slau049d*, Texas Instrument, Inc.

9.26.6 Profiling Capability

It is a performance hit to use gcc's profiling approach for this tiny target. Even more – jtag hardware facility does not perform any profiling functions. However we've got gdb's built-in simulator where we can do anything.

We define new section `‘.profiler’` which holds all profiling information. We define new pseudo operation `‘.profiler’` which will instruct assembler to add new profile entry to the object file. Profile should take place at the present address.

Pseudo operation format:

```
‘.profiler flags,function_to_profile [, cycle_corrector, extra]’
```

where:

`‘flags’` is a combination of the following characters:

| | |
|----------|--|
| s | function entry |
| x | function exit |
| i | function is in init section |
| f | function is in fini section |
| l | library call |
| c | libc standard call |
| d | stack value demand |
| I | interrupt service routine |
| P | prologue start |
| p | prologue end |
| E | epilogue start |
| e | epilogue end |
| j | long jump / sjlj unwind |
| a | an arbitrary code fragment |
| t | extra parameter saved (a constant value like frame size) |

function_to_profile

a function address

cycle_corrector

a value which should be added to the cycle counter, zero if omitted.

extra

any extra parameter, zero if omitted.

For example:

```
.global fxx
.type fxx,@function
fxx:
.LFrameOffset_fxx=0x08
.profiler "scdP", fxx      ; function entry.
    ; we also demand stack value to be saved
    push r11
    push r10
    push r9
    push r8
.profiler "cdpt",fxx,0, .LFrameOffset_fxx ; check stack value at this point
    ; (this is a prologue end)
    ; note, that spare var filled with
    ; the farne size
    mov r15,r8
...
.profiler cdE,fxx          ; check stack
    pop r8
    pop r9
    pop r10
    pop r11
.profiler xcde,fxx,3       ; exit adds 3 to the cycle counter
    ret                   ; cause 'ret' insn takes 3 cycles
```

9.27 PDP-11 Dependent Features

9.27.1 Options

The PDP-11 version of `as` has a rich set of machine dependent options.

9.27.1.1 Code Generation Options

`-mpic` | `-mno-pic`

Generate position-independent (or position-dependent) code.

The default is to generate position-independent code.

9.27.1.2 Instruction Set Extension Options

These options enables or disables the use of extensions over the base line instruction set as introduced by the first PDP-11 CPU: the KA11. Most options come in two variants: a `-mextension` that enables *extension*, and a `-mno-extension` that disables *extension*.

The default is to enable all extensions.

`-mall` | `-mall-extensions`

Enable all instruction set extensions.

`-mno-extensions`

Disable all instruction set extensions.

`-mcis` | `-mno-cis`

Enable (or disable) the use of the commercial instruction set, which consists of these instructions: ADDNI, ADDN, ADDPI, ADDP, ASHNI, ASHN, ASHPI, ASHP, CMPCI, CMPC, CMPNI, CMPN, CMPPI, CMPP, CVTLNI, CVTLN, CVTLPI, CVTLP, CVTNLI, CVTNL, CVTNPI, CVTNP, CVTPLI, CVTPL, CVTPNI, CVTPN, DIVPI, DIVP, L2DR, L3DR, LOCCI, LOCC, MATCI, MATC, MOVCI, MOVN, MOVRCI, MOVRC, MOVTNI, MOVTN, MULPI, MULP, SCANCI, SCANC, SKPCI, SKPC, SPANCI, SPANC, SUBNI, SUBN, SUBPI, and SUBP.

`-mcsn` | `-mno-csn`

Enable (or disable) the use of the CSN instruction.

`-meis` | `-mno-eis`

Enable (or disable) the use of the extended instruction set, which consists of these instructions: ASHC, ASH, DIV, MARK, MUL, RTT, SOB SXT, and XOR.

`-mfis` | `-mkev11`

`-mno-fis` | `-mno-kev11`

Enable (or disable) the use of the KEV11 floating-point instructions: FADD, FDIV, FMUL, and FSUB.

`-mfpp` | `-mfpu` | `-mfp-11`

`-mno-fpp` | `-mno-fpu` | `-mno-fp-11`

Enable (or disable) the use of FP-11 floating-point instructions: ABSF, ADDF, CFCC, CLRF, CMPF, DIVF, LDCFF, LDCIF, LDEXP, LDF, LDFPS, MODF, MULF, NEGF, SETD, SETF, SETI, SETL, STCFF, STCFI, STEXP, STF, STFPS, STST, SUBF, and TSTF.

-mlimited-eis | -mno-limited-eis

Enable (or disable) the use of the limited extended instruction set: MARK, RTT, SOB, SXT, and XOR.

The -mno-limited-eis options also implies -mno-eis.

-mmfpt | -mno-mfpt

Enable (or disable) the use of the MFPT instruction.

-mmultiproc | -mno-multiproc

Enable (or disable) the use of multiprocessor instructions: TSTSET and WRTLCK.

-mmxps | -mno-mxps

Enable (or disable) the use of the MFPS and MTPS instructions.

-mspl | -mno-spl

Enable (or disable) the use of the SPL instruction.

Enable (or disable) the use of the microcode instructions: LDUB, MED, and XFC.

9.27.1.3 CPU Model Options

These options enable the instruction set extensions supported by a particular CPU, and disables all other extensions.

-mka11 KA11 CPU. Base line instruction set only.

-mkb11 KB11 CPU. Enable extended instruction set and SPL.

-mkd11a KD11-A CPU. Enable limited extended instruction set.

-mkd11b KD11-B CPU. Base line instruction set only.

-mkd11d KD11-D CPU. Base line instruction set only.

-mkd11e KD11-E CPU. Enable extended instruction set, MFPS, and MTPS.

-mkd11f | -mkd11h | -mkd11q

KD11-F, KD11-H, or KD11-Q CPU. Enable limited extended instruction set, MFPS, and MTPS.

-mkd11k KD11-K CPU. Enable extended instruction set, LDUB, MED, MFPS, MFPT, MTPS, and XFC.

-mkd11z KD11-Z CPU. Enable extended instruction set, CSM, MFPS, MFPT, MTPS, and SPL.

-mf11 F11 CPU. Enable extended instruction set, MFPS, MFPT, and MTPS.

-mj11 J11 CPU. Enable extended instruction set, CSM, MFPS, MFPT, MTPS, SPL, TSTSET, and WRTLCK.

-mt11 T11 CPU. Enable limited extended instruction set, MFPS, and MTPS.

9.27.1.4 Machine Model Options

These options enable the instruction set extensions supported by a particular machine model, and disables all other extensions.

-m11/03 Same as -mkd11f.

```

-m11/04    Same as -mkd11d.
-m11/05 | -m11/10
           Same as -mkd11b.
-m11/15 | -m11/20
           Same as -mka11.
-m11/21    Same as -mt11.
-m11/23 | -m11/24
           Same as -mf11.
-m11/34    Same as -mkd11e.
-m11/34a   Same as -mkd11e -mfpp.
-m11/35 | -m11/40
           Same as -mkd11a.
-m11/44    Same as -mkd11z.
-m11/45 | -m11/50 | -m11/55 | -m11/70
           Same as -mkb11.
-m11/53 | -m11/73 | -m11/83 | -m11/84 | -m11/93 | -m11/94
           Same as -mj11.
-m11/60    Same as -mkd11k.

```

9.27.2 Assembler Directives

The PDP-11 version of `as` has a few machine dependent assembler directives.

```

.bss      Switch to the bss section.
.even     Align the location counter to an even number.

```

9.27.3 PDP-11 Assembly Language Syntax

`as` supports both DEC syntax and BSD syntax. The only difference is that in DEC syntax, a `#` character is used to denote an immediate constants, while in BSD syntax the character for this purpose is `$`.

general-purpose registers are named `r0` through `r7`. Mnemonic alternatives for `r6` and `r7` are `sp` and `pc`, respectively.

Floating-point registers are named `ac0` through `ac3`, or alternatively `fr0` through `fr3`.

Comments are started with a `#` or a `/` character, and extend to the end of the line. (FIXME: clash with immediates?)

9.27.4 Instruction Naming

Some instructions have alternative names.

```

BCC      BHIS
BCS      BLO
L2DR     L2D
L3DR     L3D
SYS      TRAP

```

9.27.5 Synthetic Instructions

The JBR and JCC synthetic instructions are not supported yet.

9.28 picoJava Dependent Features

9.28.1 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.

9.29 PowerPC Dependent Features

9.29.1 Options

The PowerPC 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.

The following table lists all available PowerPC options.

| | |
|---|---|
| <code>-a32</code> | Generate ELF32 or XCOFF32. |
| <code>-a64</code> | Generate ELF64 or XCOFF64. |
| <code>-K PIC</code> | Set <code>EF_PPC_RELOCATABLE_LIB</code> in ELF flags. |
| <code>-mpwrx</code> <code>-mpwr2</code> | Generate code for POWER/2 (RIOS2). |
| <code>-mpwr</code> | Generate code for POWER (RIOS1) |
| <code>-m601</code> | Generate code for PowerPC 601. |
| <code>-mppc</code> , <code>-mppc32</code> , <code>-m603</code> , <code>-m604</code> | Generate code for PowerPC 603/604. |
| <code>-m403</code> , <code>-m405</code> | Generate code for PowerPC 403/405. |
| <code>-m440</code> | Generate code for PowerPC 440. BookE and some 405 instructions. |
| <code>-m464</code> | Generate code for PowerPC 464. |
| <code>-m476</code> | Generate code for PowerPC 476. |
| <code>-m7400</code> , <code>-m7410</code> , <code>-m7450</code> , <code>-m7455</code> | Generate code for PowerPC 7400/7410/7450/7455. |
| <code>-m750cl</code> | Generate code for PowerPC 750CL. |
| <code>-mppc64</code> , <code>-m620</code> | Generate code for PowerPC 620/625/630. |
| <code>-me500</code> , <code>-me500x2</code> | Generate code for Motorola e500 core complex. |
| <code>-me500mc</code> | Generate code for Freescale e500mc core complex. |
| <code>-me500mc64</code> | Generate code for Freescale e500mc64 core complex. |
| <code>-mspe</code> | Generate code for Motorola SPE instructions. |
| <code>-mtitan</code> | Generate code for AppliedMicro Titan core complex. |
| <code>-mppc64bridge</code> | Generate code for PowerPC 64, including bridge insns. |
| <code>-mbooke</code> | Generate code for 32-bit BookE. |

| | |
|---|---|
| <code>-ma2</code> | Generate code for A2 architecture. |
| <code>-me300</code> | Generate code for PowerPC e300 family. |
| <code>-maltivec</code> | Generate code for processors with AltiVec instructions. |
| <code>-mvsx</code> | Generate code for processors with Vector-Scalar (VSX) instructions. |
| <code>-mpower4, -mpwr4</code> | Generate code for Power4 architecture. |
| <code>-mpower5, -mpwr5, -mpwr5x</code> | Generate code for Power5 architecture. |
| <code>-mpower6, -mpwr6</code> | Generate code for Power6 architecture. |
| <code>-mpower7, -mpwr7</code> | Generate code for Power7 architecture. |
| <code>-mcell</code> | Generate code for Cell Broadband Engine architecture. |
| <code>-mcom</code> | Generate code Power/PowerPC common instructions. |
| <code>-many</code> | Generate code for any architecture (PWR/PWRX/PPC). |
| <code>-mregnames</code> | Allow symbolic names for registers. |
| <code>-mno-regnames</code> | Do not allow symbolic names for registers. |
| <code>-mrelocatable</code> | Support for GCC's <code>-mrelocatable</code> option. |
| <code>-mrelocatable-lib</code> | Support for GCC's <code>-mrelocatable-lib</code> option. |
| <code>-memb</code> | Set <code>PPC_EMB</code> bit in ELF flags. |
| <code>-mlittle, -mlittle-endian, -le</code> | Generate code for a little endian machine. |
| <code>-mbig, -mbig-endian, -be</code> | Generate code for a big endian machine. |
| <code>-msolaris</code> | Generate code for Solaris. |
| <code>-mno-solaris</code> | Do not generate code for Solaris. |
| <code>-nops=<i>count</i></code> | If an alignment directive inserts more than <i>count</i> nops, put a branch at the beginning to skip execution of the nops. |

9.29.2 PowerPC Assembler Directives

A number of assembler directives are available for PowerPC. The following table is far from complete.

`.machine "string"`

This directive allows you to change the machine for which code is generated. "string" may be any of the -m cpu selection options (without the -m) enclosed in double quotes, "push", or "pop". `.machine "push"` saves the currently selected cpu, which may be restored with `.machine "pop"`.

9.30 RX Dependent Features

9.30.1 RX Options

The Renesas RX port of `as` has a few target specific command line options:

`-m32bit-doubles`

This option controls the ABI and indicates to use a 32-bit float ABI. It has no effect on the assembled instructions, but it does influence the behaviour of the `‘.double’` pseudo-op. This is the default.

`-m64bit-doubles`

This option controls the ABI and indicates to use a 64-bit float ABI. It has no effect on the assembled instructions, but it does influence the behaviour of the `‘.double’` pseudo-op.

`-mbig-endian`

This option controls the ABI and indicates to use a big-endian data ABI. It has no effect on the assembled instructions, but it does influence the behaviour of the `‘.short’`, `‘.hword’`, `‘.int’`, `‘.word’`, `‘.long’`, `‘.quad’` and `‘.octa’` pseudo-ops.

`-mlittle-endian`

This option controls the ABI and indicates to use a little-endian data ABI. It has no effect on the assembled instructions, but it does influence the behaviour of the `‘.short’`, `‘.hword’`, `‘.int’`, `‘.word’`, `‘.long’`, `‘.quad’` and `‘.octa’` pseudo-ops. This is the default.

`-muse-conventional-section-names`

This option controls the default names given to the code (`.text`), initialised data (`.data`) and uninitialised data sections (`.bss`).

`-muse-renesas-section-names`

This option controls the default names given to the code (`.P`), initialised data (`.D_1`) and uninitialised data sections (`.B_1`). This is the default.

`-msmall-data-limit`

This option tells the assembler that the small data limit feature of the RX port of GCC is being used. This results in the assembler generating an undefined reference to a symbol called `__gp` for use by the relocations that are needed to support the small data limit feature. This option is not enabled by default as it would otherwise pollute the symbol table.

9.30.2 Symbolic Operand Modifiers

The assembler supports several modifiers when using symbol addresses in RX instruction operands. The general syntax is the following:

```
%modifier(symbol)
```

```
%gp
```

9.30.3 Assembler Directives

The RX version of `as` has the following specific assembler directives:

`.3byte` Inserts a 3-byte value into the output file at the current location.

9.30.4 Floating Point

The floating point formats generated by directives are these.

`.float` Single precision (32-bit) floating point constants.

`.double` If the ‘`-m64bit-doubles`’ command line option has been specified then the `double` directive generates `double` precision (64-bit) floating point constants, otherwise it generates `single` precision (32-bit) floating point constants. To force the generation of 64-bit floating point constants use the `dc.d` directive instead.

9.31 IBM S/390 Dependent Features

The s390 version of **as** supports two architectures modes and seven chip levels. The architecture modes are the Enterprise System Architecture (ESA) and the newer z/Architecture mode. The chip levels are g5, g6, z900, z990, z9-109, z9-ec, z10 and z196.

9.31.1 Options

The following table lists all available s390 specific options:

-m31 | -m64

Select 31- or 64-bit ABI implying a word size of 32- or 64-bit.

These options are only available with the ELF object file format, and require that the necessary BFD support has been included (on a 31-bit platform you must add `-enable-64-bit-bfd` on the call to the configure script to enable 64-bit usage and use s390x as target platform).

-mesa | -mzarch

Select the architecture mode, either the Enterprise System Architecture (esa) mode or the z/Architecture mode (zarch).

The 64-bit instructions are only available with the z/Architecture mode. The combination of `'-m64'` and `'-mesa'` results in a warning message.

-march=CPU

This option specifies the target processor. The following processor names are recognized: g5, g6, z900, z990, z9-109, z9-ec, z10 and z196. Assembling an instruction that is not supported on the target processor results in an error message. Do not specify g5 or g6 with `'-mzarch'`.

-mregnames

Allow symbolic names for registers.

-mno-regnames

Do not allow symbolic names for registers.

-mwarn-areg-zero

Warn whenever the operand for a base or index register has been specified but evaluates to zero. This can indicate the misuse of general purpose register 0 as an address register.

9.31.2 Special Characters

`'#'` is the line comment character.

9.31.3 Instruction syntax

The assembler syntax closely follows the syntax outlined in Enterprise Systems Architecture/390 Principles of Operation (SA22-7201) and the z/Architecture Principles of Operation (SA22-7832).

Each instruction has two major parts, the instruction mnemonic and the instruction operands. The instruction format varies.

9.31.3.1 Register naming

The **as** recognizes a number of predefined symbols for the various processor registers. A register specification in one of the instruction formats is an unsigned integer between 0 and 15. The specific instruction and the position of the register in the instruction format denotes the type of the register. The register symbols are prefixed with ‘%’:

%rN the 16 general purpose registers, $0 \leq N \leq 15$

%fN the 16 floating point registers, $0 \leq N \leq 15$

%aN the 16 access registers, $0 \leq N \leq 15$

%cN the 16 control registers, $0 \leq N \leq 15$

%lit an alias for the general purpose register **%r13**

%sp an alias for the general purpose register **%r15**

9.31.3.2 Instruction Mnemonics

All instructions documented in the Principles of Operation are supported with the mnemonic and order of operands as described. The instruction mnemonic identifies the instruction format ([Section 9.31.3.4 \[s390 Formats\], page 219](#)) and the specific operation code for the instruction. For example, the ‘**lr**’ mnemonic denotes the instruction format ‘**RR**’ with the operation code ‘**0x18**’.

The definition of the various mnemonics follows a scheme, where the first character usually hint at the type of the instruction:

- a add instruction, for example ‘**al**’ for add logical 32-bit
- b branch instruction, for example ‘**bc**’ for branch on condition
- c compare or convert instruction, for example ‘**cr**’ for compare register 32-bit
- d divide instruction, for example ‘**d1r**’ divide logical register 64-bit to 32-bit
- i insert instruction, for example ‘**ic**’ insert character
- l load instruction, for example ‘**ltr**’ load and test register
- mv move instruction, for example ‘**mvc**’ move character
- m multiply instruction, for example ‘**mh**’ multiply halfword
- n and instruction, for example ‘**ni**’ and immediate

- o or instruction, for example ‘oc’ or character
- sla, sll shift left single instruction
- sra, srl shift right single instruction
- st store instruction, for example ‘stm’ store multiple
- s subtract instruction, for example ‘slr’ subtract
 logical 32-bit
- t test or translate instruction, of example ‘tm’ test under mask
- x exclusive or instruction, for example ‘xc’ exclusive or
 character

Certain characters at the end of the mnemonic may describe a property of the instruction:

- c the instruction uses a 8-bit character operand
- f the instruction extends a 32-bit operand to 64 bit
- g the operands are treated as 64-bit values
- h the operand uses a 16-bit halfword operand
- i the instruction uses an immediate operand
- l the instruction uses unsigned, logical operands
- m the instruction uses a mask or operates on multiple values
- r if r is the last character, the instruction operates on registers
- y the instruction uses 20-bit displacements

There are many exceptions to the scheme outlined in the above lists, in particular for the privileged instructions. For non-privileged instruction it works quite well, for example the instruction ‘clgfr’ c: compare instruction, l: unsigned operands, g: 64-bit operands, f: 32- to 64-bit extension, r: register operands. The instruction compares an 64-bit value in a register with the zero extended 32-bit value from a second register. For a complete list of all mnemonics see appendix B in the Principles of Operation.

9.31.3.3 Instruction Operands

Instruction operands can be grouped into three classes, operands located in registers, immediate operands, and operands in storage.

A register operand can be located in general, floating-point, access, or control register. The register is identified by a four-bit field. The field containing the register operand is called the R field.

Immediate operands are contained within the instruction and can have 8, 16 or 32 bits. The field containing the immediate operand is called the I field. Dependent on the instruction the I field is either signed or unsigned.

A storage operand consists of an address and a length. The address of a storage operands can be specified in any of these ways:

- The content of a single general R
- The sum of the content of a general register called the base register B plus the content of a displacement field D
- The sum of the contents of two general registers called the index register X and the base register B plus the content of a displacement field
- The sum of the current instruction address and a 32-bit signed immediate field multiplied by two.

The length of a storage operand can be:

- Implied by the instruction
- Specified by a bitmask
- Specified by a four-bit or eight-bit length field L
- Specified by the content of a general register

The notation for storage operand addresses formed from multiple fields is as follows:

Dn(Bn) the address for operand number n is formed from the content of general register Bn called the base register and the displacement field Dn.

Dn(Xn,Bn) the address for operand number n is formed from the content of general register Xn called the index register, general register Bn called the base register and the displacement field Dn.

Dn(Ln,Bn) the address for operand number n is formed from the content of general register Bn called the base register and the displacement field Dn. The length of the operand n is specified by the field Ln.

The base registers Bn and the index registers Xn of a storage operand can be skipped. If Bn and Xn are skipped, a zero will be stored to the operand field. The notation changes as follows:

| full notation | short notation |
|---------------|----------------|
| Dn(0,Bn) | Dn(Bn) |
| Dn(0,0) | Dn |

| | |
|----------|--------|
| Dn(0) | Dn |
| Dn(Ln,0) | Dn(Ln) |

9.31.3.4 Instruction Formats

The Principles of Operation manuals lists 26 instruction formats where some of the formats have multiple variants. For the ‘`.insn`’ pseudo directive the assembler recognizes some of the formats. Typically, the most general variant of the instruction format is used by the ‘`.insn`’ directive.


The following table lists the abbreviations used in the table of instruction formats:

| | |
|---------------|---|
| OpCode / OpCd | Part of the op code. |
| Bx | Base register number for operand x. |
| Dx | Displacement for operand x. |
| DLx | Displacement lower 12 bits for operand x. |
| DHx | Displacement higher 8-bits for operand x. |
| Rx | Register number for operand x. |
| Xx | Index register number for operand x. |
| Ix | Signed immediate for operand x. |
| Ux | Unsigned immediate for operand x. |


An instruction is two, four, or six bytes in length and must be aligned on a 2 byte boundary. The first two bits of the instruction specify the length of the instruction, 00 indicates a two byte instruction, 01 and 10 indicates a four byte instruction, and 11 indicates a six byte instruction.

The following table lists the s390 instruction formats that are available with the ‘`.insn`’ pseudo directive:


E format




RI format: <insn> R1,I2




RIE format: <insn> R1,R3,I2




RIL format: <insn> R1,I2




RILU format: <insn> R1,U2




RIS format: <insn> R1,I2,M3,D4(B4)




RR format: <insn> R1,R2




RRE format: <insn> R1,R2




RRF format: <insn> R1,R2,R3,M4



RRS format: <insn> R1,R2,M3,D4(B4)



RS format: <insn> R1,R3,D2(B2)



```

+-----+-----+-----+-----+-----+
0         8      12      16      20             31

```

RSE format: <insn> R1,R3,D2(B2)

```

+-----+-----+-----+-----+-----+-----+-----+
| OpCode | R1 | R3 | B2 |      D2      |////////| OpCode |
+-----+-----+-----+-----+-----+-----+-----+
0         8      12      16      20             32      40      47

```

RSI format: <insn> R1,R3,I2

```

+-----+-----+-----+-----+-----+-----+
| OpCode | R1 | R3 |      I2      |
+-----+-----+-----+-----+-----+-----+
0         8      12      16                               47

```

RSY format: <insn> R1,R3,D2(B2)

```

+-----+-----+-----+-----+-----+-----+-----+
| OpCode | R1 | R3 | B2 |      DL2      |      DH2      | OpCode |
+-----+-----+-----+-----+-----+-----+-----+
0         8      12      16      20             32      40      47

```

RX format: <insn> R1,D2(X2,B2)

```

+-----+-----+-----+-----+-----+
| OpCode | R1 | X2 | B2 |      D2      |
+-----+-----+-----+-----+-----+
0         8      12      16      20             31

```

RXE format: <insn> R1,D2(X2,B2)

```

+-----+-----+-----+-----+-----+-----+-----+
| OpCode | R1 | X2 | B2 |      D2      |////////| OpCode |
+-----+-----+-----+-----+-----+-----+-----+
0         8      12      16      20             32      40      47

```

RXF format: <insn> R1,R3,D2(X2,B2)

```

+-----+-----+-----+-----+-----+-----+-----+
| OpCode | R3 | X2 | B2 |      D2      | R1 |////| OpCode |
+-----+-----+-----+-----+-----+-----+-----+
0         8      12      16      20             32      36      40      47

```

RXY format: <insn> R1,D2(X2,B2)

```

+-----+-----+-----+-----+-----+-----+-----+
| OpCode | R1 | X2 | B2 |      DL2      |      DH2      | OpCode |
+-----+-----+-----+-----+-----+-----+-----+
0         8      12      16      20             32      36      40      47

```

S format: <insn> D2(B2)

```

+-----+-----+-----+
|      OpCode      | B2 |      D2      |
+-----+-----+-----+
0             16      20             31

```

SI format: <insn> D1(B1),I2

```

+-----+

```

| | | | | |
|---------|---------|---------|---------|---------|
| OpCode | I2 | B1 | D1 | |
| +-----+ | +-----+ | +-----+ | +-----+ | +-----+ |
| 0 | 8 | 16 | 20 | 31 |

SIY format: <insn> D1(B1),U2

| | | | | | |
|---------|---------|---------|---------|---------|----------|
| +-----+ | +-----+ | +-----+ | +-----+ | +-----+ | +-----+ |
| OpCode | I2 | B1 | DL1 | DH1 | OpCode |
| +-----+ | +-----+ | +-----+ | +-----+ | +-----+ | +-----+ |
| 0 | 8 | 16 | 20 | 32 | 36 40 47 |

SIL format: <insn> D1(B1),I2

| | | | | |
|---------|---------|---------|---------|---------|
| +-----+ | +-----+ | +-----+ | +-----+ | +-----+ |
| | OpCode | B1 | D1 | I2 |
| +-----+ | +-----+ | +-----+ | +-----+ | +-----+ |
| 0 | | 16 | 20 | 32 47 |

SS format: <insn> D1(R1,B1),D2(B3),R3

| | | | | | |
|---------|---------|---------|---------|---------|----------|
| +-----+ | +-----+ | +-----+ | +-----+ | +-----+ | +-----+ |
| OpCode | R1 | R3 | B1 | D1 | B2 |
| +-----+ | +-----+ | +-----+ | +-----+ | +-----+ | +-----+ |
| 0 | 8 | 12 | 16 | 20 | 32 36 47 |

SSE format: <insn> D1(B1),D2(B2)

| | | | | |
|---------|---------|---------|---------|-------------|
| +-----+ | +-----+ | +-----+ | +-----+ | +-----+ |
| | OpCode | B1 | D1 | B2 |
| +-----+ | +-----+ | +-----+ | +-----+ | +-----+ |
| 0 | 8 | 12 | 16 | 20 32 36 47 |

SSF format: <insn> D1(B1),D2(B2),R3

| | | | | | |
|---------|---------|---------|---------|---------|----------|
| +-----+ | +-----+ | +-----+ | +-----+ | +-----+ | +-----+ |
| OpCode | R3 | OpCd | B1 | D1 | B2 |
| +-----+ | +-----+ | +-----+ | +-----+ | +-----+ | +-----+ |
| 0 | 8 | 12 | 16 | 20 | 32 36 47 |

For the complete list of all instruction format variants see the Principles of Operation manuals.

9.31.3.5 Instruction Aliases

A specific bit pattern can have multiple mnemonics, for example the bit pattern '0xa7000000' has the mnemonics 'tmh' and 'tmlh'. In addition, there are a number of mnemonics recognized by **as** that are not present in the Principles of Operation. These are the short forms of the branch instructions, where the condition code mask operand is encoded in the mnemonic. This is relevant for the branch instructions, the compare and branch instructions, and the compare and trap instructions.

For the branch instructions there are 20 condition code strings that can be used as part of the mnemonic in place of a mask operand in the instruction format:

instruction

short form

| | | | |
|------|--------------|-------|-----------|
| bcr | M1,R2 | b<m>r | R2 |
| bc | M1,D2(X2,B2) | b<m> | D2(X2,B2) |
| brc | M1,I2 | j<m> | I2 |
| brcl | M1,I2 | jg<m> | I2 |

In the mnemonic for a branch instruction the condition code string <m> can be any of the following:

| | |
|-----|----------------------------------|
| o | jump on overflow / if ones |
| h | jump on A high |
| p | jump on plus |
| nle | jump on not low or equal |
| l | jump on A low |
| m | jump on minus |
| nhe | jump on not high or equal |
| lh | jump on low or high |
| ne | jump on A not equal B |
| nz | jump on not zero / if not zeros |
| e | jump on A equal B |
| z | jump on zero / if zeroes |
| nlh | jump on not low or high |
| he | jump on high or equal |
| nl | jump on A not low |
| nm | jump on not minus / if not mixed |
| le | jump on low or equal |
| nh | jump on A not high |

np jump on not plus

no jump on not overflow / if not ones

For the compare and branch, and compare and trap instructions there are 12 condition code strings that can be used as part of the mnemonic in place of a mask operand in the instruction format:

| instruction | short form |
|-------------------------|-------------------------|
| crb R1,R2,M3,D4(B4) | crb<m> R1,R2,D4(B4) |
| cgrb R1,R2,M3,D4(B4) | cgrb<m> R1,R2,D4(B4) |
| crj R1,R2,M3,I4 | crj<m> R1,R2,I4 |
| cgrj R1,R2,M3,I4 | cgrj<m> R1,R2,I4 |
| cib R1,I2,M3,D4(B4) | cib<m> R1,I2,D4(B4) |
| cgib R1,I2,M3,D4(B4) | cgib<m> R1,I2,D4(B4) |
| cij R1,I2,M3,I4 | cij<m> R1,I2,I4 |
| cgij R1,I2,M3,I4 | cgij<m> R1,I2,I4 |
| crt R1,R2,M3 | crt<m> R1,R2 |
| cgrt R1,R2,M3 | cgrt<m> R1,R2 |
| cit R1,I2,M3 | cit<m> R1,I2 |
| cgit R1,I2,M3 | cgit<m> R1,I2 |
| clrb R1,R2,M3,D4(B4) | clrb<m> R1,R2,D4(B4) |
| clgrb R1,R2,M3,D4(B4) | clgrb<m> R1,R2,D4(B4) |
| clrj R1,R2,M3,I4 | clrj<m> R1,R2,I4 |
| clgrj R1,R2,M3,I4 | clgrj<m> R1,R2,I4 |
| clib R1,I2,M3,D4(B4) | clib<m> R1,I2,D4(B4) |
| clgib R1,I2,M3,D4(B4) | clgib<m> R1,I2,D4(B4) |

| | |
|--------------------------------|--------------------------------------|
| <code>clij R1,I2,M3,I4</code> | <code>clij<m> R1,I2,I4</code> |
| <code>clgij R1,I2,M3,I4</code> | <code>clgij<m> R1,I2,I4</code> |
| <code>clrt R1,R2,M3</code> | <code>clrt<m> R1,R2</code> |
| <code>clgrt R1,R2,M3</code> | <code>clgrt<m> R1,R2</code> |
| <code>clfit R1,I2,M3</code> | <code>clfit<m> R1,I2</code> |
| <code>clgit R1,I2,M3</code> | <code>clgit<m> R1,I2</code> |

In the mnemonic for a compare and branch and compare and trap instruction the condition code string `<m>` can be any of the following:

| | |
|------------------|---------------------------|
| <code>h</code> | jump on A high |
| <code>nle</code> | jump on not low or equal |
| <code>l</code> | jump on A low |
| <code>nhe</code> | jump on not high or equal |
| <code>ne</code> | jump on A not equal B |
| <code>lh</code> | jump on low or high |
| <code>e</code> | jump on A equal B |
| <code>nlh</code> | jump on not low or high |
| <code>nl</code> | jump on A not low |
| <code>he</code> | jump on high or equal |
| <code>nh</code> | jump on A not high |
| <code>le</code> | jump on low or equal |

9.31.3.6 Instruction Operand Modifier

If a symbol modifier is attached to a symbol in an expression for an instruction operand field, the symbol term is replaced with a reference to an object in the global offset table (GOT) or the procedure linkage table (PLT). The following expressions are allowed: `'symbol@modifier + constant'`, `'symbol@modifier + label + constant'`, and `'symbol@modifier - label + constant'`. The term `'symbol'` is the symbol that will be entered into the GOT or PLT, `'label'` is a local label, and `'constant'` is an arbitrary expression that the assembler can evaluate to a constant value.

The term ‘(symbol + constant1)@modifier +/- label + constant2’ is also accepted but a warning message is printed and the term is converted to ‘symbol@modifier +/- label + constant1 + constant2’.

@got

@got12 The @got modifier can be used for displacement fields, 16-bit immediate fields and 32-bit pc-relative immediate fields. The @got12 modifier is synonym to @got. The symbol is added to the GOT. For displacement fields and 16-bit immediate fields the symbol term is replaced with the offset from the start of the GOT to the GOT slot for the symbol. For a 32-bit pc-relative field the pc-relative offset to the GOT slot from the current instruction address is used.

@gotent The @gotent modifier can be used for 32-bit pc-relative immediate fields. The symbol is added to the GOT and the symbol term is replaced with the pc-relative offset from the current instruction to the GOT slot for the symbol.

@gotoff The @gotoff modifier can be used for 16-bit immediate fields. The symbol term is replaced with the offset from the start of the GOT to the address of the symbol.

@gotplt The @gotplt modifier can be used for displacement fields, 16-bit immediate fields, and 32-bit pc-relative immediate fields. A procedure linkage table entry is generated for the symbol and a jump slot for the symbol is added to the GOT. For displacement fields and 16-bit immediate fields the symbol term is replaced with the offset from the start of the GOT to the jump slot for the symbol. For a 32-bit pc-relative field the pc-relative offset to the jump slot from the current instruction address is used.

@plt The @plt modifier can be used for 16-bit and 32-bit pc-relative immediate fields. A procedure linkage table entry is generated for the symbol. The symbol term is replaced with the relative offset from the current instruction to the PLT entry for the symbol.

@pltoff The @pltoff modifier can be used for 16-bit immediate fields. The symbol term is replaced with the offset from the start of the PLT to the address of the symbol.

@gotntpoff

The @gotntpoff modifier can be used for displacement fields. The symbol is added to the static TLS block and the negated offset to the symbol in the static TLS block is added to the GOT. The symbol term is replaced with the offset to the GOT slot from the start of the GOT.

@indntpoff

The @indntpoff modifier can be used for 32-bit pc-relative immediate fields. The symbol is added to the static TLS block and the negated offset to the symbol in the static TLS block is added to the GOT. The symbol term is replaced with the pc-relative offset to the GOT slot from the current instruction address.

For more information about the thread local storage modifiers ‘gotntpoff’ and ‘indntpoff’ see the ELF extension documentation ‘ELF Handling For Thread-Local Storage’.

9.31.3.7 Instruction Marker

The thread local storage instruction markers are used by the linker to perform code optimization.

:tls_load

The `:tls_load` marker is used to flag the load instruction in the initial exec TLS model that retrieves the offset from the thread pointer to a thread local storage variable from the GOT.

:tls_gdcall

The `:tls_gdcall` marker is used to flag the branch-and-save instruction to the `--tls_get_offset` function in the global dynamic TLS model.

:tls_ldcall

The `:tls_ldcall` marker is used to flag the branch-and-save instruction to the `--tls_get_offset` function in the local dynamic TLS model.

For more information about the thread local storage instruction marker and the linker optimizations see the ELF extension documentation ‘[ELF Handling For Thread-Local Storage](#)’.

9.31.3.8 Literal Pool Entries

A literal pool is a collection of values. To access the values a pointer to the literal pool is loaded to a register, the literal pool register. Usually, register `%r13` is used as the literal pool register ([Section 9.31.3.1 \[s390 Register\]](#), [page 216](#)). Literal pool entries are created by adding the suffix `:lit1`, `:lit2`, `:lit4`, or `:lit8` to the end of an expression for an instruction operand. The expression is added to the literal pool and the operand is replaced with the offset to the literal in the literal pool.

:lit1 The literal pool entry is created as an 8-bit value. An operand modifier must not be used for the original expression.

:lit2 The literal pool entry is created as a 16 bit value. The operand modifier `@got` may be used in the original expression. The term ‘`x@got:lit2`’ will put the got offset for the global symbol `x` to the literal pool as 16 bit value.

:lit4 The literal pool entry is created as a 32-bit value. The operand modifier `@got` and `@plt` may be used in the original expression. The term ‘`x@got:lit4`’ will put the got offset for the global symbol `x` to the literal pool as a 32-bit value. The term ‘`x@plt:lit4`’ will put the plt offset for the global symbol `x` to the literal pool as a 32-bit value.

:lit8 The literal pool entry is created as a 64-bit value. The operand modifier `@got` and `@plt` may be used in the original expression. The term ‘`x@got:lit8`’ will put the got offset for the global symbol `x` to the literal pool as a 64-bit value. The term ‘`x@plt:lit8`’ will put the plt offset for the global symbol `x` to the literal pool as a 64-bit value.

The assembler directive ‘`.ltorg`’ is used to emit all literal pool entries to the current position.

9.31.4 Assembler Directives

as for s390 supports all of the standard ELF assembler directives as outlined in the main part of this document. Some directives have been extended and there are some additional directives, which are only available for the s390 as.

.insn This directive permits the numeric representation of an instructions and makes the assembler insert the operands according to one of the instructions formats for ‘.insn’ (Section 9.31.3.4 [s390 Formats], page 219). For example, the instruction ‘l %r1,24(%r15)’ could be written as ‘.insn rx,0x58000000,%r1,24(%r15)’.

.short

.long

.quad

This directive places one or more 16-bit (.short), 32-bit (.long), or 64-bit (.quad) values into the current section. If an ELF or TLS modifier is used only the following expressions are allowed: ‘symbol@modifier + constant’, ‘symbol@modifier + label + constant’, and ‘symbol@modifier - label + constant’. The following modifiers are available:

@got

@got12 The @got modifier can be used for .short, .long and .quad. The @got12 modifier is synonym to @got. The symbol is added to the GOT. The symbol term is replaced with offset from the start of the GOT to the GOT slot for the symbol.

@gotoff The @gotoff modifier can be used for .short, .long and .quad. The symbol term is replaced with the offset from the start of the GOT to the address of the symbol.

@gotplt The @gotplt modifier can be used for .long and .quad. A procedure linkage table entry is generated for the symbol and a jump slot for the symbol is added to the GOT. The symbol term is replaced with the offset from the start of the GOT to the jump slot for the symbol.

@plt The @plt modifier can be used for .long and .quad. A procedure linkage table entry is generated for the symbol. The symbol term is replaced with the address of the PLT entry for the symbol.

@pltoff The @pltoff modifier can be used for .short, .long and .quad. The symbol term is replaced with the offset from the start of the PLT to the address of the symbol.

@tlsld

@tlsldm The @tlsld and @tlsldm modifier can be used for .long and .quad. A tls_index structure for the symbol is added to the GOT. The symbol term is replaced with the offset from the start of the GOT to the tls_index structure.

@gotntpoff
@indntpoff

The @gotntpoff and @indntpoff modifier can be used for .long and .quad. The symbol is added to the static TLS block and the negated offset to the symbol in the static TLS block is added to the GOT. For @gotntpoff the symbol term is replaced with the offset from the start of the GOT to the GOT slot, for @indntpoff the symbol term is replaced with the address of the GOT slot.

@dtpoff The @dtpoff modifier can be used for .long and .quad. The symbol term is replaced with the offset of the symbol relative to the start of the TLS block it is contained in.

@ntpoff The @ntpoff modifier can be used for .long and .quad. The symbol term is replaced with the offset of the symbol relative to the TCB pointer.

For more information about the thread local storage modifiers see the ELF extension documentation ‘[ELF Handling For Thread-Local Storage](#)’.

.ltorg This directive causes the current contents of the literal pool to be dumped to the current location ([Section 9.31.3.8 \[s390 Literal Pool Entries\]](#), [page 227](#)).

.machine string

This directive allows you to change the machine for which code is generated. **string** may be any of the **-march=** selection options (without the **-march=**), **push**, or **pop**. **.machine push** saves the currently selected cpu, which may be restored with **.machine pop**. Be aware that the cpu string has to be put into double quotes in case it contains characters not appropriate for identifiers. So you have to write "z9-109" instead of just z9-109.

9.31.5 Floating Point

The assembler recognizes both the IEEE floating-point instruction and the hexadecimal floating-point instructions. The floating-point constructors ‘.float’, ‘.single’, and ‘.double’ always emit the IEEE format. To assemble hexadecimal floating-point constants the ‘.long’ and ‘.quad’ directives must be used.

9.32 SCORE Dependent Features

9.32.1 Options

The following table lists all available SCORE options.

| | |
|----------------------------|---|
| <code>-G num</code> | This option sets the largest size of an object that can be referenced implicitly with the <code>gp</code> register. The default value is 8. |
| <code>-EB</code> | Assemble code for a big-endian cpu |
| <code>-EL</code> | Assemble code for a little-endian cpu |
| <code>-FIXDD</code> | Assemble code for fix data dependency |
| <code>-NWARN</code> | Assemble code for no warning message for fix data dependency |
| <code>-SCORE5</code> | Assemble code for target is SCORE5 |
| <code>-SCORE5U</code> | Assemble code for target is SCORE5U |
| <code>-SCORE7</code> | Assemble code for target is SCORE7, this is default setting |
| <code>-SCORE3</code> | Assemble code for target is SCORE3 |
| <code>-march=score7</code> | Assemble code for target is SCORE7, this is default setting |
| <code>-march=score3</code> | Assemble code for target is SCORE3 |
| <code>-USE_R1</code> | Assemble code for no warning message when using temp register r1 |
| <code>-KPIC</code> | Generate code for PIC. This option tells the assembler to generate score position-independent macro expansions. It also tells the assembler to mark the output file as PIC. |
| <code>-O0</code> | Assembler will not perform any optimizations |
| <code>-V</code> | Sunplus release version |

9.32.2 SCORE Assembler Directives

A number of assembler directives are available for SCORE. The following table is far from complete.

| | |
|---------------------------|--|
| <code>.set nwarn</code> | Let the assembler not to generate warnings if the source machine language instructions happen data dependency. |
| <code>.set fixdd</code> | Let the assembler to insert bubbles (32 bit nop instruction / 16 bit nop! Instruction) if the source machine language instructions happen data dependency. |
| <code>.set nofixdd</code> | Let the assembler to generate warnings if the source machine language instructions happen data dependency. (Default) |
| <code>.set r1</code> | Let the assembler not to generate warnings if the source program uses r1. allow user to use r1 |

- set nor1** Let the assembler to generate warnings if the source program uses r1. (Default)
- .sdata** Tell the assembler to add subsequent data into the sdata section
- .rdata** Tell the assembler to add subsequent data into the rdata section
- .frame "frame-register", "offset", "return-pc-register"**
Describe a stack frame. "frame-register" is the frame register, "offset" is the distance from the frame register to the virtual frame pointer, "return-pc-register" is the return program register. You must use ".ent" before ".frame" and only one ".frame" can be used per ".ent".
- .mask "bitmask", "frameoffset"**
Indicate which of the integer registers are saved in the current function's stack frame, this is for the debugger to explain the frame chain.
- .ent "proc-name"**
Set the beginning of the procedure "proc_name". Use this directive when you want to generate information for the debugger.
- .end proc-name**
Set the end of a procedure. Use this directive to generate information for the debugger.
- .bss** Switch the destination of following statements into the bss section, which is used for data that is uninitialized anywhere.

9.33 Renesas / SuperH SH Dependent Features

9.33.1 Options

as has following command-line options for the Renesas (formerly Hitachi) / SuperH SH family.

- little** Generate little endian code.
- big** Generate big endian code.
- relax** Alter jump instructions for long displacements.
- small** Align sections to 4 byte boundaries, not 16.
- dsp** Enable sh-dsp insns, and disable sh3e / sh4 insns.
- renesas** Disable optimization with section symbol for compatibility with Renesas assembler.
- allow-reg-prefix** Allow '\$' as a register name prefix.
- fdpic** Generate an FDPIC object file.
- isa=sh4 | sh4a** Specify the sh4 or sh4a instruction set.
- isa=dsp** Enable sh-dsp insns, and disable sh3e / sh4 insns.
- isa=fp** Enable sh2e, sh3e, sh4, and sh4a insn sets.
- isa=all** Enable sh1, sh2, sh2e, sh3, sh3e, sh4, sh4a, and sh-dsp insn sets.
- h-tick-hex** Support H'00 style hex constants in addition to 0x00 style.

9.33.2 Syntax

9.33.2.1 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.

9.33.2.2 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:

- pr** procedure register (holds return address)
- pc** 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) |

9.33.2.3 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.

| | |
|---------------------------|---|
| <code>Rn</code> | Register direct |
| <code>@Rn</code> | Register indirect |
| <code>@-Rn</code> | Register indirect with pre-decrement |
| <code>@Rn+</code> | Register indirect with post-increment |
| <code>@(disp, Rn)</code> | Register indirect with displacement |
| <code>@(R0, Rn)</code> | Register indexed |
| <code>@(disp, GBR)</code> | GBR offset |
| <code>@(R0, GBR)</code> | GBR indexed |
| <code>addr</code> | |
| <code>@(disp, PC)</code> | PC relative address (for branch or for addressing memory). The <code>as</code> implementation allows you to use the simpler form <code>addr</code> anywhere a PC relative address is called for; the alternate form is supported for compatibility with other assemblers. |
| <code>#imm</code> | Immediate data |

9.33.3 Floating Point

SH2E, SH3E and SH4 groups have on-chip floating-point unit (FPU). Other SH groups can use `.float` directive to generate IEEE floating-point numbers.

SH2E and SH3E support single-precision floating point calculations as well as entirely PCAPI compatible emulation of double-precision floating point calculations. SH2E and SH3E instructions are a subset of the floating point calculations conforming to the IEEE754 standard.

In addition to single-precision and double-precision floating-point operation capability, the on-chip FPU of SH4 has a 128-bit graphic engine that enables 32-bit floating-point data to be processed 128 bits at a time. It also supports 4 * 4 array operations and inner product operations. Also, a superscalar architecture is employed that enables simultaneous execution of two instructions (including FPU instructions), providing performance of up to twice that of conventional architectures at the same frequency.

9.33.4 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.

9.33.5 Opcodes

For detailed information on the SH machine instruction set, see *SH-Microcomputer User's Manual* (Renesas) or *SH-4 32-bit CPU Core Architecture* (SuperH) and *SuperH (SH) 64-Bit RISC Series* (SuperH).

`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)
```


9.34 SuperH SH64 Dependent Features

9.34.1 Options

- `-isa=sh4 | sh4a`
Specify the sh4 or sh4a instruction set.
- `-isa=dsp` Enable sh-dsp insns, and disable sh3e / sh4 insns.
- `-isa=fp` Enable sh2e, sh3e, sh4, and sh4a insn sets.
- `-isa=all` Enable sh1, sh2, sh2e, sh3, sh3e, sh4, sh4a, and sh-dsp insn sets.
- `-isa=shmedia | -isa=shcompact`
Specify the default instruction set. **SHmedia** specifies the 32-bit opcodes, and **SHcompact** specifies the 16-bit opcodes compatible with previous SH families. The default depends on the ABI selected; the default for the 64-bit ABI is SHmedia, and the default for the 32-bit ABI is SHcompact. If neither the ABI nor the ISA is specified, the default is 32-bit SHcompact.
Note that the `.mode` pseudo-op is not permitted if the ISA is not specified on the command line.
- `-abi=32 | -abi=64`
Specify the default ABI. If the ISA is specified and the ABI is not, the default ABI depends on the ISA, with SHmedia defaulting to 64-bit and SHcompact defaulting to 32-bit.
Note that the `.abi` pseudo-op is not permitted if the ABI is not specified on the command line. When the ABI is specified on the command line, any `.abi` pseudo-ops in the source must match it.
- `-shcompact-const-crange`
Emit code-range descriptors for constants in SHcompact code sections.
- `-no-mix` Disallow SHmedia code in the same section as constants and SHcompact code.
- `-no-expand`
Do not expand MOVI, PT, PTA or PTB instructions.
- `-expand-pt32`
With `-abi=64`, expand PT, PTA and PTB instructions to 32 bits only.
- `-h-tick-hex`
Support H'00 style hex constants in addition to 0x00 style.

9.34.2 Syntax

9.34.2.1 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.

9.34.2.2 Register Names

You can use the predefined symbols ‘r0’ through ‘r63’ to refer to the SH64 general registers, ‘cr0’ through ‘cr63’ for control registers, ‘tr0’ through ‘tr7’ for target address registers, ‘fr0’ through ‘fr63’ for single-precision floating point registers, ‘dr0’ through ‘dr62’ (even numbered registers only) for double-precision floating point registers, ‘fv0’ through ‘fv60’ (multiples of four only) for single-precision floating point vectors, ‘fp0’ through ‘fp62’ (even numbered registers only) for single-precision floating point pairs, ‘mtrx0’ through ‘mtrx48’ (multiples of 16 only) for 4x4 matrices of single-precision floating point registers, ‘pc’ for the program counter, and ‘fpscr’ for the floating point status and control register.

You can also refer to the control registers by the mnemonics ‘sr’, ‘ssr’, ‘pssr’, ‘intevt’, ‘expevt’, ‘pexpevt’, ‘tra’, ‘spc’, ‘pspc’, ‘resvec’, ‘vbr’, ‘tea’, ‘dcr’, ‘kcr0’, ‘kcr1’, ‘ctc’, and ‘usr’.

9.34.2.3 Addressing Modes

SH64 operands consist of either a register or immediate value. The immediate value can be a constant or label reference (or portion of a label reference), as in this example:

```
movi 4,r2
pt function, tr4
movi (function >> 16) & 65535,r0
shori function & 65535, r0
ld.l r0,4,r0
```

Instruction label references can reference labels in either SHmedia or SHcompact. To differentiate between the two, labels in SHmedia sections will always have the least significant bit set (i.e. they will be odd), which SHcompact labels will have the least significant bit reset (i.e. they will be even). If you need to reference the actual address of a label, you can use the `datalabel` modifier, as in this example:

```
.long function
.long datalabel function
```

In that example, the first longword may or may not have the least significant bit set depending on whether the label is an SHmedia label or an SHcompact label. The second longword will be the actual address of the label, regardless of what type of label it is.

9.34.3 SH64 Machine Directives

In addition to the SH directives, the SH64 provides the following directives:

```
.mode [shmedia|shcompact]
.isa [shmedia|shcompact]
```

Specify the ISA for the following instructions (the two directives are equivalent). Note that programs such as `objdump` rely on symbolic labels to determine when such mode switches occur (by checking the least significant bit of the label’s address), so such mode/isa changes should always be followed by a label (in practice, this is true anyway). Note that you cannot use these directives if you didn’t specify an ISA on the command line.

.abi [32|64]

Specify the ABI for the following instructions. Note that you cannot use this directive unless you specified an ABI on the command line, and the ABIs specified must match.

.uaquad Like **.uaword** and **.ualong**, this allows you to specify an intentionally unaligned quadword (64 bit word).

9.34.4 Opcodes

For detailed information on the SH64 machine instruction set, see *SuperH 64 bit RISC Series Architecture Manual* (SuperH, Inc.).

as implements all the standard SH64 opcodes. In addition, the following pseudo-opcodes may be expanded into one or more alternate opcodes:

movi If the value doesn't fit into a standard **movi** opcode, **as** will replace the **movi** with a sequence of **movi** and **shori** opcodes.

pt This expands to a sequence of **movi** and **shori** opcode, followed by a **ptrel** opcode, or to a **pta** or **ptb** opcode, depending on the label referenced.

9.35 SPARC Dependent Features

9.35.1 Options

The SPARC chip family includes several successive versions, using the same core instruction set, but including a few additional instructions at each version. 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 past 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 SunOS 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.

9.35.2 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 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 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.

9.35.3 Sparc Syntax

The assembler syntax closely follows The Sparc Architecture Manual, versions 8 and 9, as well as most extensions defined by Sun for their UltraSPARC and Niagara line of processors.

9.35.3.1 Special Characters

‘#’ is the line comment character.

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

9.35.3.2 Register Names

The Sparc integer register file is broken down into global, outgoing, local, and incoming.

- The 8 global registers are referred to as ‘%gn’.
- The 8 outgoing registers are referred to as ‘%on’.
- The 8 local registers are referred to as ‘%ln’.
- The 8 incoming registers are referred to as ‘%in’.
- The frame pointer register ‘%i6’ can be referenced using the alias ‘%fp’.
- The stack pointer register ‘%o6’ can be referenced using the alias ‘%sp’.

Floating point registers are simply referred to as ‘%fn’. When assembling for pre-V9, only 32 floating point registers are available. For V9 and later there are 64, but there are restrictions when referencing the upper 32 registers. They can only be accessed as double or quad, and thus only even or quad numbered accesses are allowed. For example, ‘%f34’ is a legal floating point register, but ‘%f35’ is not.

Certain V9 instructions allow access to ancillary state registers. Most simply they can be referred to as ‘%asrn’ where *n* can be from 16 to 31. However, there are some aliases defined to reference ASR registers defined for various UltraSPARC processors:

- The tick compare register is referred to as ‘%tick_cmpr’.
- The system tick register is referred to as ‘%stick’. An alias, ‘%sys_tick’, exists but is deprecated and should not be used by new software.
- The system tick compare register is referred to as ‘%stick_cmpr’. An alias, ‘%sys_tick_cmpr’, exists but is deprecated and should not be used by new software.
- The software interrupt register is referred to as ‘%softint’.
- The set software interrupt register is referred to as ‘%set_softint’. The mnemonic ‘%softint_set’ is provided as an alias.
- The clear software interrupt register is referred to as ‘%clear_softint’. The mnemonic ‘%softint_clear’ is provided as an alias.
- The performance instrumentation counters register is referred to as ‘%pic’.
- The performance control register is referred to as ‘%pcr’.
- The graphics status register is referred to as ‘%gsr’.
- The V9 dispatch control register is referred to as ‘%dcr’.

Various V9 branch and conditional move instructions allow specification of which set of integer condition codes to test. These are referred to as ‘%xcc’ and ‘%icc’.

In V9, there are 4 sets of floating point condition codes which are referred to as ‘%fccn’.

Several special privileged and non-privileged registers exist:

- The V9 address space identifier register is referred to as ‘%asi’.
- The V9 restorable windows register is referred to as ‘%canrestore’.
- The V9 savable windows register is referred to as ‘%cansave’.
- The V9 clean windows register is referred to as ‘%cleanwin’.
- The V9 current window pointer register is referred to as ‘%cwp’.
- The floating-point queue register is referred to as ‘%fq’.
- The V8 co-processor queue register is referred to as ‘%cq’.
- The floating point status register is referred to as ‘%fsr’.
- The other windows register is referred to as ‘%otherwin’.
- The V9 program counter register is referred to as ‘%pc’.
- The V9 next program counter register is referred to as ‘%npc’.
- The V9 processor interrupt level register is referred to as ‘%pil’.
- The V9 processor state register is referred to as ‘%pstate’.
- The trap base address register is referred to as ‘%tba’.
- The V9 tick register is referred to as ‘%tick’.
- The V9 trap level is referred to as ‘%tl’.
- The V9 trap program counter is referred to as ‘%tpc’.
- The V9 trap next program counter is referred to as ‘%tnpc’.
- The V9 trap state is referred to as ‘%tstate’.
- The V9 trap type is referred to as ‘%tt’.
- The V9 condition codes is referred to as ‘%ccr’.
- The V9 floating-point registers state is referred to as ‘%fprs’.
- The V9 version register is referred to as ‘%ver’.
- The V9 window state register is referred to as ‘%wstate’.
- The Y register is referred to as ‘%y’.
- The V8 window invalid mask register is referred to as ‘%wim’.
- The V8 processor state register is referred to as ‘%psr’.
- The V9 global register level register is referred to as ‘%gl’.

Several special register names exist for hypervisor mode code:

- The hyperprivileged processor state register is referred to as ‘%hpstate’.
 - The hyperprivileged trap state register is referred to as ‘%htstate’.
 - The hyperprivileged interrupt pending register is referred to as ‘%hintp’.
 - The hyperprivileged trap base address register is referred to as ‘%htba’.
 - The hyperprivileged implementation version register is referred to as ‘%hver’.
 - The hyperprivileged system tick compare register is referred to as ‘%hstick_cmpr’.
- Note that there is no ‘%hstick’ register, the normal ‘%stick’ is used.

9.35.3.3 Constants

Several Sparc instructions take an immediate operand field for which mnemonic names exist. Two such examples are ‘`membar`’ and ‘`prefetch`’. Another example are the set of V9 memory access instruction that allow specification of an address space identifier.

The ‘`membar`’ instruction specifies a memory barrier that is the defined by the operand which is a bitmask. The supported mask mnemonics are:

- ‘`#Sync`’ requests that all operations (including nonmemory reference operations) appearing prior to the `membar` must have been performed and the effects of any exceptions become visible before any instructions after the `membar` may be initiated. This corresponds to `membar` cmask field bit 2.
- ‘`#MemIssue`’ requests that all memory reference operations appearing prior to the `membar` must have been performed before any memory operation after the `membar` may be initiated. This corresponds to `membar` cmask field bit 1.
- ‘`#Lookaside`’ requests that a store appearing prior to the `membar` must complete before any load following the `membar` referencing the same address can be initiated. This corresponds to `membar` cmask field bit 0.
- ‘`#StoreStore`’ defines that the effects of all stores appearing prior to the `membar` instruction must be visible to all processors before the effect of any stores following the `membar`. Equivalent to the deprecated `stbar` instruction. This corresponds to `membar` mmask field bit 3.
- ‘`#LoadStore`’ defines all loads appearing prior to the `membar` instruction must have been performed before the effect of any stores following the `membar` is visible to any other processor. This corresponds to `membar` mmask field bit 2.
- ‘`#StoreLoad`’ defines that the effects of all stores appearing prior to the `membar` instruction must be visible to all processors before loads following the `membar` may be performed. This corresponds to `membar` mmask field bit 1.
- ‘`#LoadLoad`’ defines that all loads appearing prior to the `membar` instruction must have been performed before any loads following the `membar` may be performed. This corresponds to `membar` mmask field bit 0.

These values can be ored together, for example:

```
membar #Sync
membar #StoreLoad | #LoadLoad
membar #StoreLoad | #StoreStore
```

The `prefetch` and `prefetcha` instructions take a prefetch function code. The following prefetch function code constant mnemonics are available:

- ‘`#n_reads`’ requests a prefetch for several reads, and corresponds to a prefetch function code of 0.
- ‘`#one_read`’ requests a prefetch for one read, and corresponds to a prefetch function code of 1.
- ‘`#n_writes`’ requests a prefetch for several writes (and possibly reads), and corresponds to a prefetch function code of 2.
- ‘`#one_write`’ requests a prefetch for one write, and corresponds to a prefetch function code of 3.

‘#page’ requests a prefetch page, and corresponds to a prefetch function code of 4.
 ‘#invalidate’ requests a prefetch invalidate, and corresponds to a prefetch function code of 16.
 ‘#unified’ requests a prefetch to the nearest unified cache, and corresponds to a prefetch function code of 17.
 ‘#n_reads_strong’ requests a strong prefetch for several reads, and corresponds to a prefetch function code of 20.
 ‘#one_read_strong’ requests a strong prefetch for one read, and corresponds to a prefetch function code of 21.
 ‘#n_writes_strong’ requests a strong prefetch for several writes, and corresponds to a prefetch function code of 22.
 ‘#one_write_strong’ requests a strong prefetch for one write, and corresponds to a prefetch function code of 23.

Only one prefetch code may be specified. Here are some examples:

```
prefetch [%l0 + %l2], #one_read
prefetch [%g2 + 8], #n_writes
prefetcha [%g1] 0x8, #unified
prefetcha [%o0 + 0x10] %asi, #n_reads
```

The actual behavior of a given prefetch function code is processor specific. If a processor does not implement a given prefetch function code, it will treat the prefetch instruction as a nop.

For instructions that accept an immediate address space identifier, **as** provides many mnemonics corresponding to V9 defined as well as UltraSPARC and Niagara extended values. For example, ‘#ASI_P’ and ‘#ASI_BLK_INIT_QUAD_LDD_AIUS’. See the V9 and processor specific manuals for details.

9.35.3.4 Relocations

ELF relocations are available as defined in the 32-bit and 64-bit Sparc ELF specifications.

R_SPARC_HI22 is obtained using ‘%hi’ and R_SPARC_LO10 is obtained using ‘%lo’. Likewise R_SPARC_HIX22 is obtained from ‘%hix’ and R_SPARC_LOX10 is obtained using ‘%lox’. For example:

```
sethi %hi(symbol), %g1
or    %g1, %lo(symbol), %g1

sethi %hix(symbol), %g1
xor   %g1, %lox(symbol), %g1
```

These “high” mnemonics extract bits 31:10 of their operand, and the “low” mnemonics extract bits 9:0 of their operand.

V9 code model relocations can be requested as follows:

- R_SPARC_HH22 is requested using ‘%hh’. It can also be generated using ‘%uhi’.
- R_SPARC_HM10 is requested using ‘%hm’. It can also be generated using ‘%ulo’.
- R_SPARC_LM22 is requested using ‘%lm’.
- R_SPARC_H44 is requested using ‘%h44’.

- R_SPARC_M44 is requested using ‘%m44’.
- R_SPARC_L44 is requested using ‘%l44’.

The PC relative relocation R_SPARC_PC22 can be obtained by enclosing an operand inside of ‘%pc22’. Likewise, the R_SPARC_PC10 relocation can be obtained using ‘%pc10’. These are mostly used when assembling PIC code. For example, the standard PIC sequence on Sparc to get the base of the global offset table, PC relative, into a register, can be performed as:

```
sethi %pc22(_GLOBAL_OFFSET_TABLE_-4), %l7
add   %l7, %pc10(_GLOBAL_OFFSET_TABLE_+4), %l7
```

Several relocations exist to allow the link editor to potentially optimize GOT data references. The R_SPARC_GOTDATA_OP_HIX22 relocation can be obtained by enclosing an operand inside of ‘%gdop_hix22’. The R_SPARC_GOTDATA_OP_LOX10 relocation can be obtained by enclosing an operand inside of ‘%gdop_lox10’. Likewise, R_SPARC_GOTDATA_OP can be obtained by enclosing an operand inside of ‘%gdop’. For example, assuming the GOT base is in register %l7:

```
sethi %gdop_hix22(symbol), %l1
xor   %l1, %gdop_lox10(symbol), %l1
ld    [%l7 + %l1], %l2, %gdop(symbol)
```

There are many relocations that can be requested for access to thread local storage variables. All of the Sparc TLS mnemonics are supported:

- R_SPARC_TLS_GD_HI22 is requested using ‘%tgd_hi22’.
- R_SPARC_TLS_GD_LO10 is requested using ‘%tgd_lo10’.
- R_SPARC_TLS_GD_ADD is requested using ‘%tgd_add’.
- R_SPARC_TLS_GD_CALL is requested using ‘%tgd_call’.
- R_SPARC_TLS_LDM_HI22 is requested using ‘%tldm_hi22’.
- R_SPARC_TLS_LDM_LO10 is requested using ‘%tldm_lo10’.
- R_SPARC_TLS_LDM_ADD is requested using ‘%tldm_add’.
- R_SPARC_TLS_LDM_CALL is requested using ‘%tldm_call’.
- R_SPARC_TLS_LDO_HIX22 is requested using ‘%tlldo_hix22’.
- R_SPARC_TLS_LDO_LOX10 is requested using ‘%tlldo_lox10’.
- R_SPARC_TLS_LDO_ADD is requested using ‘%tlldo_add’.
- R_SPARC_TLS_IE_HI22 is requested using ‘%tie_hi22’.
- R_SPARC_TLS_IE_LO10 is requested using ‘%tie_lo10’.
- R_SPARC_TLS_IE_LD is requested using ‘%tie_ld’.
- R_SPARC_TLS_IE_LDX is requested using ‘%tie_ldx’.
- R_SPARC_TLS_IE_ADD is requested using ‘%tie_add’.
- R_SPARC_TLS_LE_HIX22 is requested using ‘%tle_hix22’.
- R_SPARC_TLS_LE_LOX10 is requested using ‘%tle_lox10’.

Here are some example TLS model sequences.

First, General Dynamic:

```

sethi %tgd_hi22(symbol), %l1
add   %l1, %tgd_lo10(symbol), %l1
add   %l7, %l1, %o0, %tgd_add(symbol)
call  __tls_get_addr, %tgd_call(symbol)
nop

```

Local Dynamic:

```

sethi %tldm_hi22(symbol), %l1
add   %l1, %tldm_lo10(symbol), %l1
add   %l7, %l1, %o0, %tldm_add(symbol)
call  __tls_get_addr, %tldm_call(symbol)
nop

```

```

sethi %tldo_hix22(symbol), %l1
xor   %l1, %tldo_lox10(symbol), %l1
add   %o0, %l1, %l1, %tldo_add(symbol)

```

Initial Exec:

```

sethi %tie_hi22(symbol), %l1
add   %l1, %tie_lo10(symbol), %l1
ld     [%l7 + %l1], %o0, %tie_ld(symbol)
add   %g7, %o0, %o0, %tie_add(symbol)

sethi %tie_hi22(symbol), %l1
add   %l1, %tie_lo10(symbol), %l1
ldx   [%l7 + %l1], %o0, %tie_ldx(symbol)
add   %g7, %o0, %o0, %tie_add(symbol)

```

And finally, Local Exec:

```

sethi %tle_hix22(symbol), %l1
add   %l1, %tle_lox10(symbol), %l1
add   %g7, %l1, %l1

```

When assembling for 64-bit, and a secondary constant addend is specified in an address expression that would normally generate an R_SPARC_L010 relocation, the assembler will emit an R_SPARC_OL010 instead.

9.35.3.5 Size Translations

Often it is desirable to write code in an operand size agnostic manner. **as** provides support for this via operand size opcode translations. Translations are supported for loads, stores, shifts, compare-and-swap atomics, and the ‘clr’ synthetic instruction.

If generating 32-bit code, **as** will generate the 32-bit opcode. Whereas if 64-bit code is being generated, the 64-bit opcode will be emitted. For example **ldn** will be transformed into **ld** for 32-bit code and **ldx** for 64-bit code.

Here is an example meant to demonstrate all the supported opcode translations:

```

ldn    [%o0], %o1
ldna   [%o0] %asi, %o2
stn    %o1, [%o0]

```

```

stna  %o2, [%o0] %asi
slln  %o3, 3, %o3
srln  %o4, 8, %o4
sran  %o5, 12, %o5
casn  [%o0], %o1, %o2
casna [%o0] %asi, %o1, %o2
clrn  %g1

```

In 32-bit mode as will emit:

```

ld    [%o0], %o1
lda   [%o0] %asi, %o2
st    %o1, [%o0]
sta   %o2, [%o0] %asi
sll   %o3, 3, %o3
srl   %o4, 8, %o4
sra   %o5, 12, %o5
cas   [%o0], %o1, %o2
casa  [%o0] %asi, %o1, %o2
clr   %g1

```

And in 64-bit mode as will emit:

```

ldx   [%o0], %o1
ldxa  [%o0] %asi, %o2
stx   %o1, [%o0]
stxa  %o2, [%o0] %asi
sllx  %o3, 3, %o3
srlx  %o4, 8, %o4
srax  %o5, 12, %o5
casx  [%o0], %o1, %o2
casxa [%o0] %asi, %o1, %o2
clrx  %g1

```

Finally, the `.nword` translating directive is supported as well. It is documented in the section on Sparc machine directives.

9.35.4 Floating Point

The Sparc uses IEEE floating-point numbers.

9.35.5 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`.

| | |
|------------------------|---|
| <code>.proc</code> | This directive is ignored. Any text following it on the same line is also ignored. |
| <code>.register</code> | This directive declares use of a global application or system register. It must be followed by a register name <code>%g2</code> , <code>%g3</code> , <code>%g6</code> or <code>%g7</code> , comma and the symbol name for that register. If symbol name is <code>#scratch</code> , it is a scratch register, if it is <code>#ignore</code> , 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. |
| <code>.reserve</code> | This must be followed by a symbol name, a positive number, and <code>"bss"</code> . This behaves somewhat like <code>.lcomm</code> , but the syntax is different. |
| <code>.seg</code> | This must be followed by <code>"text"</code> , <code>"data"</code> , or <code>"data1"</code> . It behaves like <code>.text</code> , <code>.data</code> , or <code>.data 1</code> . |
| <code>.skip</code> | This is functionally identical to the <code>.space</code> directive. |
| <code>.word</code> | On the Sparc, the <code>.word</code> directive produces 32 bit values, instead of the 16 bit values it produces on many other machines. |
| <code>.xword</code> | On the Sparc V9 processor, the <code>.xword</code> directive produces 64 bit values. |

9.36 TIC54X Dependent Features

9.36.1 Options

The TMS320C54X version of **as** has a few machine-dependent options.

You can use the `‘-mfar-mode’` option to enable extended addressing mode. All addresses will be assumed to be > 16 bits, and the appropriate relocation types will be used. This option is equivalent to using the `‘.far_mode’` directive in the assembly code. If you do not use the `‘-mfar-mode’` option, all references will be assumed to be 16 bits. This option may be abbreviated to `‘-mf’`.

You can use the `‘-mcpu’` option to specify a particular CPU. This option is equivalent to using the `‘.version’` directive in the assembly code. For recognized CPU codes, see See [Section 9.36.9 \[.version\], page 250](#). The default CPU version is `‘542’`.

You can use the `‘-merrors-to-file’` option to redirect error output to a file (this provided for those deficient environments which don’t provide adequate output redirection). This option may be abbreviated to `‘-me’`.

9.36.2 Blocking

A blocked section or memory block is guaranteed not to cross the blocking boundary (usually a page, or 128 words) if it is smaller than the blocking size, or to start on a page boundary if it is larger than the blocking size.

9.36.3 Environment Settings

`‘C54XDSP_DIR’` and `‘A_DIR’` are semicolon-separated paths which are added to the list of directories normally searched for source and include files. `‘C54XDSP_DIR’` will override `‘A_DIR’`.

9.36.4 Constants Syntax

The TIC54X version of **as** allows the following additional constant formats, using a suffix to indicate the radix:

| | |
|-------------|------------------|
| Binary | 000000B, 011000b |
| Octal | 10Q, 224q |
| Hexadecimal | 45h, 0FH |

9.36.5 String Substitution

A subset of allowable symbols (which we’ll call subsyms) may be assigned arbitrary string values. This is roughly equivalent to C preprocessor `#define` macros. When **as** encounters one of these symbols, the symbol is replaced in the input stream by its string value. Subsym names **must** begin with a letter.

Subsyms may be defined using the `.asg` and `.eval` directives (See [Section 9.36.9 \[.asg\], page 250](#), See [Section 9.36.9 \[.eval\], page 250](#)).

Expansion is recursive until a previously encountered symbol is seen, at which point substitution stops.

In this example, `x` is replaced with `SYM2`; `SYM2` is replaced with `SYM1`, and `SYM1` is replaced with `x`. At this point, `x` has already been encountered and the substitution stops.

```
.asg    "x",SYM1
.asg    "SYM1",SYM2
.asg    "SYM2",x
add     x,a           ; final code assembled is "add x, a"
```

Macro parameters are converted to subsyms; a side effect of this is the normal `as '\ARG'` dereferencing syntax is unnecessary. Subsyzms defined within a macro will have global scope, unless the `.var` directive is used to identify the subsym as a local macro variable see [Section 9.36.9 \[.var\], page 250](#).

Substitution may be forced in situations where replacement might be ambiguous by placing colons on either side of the subsym. The following code:

```
.eval  "10",x
LAB:X: add    #x, a
```

When assembled becomes:

```
LAB10 add    #10, a
```

Smaller parts of the string assigned to a subsym may be accessed with the following syntax:

`:symbol(char_index):`

Evaluates to a single-character string, the character at *char_index*.

`:symbol(start,length):`

Evaluates to a substring of *symbol* beginning at *start* with length *length*.

9.36.6 Local Labels

Local labels may be defined in two ways:

- `$N`, where `N` is a decimal number between 0 and 9
- `LABEL?`, where `LABEL` is any legal symbol name.

Local labels thus defined may be redefined or automatically generated. The scope of a local label is based on when it may be undefined or reset. This happens when one of the following situations is encountered:

- `.newblock` directive see [Section 9.36.9 \[.newblock\], page 250](#)
- The current section is changed (`.sect`, `.text`, or `.data`)
- Entering or leaving an included file
- The macro scope where the label was defined is exited

9.36.7 Math Builtins

The following built-in functions may be used to generate a floating-point value. All return a floating-point value except `$cvi`, `$int`, and `$sgn`, which return an integer value.

`$acos(expr)`

Returns the floating point arccosine of *expr*.

- `$asin(expr)`
Returns the floating point arcsine of *expr*.
- `$atan(expr)`
Returns the floating point arctangent of *expr*.
- `$atan2(expr1,expr2)`
Returns the floating point arctangent of *expr1* / *expr2*.
- `$ceil(expr)`
Returns the smallest integer not less than *expr* as floating point.
- `$cosh(expr)`
Returns the floating point hyperbolic cosine of *expr*.
- `$cos(expr)`
Returns the floating point cosine of *expr*.
- `$cvf(expr)`
Returns the integer value *expr* converted to floating-point.
- `$cvi(expr)`
Returns the floating point value *expr* converted to integer.
- `$exp(expr)`
Returns the floating point value e^{expr} .
- `$fabs(expr)`
Returns the floating point absolute value of *expr*.
- `$floor(expr)`
Returns the largest integer that is not greater than *expr* as floating point.
- `$fmod(expr1,expr2)`
Returns the floating point remainder of *expr1* / *expr2*.
- `$int(expr)`
Returns 1 if *expr* evaluates to an integer, zero otherwise.
- `$ldexp(expr1,expr2)`
Returns the floating point value $\text{expr1} * 2^{\text{expr2}}$.
- `$log10(expr)`
Returns the base 10 logarithm of *expr*.
- `$log(expr)`
Returns the natural logarithm of *expr*.
- `$max(expr1,expr2)`
Returns the floating point maximum of *expr1* and *expr2*.
- `$min(expr1,expr2)`
Returns the floating point minimum of *expr1* and *expr2*.
- `$pow(expr1,expr2)`
Returns the floating point value $\text{expr1}^{\text{expr2}}$.

\$round(*expr*)
Returns the nearest integer to *expr* as a floating point number.

\$sgn(*expr*)
Returns -1, 0, or 1 based on the sign of *expr*.

\$sin(*expr*)
Returns the floating point sine of *expr*.

\$sinh(*expr*)
Returns the floating point hyperbolic sine of *expr*.

\$sqrt(*expr*)
Returns the floating point square root of *expr*.

\$tan(*expr*)
Returns the floating point tangent of *expr*.

\$tanh(*expr*)
Returns the floating point hyperbolic tangent of *expr*.

\$trunc(*expr*)
Returns the integer value of *expr* truncated towards zero as floating point.

9.36.8 Extended Addressing

The LDX pseudo-op is provided for loading the extended addressing bits of a label or address. For example, if an address `_label` resides in extended program memory, the value of `_label` may be loaded as follows:

```
ldx    #_label,16,a    ; loads extended bits of _label
or     #_label,a       ; loads lower 16 bits of _label
bacc   a               ; full address is in accumulator A
```

9.36.9 Directives

.align [*size*]
.even Align the section program counter on the next boundary, based on *size*. *size* may be any power of 2. **.even** is equivalent to **.align** with a *size* of 2.

1 Align SPC to word boundary

2 Align SPC to longword boundary (same as **.even**)

128 Align SPC to page boundary

.asg *string*, *name*
Assign *name* the string *string*. String replacement is performed on *string* before assignment.

.eval *string*, *name*
Evaluate the contents of string *string* and assign the result as a string to the subsym *name*. String replacement is performed on *string* before assignment.

`.bss symbol, size [, [blocking_flag] [,alignment_flag]]`

Reserve space for *symbol* in the .bss section. *size* is in words. If present, *blocking_flag* indicates the allocated space should be aligned on a page boundary if it would otherwise cross a page boundary. If present, *alignment_flag* causes the assembler to allocate *size* on a long word boundary.

`.byte value [, ...,value_n]`

`.ubyte value [, ...,value_n]`

`.char value [, ...,value_n]`

`.uchar value [, ...,value_n]`

Place one or more bytes into consecutive words of the current section. The upper 8 bits of each word is zero-filled. If a label is used, it points to the word allocated for the first byte encountered.

`.clink ["section_name"]`

Set STYP_CLINK flag for this section, which indicates to the linker that if no symbols from this section are referenced, the section should not be included in the link. If *section_name* is omitted, the current section is used.

`.c_mode` TBD.

`.copy "filename" | filename`

`.include "filename" | filename`

Read source statements from *filename*. The normal include search path is used. Normally .copy will cause statements from the included file to be printed in the assembly listing and .include will not, but this distinction is not currently implemented.

`.data` Begin assembling code into the .data section.

`.double value [, ...,value_n]`

`.ldouble value [, ...,value_n]`

`.float value [, ...,value_n]`

`.xfloat value [, ...,value_n]`

Place an IEEE single-precision floating-point representation of one or more floating-point values into the current section. All but .xfloat align the result on a longword boundary. Values are stored most-significant word first.

`.drlist`

`.drnolist`

Control printing of directives to the listing file. Ignored.

`.emsg string`

`.mmsg string`

`.wmsg string`

Emit a user-defined error, message, or warning, respectively.

`.far_mode`

Use extended addressing when assembling statements. This should appear only once per file, and is equivalent to the -mfarmode option see [Section 9.36.1 \[-mfarmode\]](#), page 247.

`.fclist`
`.fcnolist`

Control printing of false conditional blocks to the listing file.

`.field value [,size]`

Initialize a bitfield of *size* bits in the current section. If *value* is relocatable, then *size* must be 16. *size* defaults to 16 bits. If *value* does not fit into *size* bits, the value will be truncated. Successive `.field` directives will pack starting at the current word, filling the most significant bits first, and aligning to the start of the next word if the field size does not fit into the space remaining in the current word. A `.align` directive with an operand of 1 will force the next `.field` directive to begin packing into a new word. If a label is used, it points to the word that contains the specified field.

`.global symbol [...,symbol_n]`

`.def symbol [...,symbol_n]`

`.ref symbol [...,symbol_n]`

`.def` nominally identifies a symbol defined in the current file and available to other files. `.ref` identifies a symbol used in the current file but defined elsewhere. Both map to the standard `.global` directive.

`.half value [...,value_n]`

`.uhalf value [...,value_n]`

`.short value [...,value_n]`

`.ushort value [...,value_n]`

`.int value [...,value_n]`

`.uint value [...,value_n]`

`.word value [...,value_n]`

`.uword value [...,value_n]`

Place one or more values into consecutive words of the current section. If a label is used, it points to the word allocated for the first value encountered.

`.label symbol`

Define a special *symbol* to refer to the load time address of the current section program counter.

`.length`

`.width` Set the page length and width of the output listing file. Ignored.

`.list`

`.nolist` Control whether the source listing is printed. Ignored.

`.long value [...,value_n]`

`.ulong value [...,value_n]`

`.xlong value [...,value_n]`

Place one or more 32-bit values into consecutive words in the current section. The most significant word is stored first. `.long` and `.ulong` align the result on a longword boundary; `xlong` does not.

`.loop` [*count*]
`.break` [*condition*]
`.endloop` Repeatedly assemble a block of code. `.loop` begins the block, and `.endloop` marks its termination. *count* defaults to 1024, and indicates the number of times the block should be repeated. `.break` terminates the loop so that assembly begins after the `.endloop` directive. The optional *condition* will cause the loop to terminate only if it evaluates to zero.

macro_name `.macro` [*param1*][, ...*param_n*]
`[.mexit]`
`.endm` See the section on macros for more explanation (See [Section 9.36.10 \[TIC54X-Macros\]](#), page 255).

`.mlib` "*filename*" | *filename*
 Load the macro library *filename*. *filename* must be an archived library (BFD ar-compatible) of text files, expected to contain only macro definitions. The standard include search path is used.

`.mlist`
`.mnolist` Control whether to include macro and loop block expansions in the listing output. Ignored.

`.mmregs` Define global symbolic names for the 'c54x registers. Supposedly equivalent to executing `.set` directives for each register with its memory-mapped value, but in reality is provided only for compatibility and does nothing.

`.newblock`
 This directive resets any TIC54X local labels currently defined. Normal `as` local labels are unaffected.

`.option` *option_list*
 Set listing options. Ignored.

`.sblock` "*section_name*" | *section_name* [, "*name_n*" | *name_n*]
 Designate *section_name* for blocking. Blocking guarantees that a section will start on a page boundary (128 words) if it would otherwise cross a page boundary. Only initialized sections may be designated with this directive. See also See [Section 9.36.2 \[TIC54X-Block\]](#), page 247.

`.sect` "*section_name*"
 Define a named initialized section and make it the current section.

symbol `.set` "*value*"
symbol `.equ` "*value*"
 Equate a constant *value* to a *symbol*, which is placed in the symbol table. *symbol* may not be previously defined.

`.space` *size_in_bits*
`.bes` *size_in_bits*
 Reserve the given number of bits in the current section and zero-fill them. If a label is used with `.space`, it points to the **first** word reserved. With `.bes`, the label points to the **last** word reserved.

`.sslist`
`.ssnolist`

Controls the inclusion of subsym replacement in the listing output. Ignored.

`.string "string" [..., "string_n"]`
`.pstring "string" [..., "string_n"]`

Place 8-bit characters from *string* into the current section. `.string` zero-fills the upper 8 bits of each word, while `.pstring` puts two characters into each word, filling the most-significant bits first. Unused space is zero-filled. If a label is used, it points to the first word initialized.

```
[stag] .struct [offset]
[name_1] element [count_1]
[name_2] element [count_2]
[tname] .tag stagx [tcount]
...
[name_n] element [count_n]
[ssize] .endstruct
label .tag [stag]
```

Assign symbolic offsets to the elements of a structure. *stag* defines a symbol to use to reference the structure. *offset* indicates a starting value to use for the first element encountered; otherwise it defaults to zero. Each element can have a named offset, *name*, which is a symbol assigned the value of the element's offset into the structure. If *stag* is missing, these become global symbols. *count* adjusts the offset that many times, as if *element* were an array. *element* may be one of `.byte`, `.word`, `.long`, `.float`, or any equivalent of those, and the structure offset is adjusted accordingly. `.field` and `.string` are also allowed; the size of `.field` is one bit, and `.string` is considered to be one word in size. Only element descriptors, structure/union tags, `.align` and conditional assembly directives are allowed within `.struct/.endstruct`. `.align` aligns member offsets to word boundaries only. *ssize*, if provided, will always be assigned the size of the structure.

The `.tag` directive, in addition to being used to define a structure/union element within a structure, may be used to apply a structure to a symbol. Once applied to *label*, the individual structure elements may be applied to *label* to produce the desired offsets using *label* as the structure base.

`.tab` Set the tab size in the output listing. Ignored.

```
[utag] .union
[name_1] element [count_1]
[name_2] element [count_2]
[tname] .tag utagx[,tcount]
...
[name_n] element [count_n]
[usize] .endstruct
label .tag [utag]
```

Similar to `.struct`, but the offset after each element is reset to zero, and the `usize` is set to the maximum of all defined elements. Starting offset for the union is always zero.

```
[symbol] .usect "section_name", size, [, [blocking_flag] [, alignment_flag]]
```

Reserve space for variables in a named, uninitialized section (similar to `.bss`). `.usect` allows definitions sections independent of `.bss`. `symbol` points to the first location reserved by this allocation. The symbol may be used as a variable name. `size` is the allocated size in words. `blocking_flag` indicates whether to block this section on a page boundary (128 words) (see [Section 9.36.2 \[TIC54X-Block\]](#), page 247). `alignment_flag` indicates whether the section should be longword-aligned.

```
.var sym[, ..., sym_n]
```

Define a subsym to be a local variable within a macro. See [Section 9.36.10 \[TIC54X-Macros\]](#), page 255.

```
.version version
```

Set which processor to build instructions for. Though the following values are accepted, the op is ignored.

```
541
542
543
545
545LP
546LP
548
549
```

9.36.10 Macros

Macros do not require explicit dereferencing of arguments (i.e., `\ARG`).

During macro expansion, the macro parameters are converted to subsyms. If the number of arguments passed the macro invocation exceeds the number of parameters defined, the last parameter is assigned the string equivalent of all remaining arguments. If fewer arguments are given than parameters, the missing parameters are assigned empty strings. To include a comma in an argument, you must enclose the argument in quotes.

The following built-in subsym functions allow examination of the string value of subsyms (or ordinary strings). The arguments are strings unless otherwise indicated (subsyms passed as args will be replaced by the strings they represent).

- `$symlen(str)`
Returns the length of *str*.
- `$symcmp(str1, str2)`
Returns 0 if *str1* == *str2*, non-zero otherwise.
- `$firstch(str, ch)`
Returns index of the first occurrence of character constant *ch* in *str*.
- `$lastch(str, ch)`
Returns index of the last occurrence of character constant *ch* in *str*.
- `$isdefed(symbol)`
Returns zero if the symbol *symbol* is not in the symbol table, non-zero otherwise.
- `$ismember(symbol, list)`
Assign the first member of comma-separated string *list* to *symbol*; *list* is re-assigned the remainder of the list. Returns zero if *list* is a null string. Both arguments must be subsyms.
- `$iscons(expr)`
Returns 1 if string *expr* is binary, 2 if octal, 3 if hexadecimal, 4 if a character, 5 if decimal, and zero if not an integer.
- `$isname(name)`
Returns 1 if *name* is a valid symbol name, zero otherwise.
- `$isreg(reg)`
Returns 1 if *reg* is a valid predefined register name (AR0-AR7 only).
- `$structsz(stag)`
Returns the size of the structure or union represented by *stag*.
- `$structacc(stag)`
Returns the reference point of the structure or union represented by *stag*. Always returns zero.

9.36.11 Memory-mapped Registers

The following symbols are recognized as memory-mapped registers:

9.37 TIC6X Dependent Features

9.37.1 TIC6X Options

`-march=arch`

Enable (only) instructions from architecture *arch*. By default, all instructions are permitted.

The following values of *arch* are accepted: `c62x`, `c64x`, `c64x+`, `c67x`, `c67x+`, `c674x`.

`-matomic`

`-mno-atomic`

Enable or disable the optional C64x+ atomic operation instructions. By default, they are enabled if no `'-march'` option is given, or if an architecture is specified with `'-march'` that implies these instructions are present (currently, there are no such architectures); they are disabled if an architecture is specified with `'-march'` on which the instructions are optional or not present. This option overrides such a default from the architecture, independent of the order in which the `'-march'` or `'-matomic'` or `'-mno-atomic'` options are passed.

`-mdsbt`

`-mno-dsbt`

The `'-mdsbt'` option causes the assembler to generate the `Tag_ABI_DSBT` attribute with a value of 1, indicating that the code is using DSBT addressing. The `'-mno-dsbt'` option, the default, causes the tag to have a value of 0, indicating that the code does not use DSBT addressing. The linker will emit a warning if objects of different type (DSBT and non-DSBT) are linked together.

`-mpid=no`

`-mpid=near`

`-mpid=far`

The `'-mpid='` option causes the assembler to generate the `Tag_ABI_PID` attribute with a value indicating the form of data addressing used by the code. `'-mpid=no'`, the default, indicates position-dependent data addressing, `'-mpid=near'` indicates position-independent addressing with GOT accesses using near DP addressing, and `'-mpid=far'` indicates position-independent addressing with GOT accesses using far DP addressing. The linker will emit a warning if objects built with different settings of this option are linked together.

`-mpic`

`-mno-pic`

The `'-mpic'` option causes the assembler to generate the `Tag_ABI_PIC` attribute with a value of 1, indicating that the code is using position-independent code addressing. The `-mno-pic` option, the default, causes the tag to have a value of 0, indicating position-dependent code addressing. The linker will emit a warning if objects of different type (position-dependent and position-independent) are linked together.

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

Generate code for the specified endianness. The default is little-endian.

9.37.2 TIC6X Syntax

The presence of a `;` on a line indicates the start of a comment that extends to the end of the current line. If a `#` or `*` appears as the first character of a line, the whole line is treated as a comment.

The `@` character can be used instead of a newline to separate statements.

Instruction, register and functional unit names are case-insensitive. `as` requires fully-specified functional unit names, such as `.S1`, `.L1X` or `.D1T2`, on all instructions using a functional unit.

For some instructions, there may be syntactic ambiguity between register or functional unit names and the names of labels or other symbols. To avoid this, enclose the ambiguous symbol name in parentheses; register and functional unit names may not be enclosed in parentheses.

9.37.3 TIC6X Directives

Directives controlling the set of instructions accepted by the assembler have effect for instructions between the directive and any subsequent directive overriding it.

`.arch arch`

This has the same effect as `-march=arch`.

`.atomic`

`.noatomic`

These have the same effects as `-matomic` and `-mno-atomic`.

`.c6xabi_attribute tag, value`

Set the C6000 EABI build attribute *tag* to *value*.

The *tag* is either an attribute number or one of `Tag_ISA`, `Tag_ABI_wchar_t`, `Tag_ABI_stack_align_needed`, `Tag_ABI_stack_align_preserved`, `Tag_ABI_DSBT`, `Tag_ABI_PID`, `Tag_ABI_PIC`, `Tag_ABI_array_object_alignment`, `Tag_ABI_array_object_align_expected`, `Tag_ABI_compatibility` and `Tag_ABI_conformance`. The *value* is either a number, "string", or number, "string" depending on the tag.

`.nocmp`

Disallow use of C64x+ compact instructions in the current text section.

9.38 Z80 Dependent Features

9.38.1 Options

The Zilog Z80 and Ascii R800 version of **as** have a few machine dependent options.

- `'-z80'` Produce code for the Z80 processor. There are additional options to request warnings and error messages for undocumented instructions.
- `'-ignore-undocumented-instructions'`
- `'-Wnud'` Silently assemble undocumented Z80-instructions that have been adopted as documented R800-instructions.
- `'-ignore-unportable-instructions'`
- `'-Wnup'` Silently assemble all undocumented Z80-instructions.
- `'-warn-undocumented-instructions'`
- `'-Wud'` Issue warnings for undocumented Z80-instructions that work on R800, do not assemble other undocumented instructions without warning.
- `'-warn-unportable-instructions'`
- `'-Wup'` Issue warnings for other undocumented Z80-instructions, do not treat any undocumented instructions as errors.
- `'-forbid-undocumented-instructions'`
- `'-Fud'` Treat all undocumented z80-instructions as errors.
- `'-forbid-unportable-instructions'`
- `'-Fup'` Treat undocumented z80-instructions that do not work on R800 as errors.
- `'-r800'` Produce code for the R800 processor. The assembler does not support undocumented instructions for the R800. In line with common practice, **as** uses Z80 instruction names for the R800 processor, as far as they exist.

9.38.2 Syntax

The assembler syntax closely follows the 'Z80 family CPU User Manual' by Zilog. In expressions a single '=' may be used as "is equal to" comparison operator.

Suffices can be used to indicate the radix of integer constants; 'H' or 'h' for hexadecimal, 'D' or 'd' for decimal, 'Q', 'O', 'q' or 'o' for octal, and 'B' for binary.

The suffix 'b' denotes a backreference to local label.

9.38.2.1 Special Characters

The semicolon ';' is the line comment character;

The dollar sign '\$' can be used as a prefix for hexadecimal numbers and as a symbol denoting the current location counter.

A backslash '\' is an ordinary character for the Z80 assembler.

The single quote "'" must be followed by a closing quote. If there is one character in between, it is a character constant, otherwise it is a string constant.

9.38.2.2 Register Names

The registers are referred to with the letters assigned to them by Zilog. In addition `as` recognizes ‘`ixl`’ and ‘`ixh`’ as the least and most significant octet in ‘`ix`’, and similarly ‘`iy1`’ and ‘`iyh`’ as parts of ‘`iy`’.

9.38.2.3 Case Sensitivity

Upper and lower case are equivalent in register names, opcodes, condition codes and assembler directives. The case of letters is significant in labels and symbol names. The case is also important to distinguish the suffix ‘`b`’ for a backward reference to a local label from the suffix ‘`B`’ for a number in binary notation.

9.38.3 Floating Point

Floating-point numbers are not supported.

9.38.4 Z80 Assembler Directives

`as` for the Z80 supports some additional directives for compatibility with other assemblers.

These are the additional directives in `as` for the Z80:

`db expression | string [, expression | string ...]`

`defb expression | string [, expression | string ...]`

For each *string* the characters are copied to the object file, for each other *expression* the value is stored in one byte. A warning is issued in case of an overflow.

`dw expression [, expression ...]`

`defw expression [, expression ...]`

For each *expression* the value is stored in two bytes, ignoring overflow.

`d24 expression [, expression ...]`

`def24 expression [, expression ...]`

For each *expression* the value is stored in three bytes, ignoring overflow.

`d32 expression [, expression ...]`

`def32 expression [, expression ...]`

For each *expression* the value is stored in four bytes, ignoring overflow.

`ds count [, value]`

`defs count [, value]`

Fill *count* bytes in the object file with *value*, if *value* is omitted it defaults to zero.

`symbol equ expression`

`symbol defl expression`

These directives set the value of *symbol* to *expression*. If `equ` is used, it is an error if *symbol* is already defined. Symbols defined with `equ` are not protected from redefinition.

`set` This is a normal instruction on Z80, and not an assembler directive.

`psect name`

A synonym for See [Section 7.96 \[Section\], page 68](#), no second argument should be given.

9.38.5 Opcodes

In line with common practice, Z80 mnemonics are used for both the Z80 and the R800.

In many instructions it is possible to use one of the half index registers ('ixl', 'ixh', 'iy1', 'iyh') in stead of an 8-bit general purpose register. This yields instructions that are documented on the R800 and undocumented on the Z80. Similarly `in f, (c)` is documented on the R800 and undocumented on the Z80.

The assembler also supports the following undocumented Z80-instructions, that have not been adopted in the R800 instruction set:

`out (c), 0` Sends zero to the port pointed to by register `c`.

`sli m` Equivalent to $m = (m \ll 1) + 1$, the operand `m` can be any operand that is valid for 'sla'. One can use 'sll' as a synonym for 'sli'.

`op (ix+d), r`

This is equivalent to

```
ld r, (ix+d)
opc r
ld (ix+d), r
```

The operation 'opc' may be any of 'res b,', 'set b,', 'rl', 'rlc', 'rr', 'rrc', 'sla', 'sli', 'sra' and 'srl', and the register 'r' may be any of 'a', 'b', 'c', 'd', 'e', 'h' and 'l'.

`opc (iy+d), r`

As above, but with 'iy' instead of 'ix'.

The web site at <http://www.z80.info> is a good starting place to find more information on programming the Z80.

9.39 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 9.39.3 \[Assembler Directives for the Z8000\], page 263](#), for a list of other Z8000 specific assembler directives.

9.39.1 Options

`‘-z8001’` Generate segmented code by default.

`‘-z8002’` Generate unsegmented code by default.

9.39.2 Syntax

9.39.2.1 Special Characters

`‘!’` is the line comment character.

You can use `‘;’` instead of a newline to separate statements.

9.39.2.2 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 `‘rln’` and `‘rhn’`.

byte registers

```
r10 rh0 r11 rh1 r12 rh2 r13 rh3
r14 rh4 r15 rh5 r16 rh6 r17 rh7
```

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
```

9.39.2.3 Addressing Modes

as understands the following addressing modes for the Z8000:

`rln`

`rhn`

`rn`

`rrn`

`rqn` Register direct: 8bit, 16bit, 32bit, and 64bit registers.

`@rn`

`@rrn` Indirect register: `@rrn` in segmented mode, `@rn` in unsegmented mode.

| | |
|---|--|
| <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> <code>rrn(#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> <code>rrn(rm)</code> | Base Index: the 16 or 24 bit register <code>rn</code> or <code>rrn</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> . |

9.39.3 Assembler Directives for the Z8000

The Z8000 port of as includes additional assembler directives, for compatibility with other Z8000 assemblers. These do not begin with ‘.’ (unlike the ordinary as directives).

| | |
|---------------------|---|
| <code>segm</code> | |
| <code>.z8001</code> | Generate code for the segmented Z8001. |
| <code>unsegm</code> | |
| <code>.z8002</code> | Generate 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 ‘%xx’ (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. |

9.39.4 Opcodes

For detailed information on the Z8000 machine instruction set, see *Z8000 Technical Manual*.

9.40 VAX Dependent Features

9.40.1 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.

9.40.2 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.

9.40.3 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.

- | | |
|----------------------|---|
| <code>.dfloat</code> | This expects zero or more flonums, separated by commas, and assembles Vax <code>d</code> format 64-bit floating point constants. |
| <code>.ffloat</code> | This expects zero or more flonums, separated by commas, and assembles Vax <code>f</code> format 32-bit floating point constants. |
| <code>.gfloat</code> | This expects zero or more flonums, separated by commas, and assembles Vax <code>g</code> format 64-bit floating point constants. |
| <code>.hfloat</code> | This expects zero or more flonums, separated by commas, and assembles Vax <code>h</code> format 128-bit floating point constants. |

9.40.4 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.

9.40.5 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 | <p>‘Jsb’ is already an instruction mnemonic, so we chose ‘jbsb’.</p> <p>(byte displacement)</p> <p style="padding-left: 40px;"><i>bsbb ...</i></p> <p>(word displacement)</p> <p style="padding-left: 40px;"><i>bsbw ...</i></p> <p>(long displacement)</p> <p style="padding-left: 40px;"><i>jsb ...</i></p> |
| jbr | |
| jr | <p>Unconditional branch.</p> <p>(byte displacement)</p> <p style="padding-left: 40px;"><i>brb ...</i></p> <p>(word displacement)</p> <p style="padding-left: 40px;"><i>brw ...</i></p> <p>(long displacement)</p> <p style="padding-left: 40px;"><i>jmp ...</i></p> |
| jCOND | <p>COND may be any one of the conditional branches <i>neq</i>, <i>nequ</i>, <i>eql</i>, <i>eqlu</i>, <i>gtr</i>, <i>geq</i>, <i>lss</i>, <i>gtru</i>, <i>lequ</i>, <i>vc</i>, <i>vs</i>, <i>gequ</i>, <i>cc</i>, <i>lssu</i>, <i>cs</i>. COND may also be one of the bit tests <i>bs</i>, <i>bc</i>, <i>bss</i>, <i>bcs</i>, <i>bsc</i>, <i>bcc</i>, <i>bssi</i>, <i>bcci</i>, <i>lbs</i>, <i>lbc</i>. NOTCOND is the opposite condition to COND.</p> <p>(byte displacement)</p> <p style="padding-left: 40px;"><i>bCOND ...</i></p> <p>(word displacement)</p> <p style="padding-left: 40px;"><i>bNOTCOND foo ; brw ... ; foo:</i></p> <p>(long displacement)</p> <p style="padding-left: 40px;"><i>bNOTCOND foo ; jmp ... ; foo:</i></p> |
| jacbX | <p>X may be one of <i>b d f g h l w</i>.</p> <p>(word displacement)</p> <p style="padding-left: 40px;"><i>OPCODE ...</i></p> <p>(long displacement)</p> <p style="padding-left: 40px;"> <i>OPCODE ..., foo ;</i> <i>brb bar ;</i> <i>foo: jmp ... ;</i> <i>bar:</i> </p> |
| jaobYYY | <p>YYY may be one of <i>lss leq</i>.</p> |

`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:`

9.40.6 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.

9.40.7 Not Supported on VAX

Vax bit fields can not be assembled with **as**. Someone can add the required code if they really need it.

9.41 v850 Dependent Features

9.41.1 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.

-mv850e1 Specifies that the assembled code should be marked as being targeted at the V850E1 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.

-mv850e2 Specifies that the assembled code should be marked as being targeted at the V850E2 processor. This allows the linker to detect attempts to link such code with code assembled for other processors.

-mv850e2v3

Specifies that the assembled code should be marked as being targeted at the V850E2V3 processor. This allows the linker to detect attempts to link such code with code assembled for other processors.

`-mrelax` Enables relaxation. This allows the `.longcall` and `.longjump` pseudo ops to be used in the assembler source code. These ops label sections of code which are either a long function call or a long branch. The assembler will then flag these sections of code and the linker will attempt to relax them.

9.41.2 Syntax

9.41.2.1 Special Characters

`#` is the line comment character.

9.41.2.2 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

```

9.41.3 Floating Point

The V850 family uses IEEE floating-point numbers.

9.41.4 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.

.v850e1
    Specifies that the assembled code should be marked as being targeted at the
    V850E1 processor. This allows the linker to detect attempts to link such code
    with code assembled for other processors.

```

- `.v850e2` Specifies that the assembled code should be marked as being targeted at the V850E2 processor. This allows the linker to detect attempts to link such code with code assembled for other processors.
- `.v850e2v3` Specifies that the assembled code should be marked as being targeted at the V850E2V3 processor. This allows the linker to detect attempts to link such code with code assembled for other processors.

9.41.5 Opcodes

`as` implements all the standard V850 opcodes.

`as` also implements the following pseudo ops:

- `hi()` 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 (who's 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 (who's 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 (who's 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'.

.longcall name

Indicates that the following sequence of instructions is a long call to function **name**. The linker will attempt to shorten this call sequence if **name** is within a 22bit offset of the call. Only valid if the **-mrelax** command line switch has been enabled.

.longjump name

Indicates that the following sequence of instructions is a long jump to label **name**. The linker will attempt to shorten this code sequence if **name** is within a 22bit offset of the jump. Only valid if the **-mrelax** command line switch has been enabled.

For information on the V850 instruction set, see *V850 Family 32-/16-Bit single-Chip Microcontroller Architecture Manual* from NEC. Ltd.

9.42 Xtensa Dependent Features

This chapter covers features of the GNU assembler that are specific to the Xtensa architecture. For details about the Xtensa instruction set, please consult the *Xtensa Instruction Set Architecture (ISA) Reference Manual*.

9.42.1 Command Line Options

The Xtensa version of the GNU assembler supports these special options:

--text-section-literals | --no-text-section-literals

Control the treatment of literal pools. The default is ‘**--no-text-section-literals**’, which places literals in separate sections in the output file. This allows the literal pool to be placed in a data RAM/ROM. With ‘**--text-section-literals**’, the literals are interspersed in the text section in order to keep them as close as possible to their references. This may be necessary for large assembly files, where the literals would otherwise be out of range of the L32R instructions in the text section. These options only affect literals referenced via PC-relative L32R instructions; literals for absolute mode L32R instructions are handled separately. See [Section 9.42.5.4 \[literal\]](#), page 281.

--absolute-literals | --no-absolute-literals

Indicate to the assembler whether L32R instructions use absolute or PC-relative addressing. If the processor includes the absolute addressing option, the default is to use absolute L32R relocations. Otherwise, only the PC-relative L32R relocations can be used.

--target-align | --no-target-align

Enable or disable automatic alignment to reduce branch penalties at some expense in code size. See [Section 9.42.3.2 \[Automatic Instruction Alignment\]](#), page 277. This optimization is enabled by default. Note that the assembler will always align instructions like LOOP that have fixed alignment requirements.

--longcalls | --no-longcalls

Enable or disable transformation of call instructions to allow calls across a greater range of addresses. See [Section 9.42.4.2 \[Function Call Relaxation\]](#), page 279. This option should be used when call targets can potentially be out of range. It may degrade both code size and performance, but the linker can generally optimize away the unnecessary overhead when a call ends up within range. The default is ‘**--no-longcalls**’.

--transform | --no-transform

Enable or disable all assembler transformations of Xtensa instructions, including both relaxation and optimization. The default is ‘**--transform**’; ‘**--no-transform**’ should only be used in the rare cases when the instructions must be exactly as specified in the assembly source. Using ‘**--no-transform**’ causes out of range instruction operands to be errors.

--rename-section oldname=newname

Rename the *oldname* section to *newname*. This option can be used multiple times to rename multiple sections.

9.42.2 Assembler Syntax

Block comments are delimited by ‘/*’ and ‘*/’. End of line comments may be introduced with either ‘#’ or ‘//’.

Instructions consist of a leading opcode or macro name followed by whitespace and an optional comma-separated list of operands:

```
opcode [operand, ...]
```

Instructions must be separated by a newline or semicolon.

FLIX instructions, which bundle multiple opcodes together in a single instruction, are specified by enclosing the bundled opcodes inside braces:

```
{
  [format]
  opcode0 [operands]
  opcode1 [operands]
  opcode2 [operands]
  ...
}
```

The opcodes in a FLIX instruction are listed in the same order as the corresponding instruction slots in the TIE format declaration. Directives and labels are not allowed inside the braces of a FLIX instruction. A particular TIE format name can optionally be specified immediately after the opening brace, but this is usually unnecessary. The assembler will automatically search for a format that can encode the specified opcodes, so the format name need only be specified in rare cases where there is more than one applicable format and where it matters which of those formats is used. A FLIX instruction can also be specified on a single line by separating the opcodes with semicolons:

```
{ [format;] opcode0 [operands]; opcode1 [operands]; opcode2 [operands]; ... }
```

If an opcode can only be encoded in a FLIX instruction but is not specified as part of a FLIX bundle, the assembler will choose the smallest format where the opcode can be encoded and will fill unused instruction slots with no-ops.

9.42.2.1 Opcode Names

See the *Xtensa Instruction Set Architecture (ISA) Reference Manual* for a complete list of opcodes and descriptions of their semantics.

If an opcode name is prefixed with an underscore character (‘_’), `as` will not transform that instruction in any way. The underscore prefix disables both optimization (see [Section 9.42.3 \[Xtensa Optimizations\], page 277](#)) and relaxation (see [Section 9.42.4 \[Xtensa Relaxation\], page 278](#)) for that particular instruction. Only use the underscore prefix when it is essential to select the exact opcode produced by the assembler. Using this feature unnecessarily makes the code less efficient by disabling assembler optimization and less flexible by disabling relaxation.

Note that this special handling of underscore prefixes only applies to Xtensa opcodes, not to either built-in macros or user-defined macros. When an underscore prefix is used with a macro (e.g., `_MOV`), it refers to a different macro. The assembler generally provides built-in macros both with and without the underscore prefix, where the underscore versions behave as if the underscore carries through to the instructions in the macros. For example, `_MOV` may expand to `_MOV.N`.

The underscore prefix only applies to individual instructions, not to series of instructions. For example, if a series of instructions have underscore prefixes, the assembler will not transform the individual instructions, but it may insert other instructions between them (e.g., to align a LOOP instruction). To prevent the assembler from modifying a series of instructions as a whole, use the `no-transform` directive. See [Section 9.42.5.3 \[transform\]](#), [page 281](#).

9.42.2.2 Register Names

The assembly syntax for a register file entry is the “short” name for a TIE register file followed by the index into that register file. For example, the general-purpose AR register file has a short name of `a`, so these registers are named `a0...a15`. As a special feature, `sp` is also supported as a synonym for `a1`. Additional registers may be added by processor configuration options and by designer-defined TIE extensions. An initial ‘\$’ character is optional in all register names.

9.42.3 Xtensa Optimizations

The optimizations currently supported by `as` are generation of density instructions where appropriate and automatic branch target alignment.

9.42.3.1 Using Density Instructions

The Xtensa instruction set has a code density option that provides 16-bit versions of some of the most commonly used opcodes. Use of these opcodes can significantly reduce code size. When possible, the assembler automatically translates instructions from the core Xtensa instruction set into equivalent instructions from the Xtensa code density option. This translation can be disabled by using underscore prefixes (see [Section 9.42.2.1 \[Opcode Names\]](#), [page 276](#)), by using the ‘`--no-transform`’ command-line option (see [Section 9.42.1 \[Command Line Options\]](#), [page 275](#)), or by using the `no-transform` directive (see [Section 9.42.5.3 \[transform\]](#), [page 281](#)).

It is a good idea *not* to use the density instructions directly. The assembler will automatically select dense instructions where possible. If you later need to use an Xtensa processor without the code density option, the same assembly code will then work without modification.

9.42.3.2 Automatic Instruction Alignment

The Xtensa assembler will automatically align certain instructions, both to optimize performance and to satisfy architectural requirements.

As an optimization to improve performance, the assembler attempts to align branch targets so they do not cross instruction fetch boundaries. (Xtensa processors can be configured with either 32-bit or 64-bit instruction fetch widths.) An instruction immediately following a call is treated as a branch target in this context, because it will be the target of a return from the call. This alignment has the potential to reduce branch penalties at some expense in code size. This optimization is enabled by default. You can disable it with the ‘`--no-target-align`’ command-line option (see [Section 9.42.1 \[Command Line Options\]](#), [page 275](#)).

The target alignment optimization is done without adding instructions that could increase the execution time of the program. If there are density instructions in the code

preceding a target, the assembler can change the target alignment by widening some of those instructions to the equivalent 24-bit instructions. Extra bytes of padding can be inserted immediately following unconditional jump and return instructions. This approach is usually successful in aligning many, but not all, branch targets.

The LOOP family of instructions must be aligned such that the first instruction in the loop body does not cross an instruction fetch boundary (e.g., with a 32-bit fetch width, a LOOP instruction must be on either a 1 or 2 mod 4 byte boundary). The assembler knows about this restriction and inserts the minimal number of 2 or 3 byte no-op instructions to satisfy it. When no-op instructions are added, any label immediately preceding the original loop will be moved in order to refer to the loop instruction, not the newly generated no-op instruction. To preserve binary compatibility across processors with different fetch widths, the assembler conservatively assumes a 32-bit fetch width when aligning LOOP instructions (except if the first instruction in the loop is a 64-bit instruction).

Previous versions of the assembler automatically aligned ENTRY instructions to 4-byte boundaries, but that alignment is now the programmer's responsibility.

9.42.4 Xtensa Relaxation

When an instruction operand is outside the range allowed for that particular instruction field, **as** can transform the code to use a functionally-equivalent instruction or sequence of instructions. This process is known as *relaxation*. This is typically done for branch instructions because the distance of the branch targets is not known until assembly-time. The Xtensa assembler offers branch relaxation and also extends this concept to function calls, MOVI instructions and other instructions with immediate fields.

9.42.4.1 Conditional Branch Relaxation

When the target of a branch is too far away from the branch itself, i.e., when the offset from the branch to the target is too large to fit in the immediate field of the branch instruction, it may be necessary to replace the branch with a branch around a jump. For example,

```
    beqz    a2, L
may result in:
    bnez.n  a2, M
    j      L
M:
```

(The BNEZ.N instruction would be used in this example only if the density option is available. Otherwise, BNEZ would be used.)

This relaxation works well because the unconditional jump instruction has a much larger offset range than the various conditional branches. However, an error will occur if a branch target is beyond the range of a jump instruction. **as** cannot relax unconditional jumps. Similarly, an error will occur if the original input contains an unconditional jump to a target that is out of range.

Branch relaxation is enabled by default. It can be disabled by using underscore prefixes (see [Section 9.42.2.1 \[Opcode Names\], page 276](#)), the `--no-transform` command-line option (see [Section 9.42.1 \[Command Line Options\], page 275](#)), or the `no-transform` directive (see [Section 9.42.5.3 \[transform\], page 281](#)).

9.42.4.2 Function Call Relaxation

Function calls may require relaxation because the Xtensa immediate call instructions (`CALL0`, `CALL4`, `CALL8` and `CALL12`) provide a PC-relative offset of only 512 Kbytes in either direction. For larger programs, it may be necessary to use indirect calls (`CALLX0`, `CALLX4`, `CALLX8` and `CALLX12`) where the target address is specified in a register. The Xtensa assembler can automatically relax immediate call instructions into indirect call instructions. This relaxation is done by loading the address of the called function into the callee's return address register and then using a `CALLX` instruction. So, for example:

```
call8 func
might be relaxed to:
.literal .L1, func
l32r    a8, .L1
callx8  a8
```

Because the addresses of targets of function calls are not generally known until link-time, the assembler must assume the worst and relax all the calls to functions in other source files, not just those that really will be out of range. The linker can recognize calls that were unnecessarily relaxed, and it will remove the overhead introduced by the assembler for those cases where direct calls are sufficient.

Call relaxation is disabled by default because it can have a negative effect on both code size and performance, although the linker can usually eliminate the unnecessary overhead. If a program is too large and some of the calls are out of range, function call relaxation can be enabled using the `--longcalls` command-line option or the `longcalls` directive (see [Section 9.42.5.2 \[longcalls\]](#), page 281).

9.42.4.3 Other Immediate Field Relaxation

The assembler normally performs the following other relaxations. They can be disabled by using underscore prefixes (see [Section 9.42.2.1 \[Opcode Names\]](#), page 276), the `--no-transform` command-line option (see [Section 9.42.1 \[Command Line Options\]](#), page 275), or the `no-transform` directive (see [Section 9.42.5.3 \[transform\]](#), page 281).

The `MOVI` machine instruction can only materialize values in the range from -2048 to 2047. Values outside this range are best materialized with `L32R` instructions. Thus:

```
movi a0, 100000
is assembled into the following machine code:
.literal .L1, 100000
l32r a0, .L1
```

The `L8UI` machine instruction can only be used with immediate offsets in the range from 0 to 255. The `L16SI` and `L16UI` machine instructions can only be used with offsets from 0 to 510. The `L32I` machine instruction can only be used with offsets from 0 to 1020. A load offset outside these ranges can be materialized with an `L32R` instruction if the destination register of the load is different than the source address register. For example:

```
l32i a1, a0, 2040
is translated to:
.literal .L1, 2040
l32r a1, .L1
add a1, a0, a1
l32i a1, a1, 0
```

If the load destination and source address register are the same, an out-of-range offset causes an error.

The Xtensa `ADDI` instruction only allows immediate operands in the range from -128 to 127. There are a number of alternate instruction sequences for the `ADDI` operation. First, if the immediate is 0, the `ADDI` will be turned into a `MOV.N` instruction (or the equivalent `OR` instruction if the code density option is not available). If the `ADDI` immediate is outside of the range -128 to 127, but inside the range -32896 to 32639, an `ADDMI` instruction or `ADDMI/ADDI` sequence will be used. Finally, if the immediate is outside of this range and a free register is available, an `L32R/ADD` sequence will be used with a literal allocated from the literal pool.

For example:

```
addi    a5, a6, 0
addi    a5, a6, 512
addi    a5, a6, 513
addi    a5, a6, 50000
```

is assembled into the following:

```
.literal .L1, 50000
mov.n   a5, a6
addmi   a5, a6, 0x200
addmi   a5, a6, 0x200
addi    a5, a5, 1
l32r    a5, .L1
add     a5, a6, a5
```

9.42.5 Directives

The Xtensa assembler supports a region-based directive syntax:

```
.begin directive [options]
...
.end directive
```

All the Xtensa-specific directives that apply to a region of code use this syntax.

The directive applies to code between the `.begin` and the `.end`. The state of the option after the `.end` reverts to what it was before the `.begin`. A nested `.begin/.end` region can further change the state of the directive without having to be aware of its outer state. For example, consider:

```
.begin no-transform
L:  add a0, a1, a2
    .begin transform
M:  add a0, a1, a2
    .end transform
N:  add a0, a1, a2
    .end no-transform
```

The `ADD` opcodes at L and N in the outer `no-transform` region both result in `ADD` machine instructions, but the assembler selects an `ADD.N` instruction for the `ADD` at M in the inner `transform` region.

The advantage of this style is that it works well inside macros which can preserve the context of their callers.

The following directives are available:

9.42.5.1 `schedule`

The `schedule` directive is recognized only for compatibility with Tensilica’s assembler.

```
.begin [no-]schedule
.end [no-]schedule
```

This directive is ignored and has no effect on `as`.

9.42.5.2 `longcalls`

The `longcalls` directive enables or disables function call relaxation. See [Section 9.42.4.2 \[Function Call Relaxation\]](#), page 279.

```
.begin [no-]longcalls
.end [no-]longcalls
```

Call relaxation is disabled by default unless the ‘`--longcalls`’ command-line option is specified. The `longcalls` directive overrides the default determined by the command-line options.

9.42.5.3 `transform`

This directive enables or disables all assembler transformation, including relaxation (see [Section 9.42.4 \[Xtensa Relaxation\]](#), page 278) and optimization (see [Section 9.42.3 \[Xtensa Optimizations\]](#), page 277).

```
.begin [no-]transform
.end [no-]transform
```

Transformations are enabled by default unless the ‘`--no-transform`’ option is used. The `transform` directive overrides the default determined by the command-line options. An underscore opcode prefix, disabling transformation of that opcode, always takes precedence over both directives and command-line flags.

9.42.5.4 `literal`

The `.literal` directive is used to define literal pool data, i.e., read-only 32-bit data accessed via L32R instructions.

```
.literal label, value[, value...]
```

This directive is similar to the standard `.word` directive, except that the actual location of the literal data is determined by the assembler and linker, not by the position of the `.literal` directive. Using this directive gives the assembler freedom to locate the literal data in the most appropriate place and possibly to combine identical literals. For example, the code:

```
entry sp, 40
.literal .L1, sym
l32r    a4, .L1
```

can be used to load a pointer to the symbol `sym` into register `a4`. The value of `sym` will not be placed between the `ENTRY` and `L32R` instructions; instead, the assembler puts the data in a literal pool.

Literal pools are placed by default in separate literal sections; however, when using the ‘`--text-section-literals`’ option (see [Section 9.42.1 \[Command Line Options\]](#), page 275), the literal pools for PC-relative mode L32R instructions are placed in the

current section.¹ These text section literal pools are created automatically before `ENTRY` instructions and manually after `‘.literal_position’` directives (see [Section 9.42.5.5 \[literal_position\]](#), page 282). If there are no preceding `ENTRY` instructions, explicit `.literal_position` directives must be used to place the text section literal pools; otherwise, `as` will report an error.

When literals are placed in separate sections, the literal section names are derived from the names of the sections where the literals are defined. The base literal section names are `.literal` for PC-relative mode L32R instructions and `.lit4` for absolute mode L32R instructions (see [Section 9.42.5.7 \[absolute-literals\]](#), page 283). These base names are used for literals defined in the default `.text` section. For literals defined in other sections or within the scope of a `literal_prefix` directive (see [Section 9.42.5.6 \[literal_prefix\]](#), page 283), the following rules determine the literal section name:

1. If the current section is a member of a section group, the literal section name includes the group name as a suffix to the base `.literal` or `.lit4` name, with a period to separate the base name and group name. The literal section is also made a member of the group.
2. If the current section name (or `literal_prefix` value) begins with `“.gnu.linkonce.kind.”`, the literal section name is formed by replacing `“kind”` with the base `.literal` or `.lit4` name. For example, for literals defined in a section named `.gnu.linkonce.t.func`, the literal section will be `.gnu.linkonce.literal.func` or `.gnu.linkonce.lit4.func`.
3. If the current section name (or `literal_prefix` value) ends with `.text`, the literal section name is formed by replacing that suffix with the base `.literal` or `.lit4` name. For example, for literals defined in a section named `.iram0.text`, the literal section will be `.iram0.literal` or `.iram0.lit4`.
4. If none of the preceding conditions apply, the literal section name is formed by adding the base `.literal` or `.lit4` name as a suffix to the current section name (or `literal_prefix` value).

9.42.5.5 literal_position

When using `‘--text-section-literals’` to place literals inline in the section being assembled, the `.literal_position` directive can be used to mark a potential location for a literal pool.

```
.literal_position
```

The `.literal_position` directive is ignored when the `‘--text-section-literals’` option is not used or when L32R instructions use the absolute addressing mode.

The assembler will automatically place text section literal pools before `ENTRY` instructions, so the `.literal_position` directive is only needed to specify some other location for a literal pool. You may need to add an explicit jump instruction to skip over an inline literal pool.

For example, an interrupt vector does not begin with an `ENTRY` instruction so the assembler will be unable to automatically find a good place to put a literal pool. Moreover,

¹ Literals for the `.init` and `.fini` sections are always placed in separate sections, even when `‘--text-section-literals’` is enabled.

the code for the interrupt vector must be at a specific starting address, so the literal pool cannot come before the start of the code. The literal pool for the vector must be explicitly positioned in the middle of the vector (before any uses of the literals, due to the negative offsets used by PC-relative L32R instructions). The `.literal_position` directive can be used to do this. In the following code, the literal for ‘M’ will automatically be aligned correctly and is placed after the unconditional jump.

```
.global M
code_start:
    j continue
    .literal_position
    .align 4
continue:
    movi    a4, M
```

9.42.5.6 literal_prefix

The `literal_prefix` directive allows you to override the default literal section names, which are derived from the names of the sections where the literals are defined.

```
.begin literal_prefix [name]
.end literal_prefix
```

For literals defined within the delimited region, the literal section names are derived from the *name* argument instead of the name of the current section. The rules used to derive the literal section names do not change. See [Section 9.42.5.4 \[literal\]](#), page 281. If the *name* argument is omitted, the literal sections revert to the defaults. This directive has no effect when using the ‘`--text-section-literals`’ option (see [Section 9.42.1 \[Command Line Options\]](#), page 275).

9.42.5.7 absolute-literals

The `absolute-literals` and `no-absolute-literals` directives control the absolute vs. PC-relative mode for L32R instructions. These are relevant only for Xtensa configurations that include the absolute addressing option for L32R instructions.

```
.begin [no-]absolute-literals
.end [no-]absolute-literals
```

These directives do not change the L32R mode—they only cause the assembler to emit the appropriate kind of relocation for L32R instructions and to place the literal values in the appropriate section. To change the L32R mode, the program must write the `LITBASE` special register. It is the programmer’s responsibility to keep track of the mode and indicate to the assembler which mode is used in each region of code.

If the Xtensa configuration includes the absolute L32R addressing option, the default is to assume absolute L32R addressing unless the ‘`--no-absolute-literals`’ command-line option is specified. Otherwise, the default is to assume PC-relative L32R addressing. The `absolute-literals` directive can then be used to override the default determined by the command-line options.

10 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.

10.1 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.

10.2 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 <http://www.sourceware.org/bugzilla/>.

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?” This cannot help us fix a bug, so it is basically useless. We respond by asking for enough details to

enable us to investigate. You might as well expedite matters by sending them to begin with.

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 sync, 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.

11 Acknowledgements

If you have contributed to GAS 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 Renesas H8/300 processors (tc-z8k, tc-h8300), 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 GAS 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 Eichin 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 “IBM 370” architecture.

Richard Henderson rewrote the Alpha assembler. Klaus Kaempf wrote GAS and BFD support for openVMS/Alpha.

Timothy Wall, Michael Hayes, and Greg Smart contributed to the various tic* flavors.

David Heine, Sterling Augustine, Bob Wilson and John Ruttenberg from Tensilica, Inc. added support for Xtensa processors.

Several engineers at Cygnus Support have also provided many small bug fixes and configuration enhancements.

Jon Beniston added support for the Lattice Mico32 architecture.

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.

Appendix A GNU Free Documentation License

Version 1.3, 3 November 2008

Copyright © 2000, 2001, 2002, 2007, 2008 Free Software Foundation, Inc.

<http://fsf.org/>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document *free* in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or non-commercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of “copyleft”, which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The “Document”, below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as “you”. You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A “Modified Version” of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A “Secondary Section” is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document’s overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The “Invariant Sections” are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released

under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The “Cover Texts” are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A “Transparent” copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not “Transparent” is called “Opaque”.

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The “Title Page” means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, “Title Page” means the text near the most prominent appearance of the work’s title, preceding the beginning of the body of the text.

The “publisher” means any person or entity that distributes copies of the Document to the public.

A section “Entitled XYZ” means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as “Acknowledgements”, “Dedications”, “Endorsements”, or “History”.) To “Preserve the Title” of such a section when you modify the Document means that it remains a section “Entitled XYZ” according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any,

be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.

- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their

titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements."

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an “aggregate” if the copyright resulting from the compilation is not used to limit the legal rights of the compilation’s users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document’s Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled “Acknowledgements”, “Dedications”, or “History”, the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, or distribute it is void, and will automatically terminate your rights under this License.

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, receipt of a copy of some or all of the same material does not give you any rights to use it.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License “or any later version” applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation. If the Document specifies that a proxy can decide which future versions of this License can be used, that proxy’s public statement of acceptance of a version permanently authorizes you to choose that version for the Document.

11. RELICENSING

“Massive Multiauthor Collaboration Site” (or “MMC Site”) means any World Wide Web server that publishes copyrightable works and also provides prominent facilities for anybody to edit those works. A public wiki that anybody can edit is an example of such a server. A “Massive Multiauthor Collaboration” (or “MMC”) contained in the site means any set of copyrightable works thus published on the MMC site.

“CC-BY-SA” means the Creative Commons Attribution-Share Alike 3.0 license published by Creative Commons Corporation, a not-for-profit corporation with a principal place of business in San Francisco, California, as well as future copyleft versions of that license published by that same organization.

“Incorporate” means to publish or republish a Document, in whole or in part, as part of another Document.

An MMC is “eligible for relicensing” if it is licensed under this License, and if all works that were first published under this License somewhere other than this MMC, and subsequently incorporated in whole or in part into the MMC, (1) had no cover texts or invariant sections, and (2) were thus incorporated prior to November 1, 2008.

The operator of an MMC Site may republish an MMC contained in the site under CC-BY-SA on the same site at any time before August 1, 2009, provided the MMC is eligible for relicensing.

ADDENDUM: How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

```
Copyright (C)  year  your name.
Permission is granted to copy, distribute and/or modify this document
under the terms of the GNU Free Documentation License, Version 1.3
or any later version published by the Free Software Foundation;
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover
Texts. A copy of the license is included in the section entitled ‘‘GNU
Free Documentation License’’.
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “with...Texts.” line with this:

```
with the Invariant Sections being list their titles, with
the Front-Cover Texts being list, and with the Back-Cover Texts
being list.
```

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.

AS Index

(Index is nonexistent)

