NASM — The Netwide Assembler

version 0.98.39

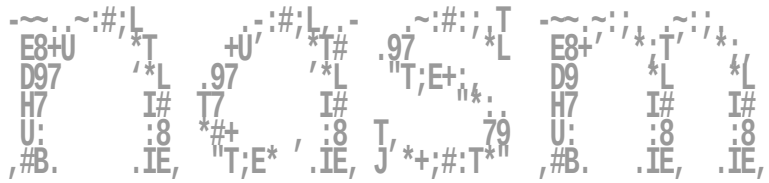

© 2003 The NASM Development Team

All rights reserved. This document is redistributable under the licence given in the file "COPYING" distributed in the NASM archive.

Contents

Chapter 1: Introduction 20
1.1 What Is NASM? 20
1.1.1 Why Yet Another Assembler? 20
1.1.2 Licence Conditions 20
1.2 Contact Information 20
1.3 Installation 21
1.3.1 Installing NASM under MS–DOS or Windows 21
1.3.2 Installing NASM under Unix 21
Chapter 2: Running NASM 23
2.1 NASM Command–Line Syntax 23
2.1.1 The `-o` Option: Specifying the Output File Name 23
2.1.2 The `-f` Option: Specifying the Output File Format 24
2.1.3 The `-l` Option: Generating a Listing File 24
2.1.4 The `-M` Option: Generate Makefile Dependencies 24
2.1.5 The `-F` Option: Selecting a Debug Information Format 24
2.1.6 The `-g` Option: Enabling Debug Information 24
2.1.7 The `-X` Option: Selecting an Error Reporting Format 25
2.1.8 The `-E` Option: Send Errors to a File 25
2.1.9 The `-s` Option: Send Errors to `stdout` 25
2.1.10 The `-i` Option: Include File Search Directories 25
2.1.11 The `-p` Option: Pre–Include a File 26
2.1.12 The `-d` Option: Pre–Define a Macro 26
2.1.13 The `-u` Option: Undefine a Macro 26
2.1.14 The `-e` Option: Preprocess Only 26
2.1.15 The `-a` Option: Don't Preprocess At All 27
2.1.16 The `-On` Option: Specifying Multipass Optimization 27
2.1.17 The `-t` option: Enable TASM Compatibility Mode 27
2.1.18 The `-w` Option: Enable or Disable Assembly Warnings 28
2.1.19 The `-v` Option: Display Version Info 28
2.1.20 The `-y` Option: Display Available Debug Info Formats 28

2.1.21 The `--prefix` and `--postfix` Options. 28
2.1.22 The `NASMENV` Environment Variable 28
2.2 Quick Start for MASM Users 29
2.2.1 NASM Is Case–Sensitive 29
2.2.2 NASM Requires Square Brackets For Memory References 29
2.2.3 NASM Doesn't Store Variable Types 30
2.2.4 NASM Doesn't `ASSUME` 30
2.2.5 NASM Doesn't Support Memory Models 30
2.2.6 Floating–Point Differences 30
2.2.7 Other Differences. 30
Chapter 3: The NASM Language 31
3.1 Layout of a NASM Source Line 31
3.2 Pseudo–Instructions 32
3.2.1 `DB` and friends: Declaring Initialised Data 32
3.2.2 `RESB` and friends: Declaring Uninitialised Data 32
3.2.3 `INCBIN`: Including External Binary Files 32
3.2.4 `EQU`: Defining Constants. 33
3.2.5 `TIMES`: Repeating Instructions or Data 33
3.3 Effective Addresses 33
3.4 Constants. 34
3.4.1 Numeric Constants 34
3.4.2 Character Constants 35
3.4.3 String Constants 35
3.4.4 Floating–Point Constants 35
3.5 Expressions 36
3.5.1 `|`: Bitwise OR Operator 36
3.5.2 `^`: Bitwise XOR Operator 36
3.5.3 `&`: Bitwise AND Operator 36
3.5.4 << and >>: Bit Shift Operators 36
3.5.5 + and −: Addition and Subtraction Operators 36
3.5.6 `*`, `/`, `//`, `%` and `%%`: Multiplication and Division 36
3.5.7 Unary Operators: +, −, ~ and `SEG` 36
3.6 `SEG` and `WRT` 37
3.7 `STRICT`: Inhibiting Optimization 37
3.8 Critical Expressions 38

3.9 Local Labels 39
Chapter 4: The NASM Preprocessor 41
4.1 Single−Line Macros 41
4.1.1 The Normal Way: `%define` 41
4.1.2 Enhancing %define: `%xdefine` 42
4.1.3 Concatenating Single Line Macro Tokens: `%+` 43
4.1.4 Undefining macros: `%undef` 43
4.1.5 Preprocessor Variables: `%assign` 43
4.2 String Handling in Macros: `%strlen` and `%substr` 44
4.2.1 String Length: `%strlen` 44
4.2.2 Sub−strings: `%substr` 44
4.3 Multi−Line Macros: `%macro` 44
4.3.1 Overloading Multi−Line Macros 45
4.3.2 Macro−Local Labels 46
4.3.3 Greedy Macro Parameters 46
4.3.4 Default Macro Parameters 47
4.3.5 `%0`: Macro Parameter Counter 48
4.3.6 `%rotate`: Rotating Macro Parameters 48
4.3.7 Concatenating Macro Parameters 49
4.3.8 Condition Codes as Macro Parameters 50
4.3.9 Disabling Listing Expansion 50
4.4 Conditional Assembly 50
4.4.1 `%ifdef`: Testing Single−Line Macro Existence 51
4.4.2 `ifmacro`: Testing Multi−Line Macro Existence 51
4.4.3 `%ifctx`: Testing the Context Stack 51
4.4.4 `%if`: Testing Arbitrary Numeric Expressions 52
4.4.5 `%ifidn` and `%ifidni`: Testing Exact Text Identity 52
4.4.6 `%ifid`, `%ifnum`, `%ifstr`: Testing Token Types 52
4.4.7 `%error`: Reporting User−Defined Errors 53
4.5 Preprocessor Loops: `%rep` 54
4.6 Including Other Files 54
4.7 The Context Stack 55
4.7.1 `%push` and `%pop`: Creating and Removing Contexts 55
4.7.2 Context−Local Labels 55
4.7.3 Context−Local Single−Line Macros 56

4.7.4 %repl: Renaming a Context 56
4.7.5 Example Use of the Context Stack: Block IFs 56
4.8 Standard Macros 58
4.8.1 __NASM_MAJOR__, __NASM_MINOR__, __NASM_SUBMINOR__ and ____NASM_PATCHLEVEL__: NASM Version 58
4.8.2 __NASM_VERSION_ID__: NASM Version ID 58
4.8.3 __NASM_VER__: NASM Version string 58
4.8.4 __FILE__ and __LINE__: File Name and Line Number 58
4.8.5 STRUC and ENDSTRUC: Declaring Structure Data Types 59
4.8.6 ISTRUC, AT and IEND: Declaring Instances of Structures 60
4.8.7 ALIGN and ALIGNB: Data Alignment 60
4.9 TASM Compatible Preprocessor Directives 61
4.9.1 %arg Directive 61
4.9.2 %stacksize Directive 62
4.9.3 %local Directive 62
4.10 Other Preprocessor Directives 63
4.10.1 %line Directive 63
4.10.2 %!<env>: Read an environment variable 63
Chapter 5: Assembler Directives 65
5.1 BITS: Specifying Target Processor Mode 65
5.1.1 USE16 & USE32: Aliases for BITS 65
5.2 SECTION or SEGMENT: Changing and Defining Sections 65
5.2.1 The __SECT__ Macro 66
5.3 ABSOLUTE: Defining Absolute Labels 66
5.4 EXTERN: Importing Symbols from Other Modules 67
5.5 GLOBAL: Exporting Symbols to Other Modules 68
5.6 COMMON: Defining Common Data Areas 68
5.7 CPU: Defining CPU Dependencies 68
Chapter 6: Output Formats 70
6.1 bin: Flat–Form Binary Output 70
6.1.1 ORG: Binary File Program Origin 70
6.1.2 bin Extensions to the SECTION Directive 70
6.1.3 Multisection support for the BIN format 71
6.1.4 Map files 71
6.2 obj: Microsoft OMF Object Files 71

6.2.1 obj Extensions to the SEGMENT Directive . . . 72
6.2.2 GROUP: Defining Groups of Segments . . . 73
6.2.3 UPPERCASE: Disabling Case Sensitivity in Output . . . 73
6.2.4 IMPORT: Importing DLL Symbols . . . 74
6.2.5 EXPORT: Exporting DLL Symbols . . . 74
6.2.6 ..start: Defining the Program Entry Point . . . 74
6.2.7 obj Extensions to the EXTERN Directive . . . 75
6.2.8 obj Extensions to the COMMON Directive . . . 75
6.3 win32: Microsoft Win32 Object Files . . . 76
6.3.1 win32 Extensions to the SECTION Directive . . . 76
6.4 coff: Common Object File Format . . . 77
6.5 elf: Executable and Linkable Format Object Files . . . 77
6.5.1 elf Extensions to the SECTION Directive . . . 77
6.5.2 Position−Independent Code: elf Special Symbols and WRT . . . 77
6.5.3 elf Extensions to the GLOBAL Directive . . . 78
6.5.4 elf Extensions to the COMMON Directive . . . 79
6.5.5 16−bit code and ELF . . . 79
6.6 aout: Linux a.out Object Files . . . 79
6.7 aoutb: NetBSD/FreeBSD/OpenBSD a.out Object Files . . . 79
6.8 as86: Minix/Linux as86 Object Files . . . 79
6.9 rdf: Relocatable Dynamic Object File Format . . . 80
6.9.1 Requiring a Library: The LIBRARY Directive . . . 80
6.9.2 Specifying a Module Name: The MODULE Directive . . . 80
6.9.3 rdf Extensions to the GLOBAL directive . . . 80
6.9.4 rdf Extensions to the EXTERN directive . . . 81
6.10 dbg: Debugging Format . . . 81
Chapter 7: Writing 16−bit Code (DOS, Windows 3/3.1) . . . 82
7.1 Producing .EXE Files . . . 82
7.1.1 Using the obj Format To Generate .EXE Files . . . 82
7.1.2 Using the bin Format To Generate .EXE Files . . . 83
7.2 Producing .COM Files . . . 84
7.2.1 Using the bin Format To Generate .COM Files . . . 84
7.2.2 Using the obj Format To Generate .COM Files . . . 85
7.3 Producing .SYS Files . . . 85
7.4 Interfacing to 16−bit C Programs . . . 85

7.4.1 External Symbol Names 85
7.4.2 Memory Models 86
7.4.3 Function Definitions and Function Calls 87
7.4.4 Accessing Data Items 89
7.4.5 `c16.mac`: Helper Macros for the 16-bit C Interface 90
7.5 Interfacing to Borland Pascal Programs 91
7.5.1 The Pascal Calling Convention 91
7.5.2 Borland Pascal Segment Name Restrictions 92
7.5.3 Using `c16.mac` With Pascal Programs 93
Chapter 8: Writing 32-bit Code (Unix, Win32, DJGPP) 94
8.1 Interfacing to 32-bit C Programs 94
8.1.1 External Symbol Names 94
8.1.2 Function Definitions and Function Calls 94
8.1.3 Accessing Data Items 96
8.1.4 `c32.mac`: Helper Macros for the 32-bit C Interface 96
8.2 Writing NetBSD/FreeBSD/OpenBSD and Linux/ELF Shared Libraries 97
8.2.1 Obtaining the Address of the GOT 97
8.2.2 Finding Your Local Data Items 98
8.2.3 Finding External and Common Data Items 99
8.2.4 Exporting Symbols to the Library User 99
8.2.5 Calling Procedures Outside the Library 100
8.2.6 Generating the Library File 100
Chapter 9: Mixing 16 and 32 Bit Code 101
9.1 Mixed-Size Jumps 101
9.2 Addressing Between Different-Size Segments 101
9.3 Other Mixed-Size Instructions 102
Chapter 10: Troubleshooting 104
10.1 Common Problems 104
10.1.1 NASM Generates Inefficient Code 104
10.1.2 My Jumps are Out of Range 104
10.1.3 `ORG` Doesn't Work 104
10.1.4 `TIMES` Doesn't Work 105
10.2 Bugs 105
Appendix A: Ndisasm 107
A.1 Introduction 107

A.2 Getting Started: Installation . . . 107
A.3 Running NDISASM . . . 107
A.3.1 COM Files: Specifying an Origin . . . 107
A.3.2 Code Following Data: Synchronisation . . . 107
A.3.3 Mixed Code and Data: Automatic (Intelligent) Synchronisation . . . 108
A.3.4 Other Options . . . 109
A.4 Bugs and Improvements . . . 109
Appendix B: x86 Instruction Reference . . . 110
B.1 Key to Operand Specifications . . . 110
B.2 Key to Opcode Descriptions . . . 110
B.2.1 Register Values . . . 111
B.2.2 Condition Codes . . . 112
B.2.3 SSE Condition Predicates . . . 112
B.2.4 Status Flags . . . 113
B.2.5 Effective Address Encoding: ModR/M and SIB . . . 114
B.3 Key to Instruction Flags . . . 115
B.4 x86 Instruction Set . . . 116
B.4.1 AAA, AAS, AAM, AAD: ASCII Adjustments . . . 116
B.4.2 ADC: Add with Carry . . . 116
B.4.3 ADD: Add Integers . . . 117
B.4.4 ADDPD: ADD Packed Double−Precision FP Values . . . 118
B.4.5 ADDPS: ADD Packed Single−Precision FP Values . . . 118
B.4.6 ADDSD: ADD Scalar Double−Precision FP Values . . . 118
B.4.7 ADDSS: ADD Scalar Single−Precision FP Values . . . 118
B.4.8 AND: Bitwise AND . . . 118
B.4.9 ANDNPD: Bitwise Logical AND NOT of Packed Double−Precision FP Values . . . 119
B.4.10 ANDNPS: Bitwise Logical AND NOT of Packed Single−Precision FP Values . . . 119
B.4.11 ANDPD: Bitwise Logical AND For Single FP . . . 119
B.4.12 ANDPS: Bitwise Logical AND For Single FP . . . 120
B.4.13 ARPL: Adjust RPL Field of Selector . . . 120
B.4.14 BOUND: Check Array Index against Bounds . . . 120
B.4.15 BSF, BSR: Bit Scan . . . 120
B.4.16 BSWAP: Byte Swap . . . 120
B.4.17 BT, BTC, BTR, BTS: Bit Test . . . 121
B.4.18 CALL: Call Subroutine . . . 121

B.4.19 CBW, CWD, CDQ, CWDE: Sign Extensions 122
B.4.20 CLC, CLD, CLI, CLTS: Clear Flags 122
B.4.21 CLFLUSH: Flush Cache Line 122
B.4.22 CMC: Complement Carry Flag 123
B.4.23 CMOVcc: Conditional Move 123
B.4.24 CMP: Compare Integers 123
B.4.25 CMPccPD: Packed Double−Precision FP Compare 123
B.4.26 CMPccPS: Packed Single−Precision FP Compare 124
B.4.27 CMPSB, CMPSW, CMPSD: Compare Strings 125
B.4.28 CMPccSD: Scalar Double−Precision FP Compare 125
B.4.29 CMPccSS: Scalar Single−Precision FP Compare 126
B.4.30 CMPXCHG, CMPXCHG486: Compare and Exchange 126
B.4.31 CMPXCHG8B: Compare and Exchange Eight Bytes 127
B.4.32 COMISD: Scalar Ordered Double−Precision FP Compare and Set EFLAGS 127
B.4.33 COMISS: Scalar Ordered Single−Precision FP Compare and Set EFLAGS. 127
B.4.34 CPUID: Get CPU Identification Code 128
B.4.35 CVTDQ2PD: Packed Signed INT32 to Packed Double−Precision FP Conversion . . 128
B.4.36 CVTDQ2PS: Packed Signed INT32 to Packed Single−Precision FP Conversion . . . 128
B.4.37 CVTPD2DQ: Packed Double−Precision FP to Packed Signed INT32 Conversion . . 128
B.4.38 CVTPD2PI: Packed Double−Precision FP to Packed Signed INT32 Conversion . . 129
B.4.39 CVTPD2PS: Packed Double−Precision FP to Packed Single−Precision FP Conversion 129
B.4.40 CVTPI2PD: Packed Signed INT32 to Packed Double−Precision FP Conversion . . 129
B.4.41 CVTPI2PS: Packed Signed INT32 to Packed Single−FP Conversion 129
B.4.42 CVTPS2DQ: Packed Single−Precision FP to Packed Signed INT32 Conversion . . . 129
B.4.43 CVTPS2PD: Packed Single−Precision FP to Packed Double−Precision FP Conversion 130
B.4.44 CVTPS2PI: Packed Single−Precision FP to Packed Signed INT32 Conversion . . . 130
B.4.45 CVTSD2SI: Scalar Double−Precision FP to Signed INT32 Conversion 130
B.4.46 CVTSD2SS: Scalar Double−Precision FP to Scalar Single−Precision FP Conversion 130
B.4.47 CVTSI2SD: Signed INT32 to Scalar Double−Precision FP Conversion 130
B.4.48 CVTSI2SS: Signed INT32 to Scalar Single−Precision FP Conversion 131
B.4.49 CVTSS2SD: Scalar Single−Precision FP to Scalar Double−Precision FP Conversion 131
B.4.50 CVTSS2SI: Scalar Single−Precision FP to Signed INT32 Conversion 131

B.4.51 CVTTPD2DQ: Packed Double−Precision FP to Packed Signed INT32 Conversion with Truncation . . . 131
B.4.52 CVTTPD2PI: Packed Double−Precision FP to Packed Signed INT32 Conversion with Truncation . . . 131
B.4.53 CVTTPS2DQ: Packed Single−Precision FP to Packed Signed INT32 Conversion with Truncation . . . 132
B.4.54 CVTTPS2PI: Packed Single−Precision FP to Packed Signed INT32 Conversion with Truncation . . . 132
B.4.55 CVTTSD2SI: Scalar Double−Precision FP to Signed INT32 Conversion with Truncation . . . 132
B.4.56 CVTTSS2SI: Scalar Single−Precision FP to Signed INT32 Conversion with Truncation . . . 132
B.4.57 DAA, DAS: Decimal Adjustments . . . 132
B.4.58 DEC: Decrement Integer . . . 133
B.4.59 DIV: Unsigned Integer Divide . . . 133
B.4.60 DIVPD: Packed Double−Precision FP Divide . . . 133
B.4.61 DIVPS: Packed Single−Precision FP Divide . . . 134
B.4.62 DIVSD: Scalar Double−Precision FP Divide . . . 134
B.4.63 DIVSS: Scalar Single−Precision FP Divide . . . 134
B.4.64 EMMS: Empty MMX State . . . 134
B.4.65 ENTER: Create Stack Frame . . . 134
B.4.66 F2XM1: Calculate 2**X−1 . . . 135
B.4.67 FABS: Floating−Point Absolute Value . . . 135
B.4.68 FADD, FADDP: Floating−Point Addition . . . 135
B.4.69 FBLD, FBSTP: BCD Floating−Point Load and Store . . . 135
B.4.70 FCHS: Floating−Point Change Sign . . . 136
B.4.71 FCLEX, FNCLEX: Clear Floating−Point Exceptions . . . 136
B.4.72 FCMOVcc: Floating−Point Conditional Move . . . 136
B.4.73 FCOM, FCOMP, FCOMPP, FCOMI, FCOMIP: Floating−Point Compare . . . 137
B.4.74 FCOS: Cosine . . . 137
B.4.75 FDECSTP: Decrement Floating−Point Stack Pointer . . . 137
B.4.76 FxDISI, FxENI: Disable and Enable Floating−Point Interrupts . . . 137
B.4.77 FDIV, FDIVP, FDIVR, FDIVRP: Floating−Point Division . . . 138
B.4.78 FEMMS: Faster Enter/Exit of the MMX or floating−point state . . . 138
B.4.79 FFREE: Flag Floating−Point Register as Unused . . . 138
B.4.80 FIADD: Floating−Point/Integer Addition . . . 139
B.4.81 FICOM, FICOMP: Floating−Point/Integer Compare . . . 139

B.4.82 FIDIV, FIDIVR: Floating–Point/Integer Division 139
B.4.83 FILD, FIST, FISTP: Floating–Point/Integer Conversion 139
B.4.84 FIMUL: Floating–Point/Integer Multiplication 139
B.4.85 FINCSTP: Increment Floating–Point Stack Pointer 139
B.4.86 FINIT, FNINIT: Initialise Floating–Point Unit 140
B.4.87 FISUB: Floating–Point/Integer Subtraction 140
B.4.88 FLD: Floating–Point Load 140
B.4.89 FLDxx: Floating–Point Load Constants 140
B.4.90 FLDCW: Load Floating–Point Control Word 140
B.4.91 FLDENV: Load Floating–Point Environment 141
B.4.92 FMUL, FMULP: Floating–Point Multiply 141
B.4.93 FNOP: Floating–Point No Operation 141
B.4.94 FPATAN, FPTAN: Arctangent and Tangent 141
B.4.95 FPREM, FPREM1: Floating–Point Partial Remainder 141
B.4.96 FRNDINT: Floating–Point Round to Integer 142
B.4.97 FSAVE, FRSTOR: Save/Restore Floating–Point State 142
B.4.98 FSCALE: Scale Floating–Point Value by Power of Two 142
B.4.99 FSETPM: Set Protected Mode 142
B.4.100 FSIN, FSINCOS: Sine and Cosine 142
B.4.101 FSQRT: Floating–Point Square Root. 142
B.4.102 FST, FSTP: Floating–Point Store 143
B.4.103 FSTCW: Store Floating–Point Control Word 143
B.4.104 FSTENV: Store Floating–Point Environment 143
B.4.105 FSTSW: Store Floating–Point Status Word 143
B.4.106 FSUB, FSUBP, FSUBR, FSUBRP: Floating–Point Subtract 143
B.4.107 FTST: Test ST0 Against Zero 144
B.4.108 FUCOMxx: Floating–Point Unordered Compare 144
B.4.109 FXAM: Examine Class of Value in ST0 145
B.4.110 FXCH: Floating–Point Exchange 145
B.4.111 FXRSTOR: Restore FP, MMX and SSE State 145
B.4.112 FXSAVE: Store FP, MMX and SSE State 145
B.4.113 FXTRACT: Extract Exponent and Significand 145
B.4.114 FYL2X, FYL2XP1: Compute Y times Log2(X) or Log2(X+1). 146
B.4.115 HLT: Halt Processor. 146
B.4.116 IBTS: Insert Bit String 146

B.4.117 IDIV: Signed Integer Divide . . . 146
B.4.118 IMUL: Signed Integer Multiply . . . 147
B.4.119 IN: Input from I/O Port . . . 147
B.4.120 INC: Increment Integer . . . 147
B.4.121 INSB, INSW, INSD: Input String from I/O Port . . . 148
B.4.122 INT: Software Interrupt . . . 148
B.4.123 INT3, INT1, ICEBP, INT01: Breakpoints . . . 148
B.4.124 INTO: Interrupt if Overflow . . . 149
B.4.125 INVD: Invalidate Internal Caches . . . 149
B.4.126 INVLPG: Invalidate TLB Entry . . . 149
B.4.127 IRET, IRETW, IRETD: Return from Interrupt . . . 149
B.4.128 Jcc: Conditional Branch . . . 149
B.4.129 JCXZ, JECXZ: Jump if CX/ECX Zero . . . 149
B.4.130 JMP: Jump . . . 150
B.4.131 LAHF: Load AH from Flags . . . 150
B.4.132 LAR: Load Access Rights . . . 150
B.4.133 LDMXCSR: Load Streaming SIMD Extension Control/Status . . . 151
B.4.134 LDS, LES, LFS, LGS, LSS: Load Far Pointer . . . 151
B.4.135 LEA: Load Effective Address . . . 151
B.4.136 LEAVE: Destroy Stack Frame . . . 151
B.4.137 LFENCE: Load Fence . . . 152
B.4.138 LGDT, LIDT, LLDT: Load Descriptor Tables . . . 152
B.4.139 LMSW: Load/Store Machine Status Word . . . 152
B.4.140 LOADALL, LOADALL286: Load Processor State . . . 152
B.4.141 LODSB, LODSW, LODSD: Load from String . . . 153
B.4.142 LOOP, LOOPE, LOOPZ, LOOPNE, LOOPNZ: Loop with Counter . . . 153
B.4.143 LSL: Load Segment Limit . . . 153
B.4.144 LTR: Load Task Register . . . 154
B.4.145 MASKMOVDQU: Byte Mask Write . . . 154
B.4.146 MASKMOVQ: Byte Mask Write . . . 154
B.4.147 MAXPD: Return Packed Double−Precision FP Maximum . . . 154
B.4.148 MAXPS: Return Packed Single−Precision FP Maximum . . . 154
B.4.149 MAXSD: Return Scalar Double−Precision FP Maximum . . . 154
B.4.150 MAXSS: Return Scalar Single−Precision FP Maximum . . . 154
B.4.151 MFENCE: Memory Fence . . . 155

B.4.152 MINPD: Return Packed Double−Precision FP Minimum 155
B.4.153 MINPS: Return Packed Single−Precision FP Minimum 155
B.4.154 MINSD: Return Scalar Double−Precision FP Minimum 155
B.4.155 MINSS: Return Scalar Single−Precision FP Minimum 155
B.4.156 MOV: Move Data 156
B.4.157 MOVAPD: Move Aligned Packed Double−Precision FP Values. 156
B.4.158 MOVAPS: Move Aligned Packed Single−Precision FP Values 157
B.4.159 MOVD: Move Doubleword to/from MMX Register 157
B.4.160 MOVDQ2Q: Move Quadword from XMM to MMX register. 157
B.4.161 MOVDQA: Move Aligned Double Quadword. 157
B.4.162 MOVDQU: Move Unaligned Double Quadword 157
B.4.163 MOVHLPS: Move Packed Single−Precision FP High to Low 157
B.4.164 MOVHPD: Move High Packed Double−Precision FP 158
B.4.165 MOVHPS: Move High Packed Single−Precision FP 158
B.4.166 MOVLHPS: Move Packed Single−Precision FP Low to High 158
B.4.167 MOVLPD: Move Low Packed Double−Precision FP. 158
B.4.168 MOVLPS: Move Low Packed Single−Precision FP 159
B.4.169 MOVMSKPD: Extract Packed Double−Precision FP Sign Mask 159
B.4.170 MOVMSKPS: Extract Packed Single−Precision FP Sign Mask 159
B.4.171 MOVNTDQ: Move Double Quadword Non Temporal 159
B.4.172 MOVNTI: Move Doubleword Non Temporal 159
B.4.173 MOVNTPD: Move Aligned Four Packed Single−Precision FP Values Non Temporal 159
B.4.174 MOVNTPS: Move Aligned Four Packed Single−Precision FP Values Non Temporal 160
B.4.175 MOVNTQ: Move Quadword Non Temporal 160
B.4.176 MOVQ: Move Quadword to/from MMX Register 160
B.4.177 MOVQ2DQ: Move Quadword from MMX to XMM register. 160
B.4.178 MOVSB, MOVSW, MOVSD: Move String 160
B.4.179 MOVSD: Move Scalar Double−Precision FP Value 160
B.4.180 MOVSS: Move Scalar Single−Precision FP Value 161
B.4.181 MOVSX, MOVZX: Move Data with Sign or Zero Extend. 161
B.4.182 MOVUPD: Move Unaligned Packed Double−Precision FP Values 161
B.4.183 MOVUPS: Move Unaligned Packed Single−Precision FP Values 161
B.4.184 MUL: Unsigned Integer Multiply 161

B.4.185 MULPD: Packed Single−FP Multiply 162
B.4.186 MULPS: Packed Single−FP Multiply 162
B.4.187 MULSD: Scalar Single−FP Multiply 162
B.4.188 MULSS: Scalar Single−FP Multiply 162
B.4.189 NEG, NOT: Two's and One's Complement. 162
B.4.190 NOP: No Operation 162
B.4.191 OR: Bitwise OR 162
B.4.192 ORPD: Bit−wise Logical OR of Double−Precision FP Data 163
B.4.193 ORPS: Bit−wise Logical OR of Single−Precision FP Data 163
B.4.194 OUT: Output Data to I/O Port. 163
B.4.195 OUTSB, OUTSW, OUTSD: Output String to I/O Port. 163
B.4.196 PACKSSDW, PACKSSWB, PACKUSWB: Pack Data 164
B.4.197 PADDB, PADDW, PADDD: Add Packed Integers 164
B.4.198 PADDQ: Add Packed Quadword Integers 165
B.4.199 PADDSB, PADDSW: Add Packed Signed Integers With Saturation 165
B.4.200 PADDSIW: MMX Packed Addition to Implicit Destination. 165
B.4.201 PADDUSB, PADDUSW: Add Packed Unsigned Integers With Saturation. 165
B.4.202 PAND, PANDN: MMX Bitwise AND and AND−NOT 165
B.4.203 PAUSE: Spin Loop Hint 166
B.4.204 PAVEB: MMX Packed Average 166
B.4.205 PAVGB PAVGW: Average Packed Integers 166
B.4.206 PAVGUSB: Average of unsigned packed 8−bit values 166
B.4.207 PCMPxx: Compare Packed Integers.. 166
B.4.208 PDISTIB: MMX Packed Distance and Accumulate with Implied Register 167
B.4.209 PEXTRW: Extract Word 167
B.4.210 PF2ID: Packed Single−Precision FP to Integer Convert 168
B.4.211 PF2IW: Packed Single−Precision FP to Integer Word Convert 168
B.4.212 PFACC: Packed Single−Precision FP Accumulate 168
B.4.213 PFADD: Packed Single−Precision FP Addition 168
B.4.214 PFCMPxx: Packed Single−Precision FP Compare 168
B.4.215 PFMAX: Packed Single−Precision FP Maximum 169
B.4.216 PFMIN: Packed Single−Precision FP Minimum. 169
B.4.217 PFMUL: Packed Single−Precision FP Multiply 169
B.4.218 PFNACC: Packed Single−Precision FP Negative Accumulate 169
B.4.219 PFPNACC: Packed Single−Precision FP Mixed Accumulate 169

B.4.220 PFRCP: Packed Single−Precision FP Reciprocal Approximation 169
B.4.221 PFRCPIT1: Packed Single−Precision FP Reciprocal, First Iteration Step. 170
B.4.222 PFRCPIT2: Packed Single−Precision FP Reciprocal/ Reciprocal Square Root, Second Iteration Step 170
B.4.223 PFRSQIT1: Packed Single−Precision FP Reciprocal Square Root, First Iteration Step 170
B.4.224 PFRSQRT: Packed Single−Precision FP Reciprocal Square Root Approximation . 170
B.4.225 PFSUB: Packed Single−Precision FP Subtract 170
B.4.226 PFSUBR: Packed Single−Precision FP Reverse Subtract 170
B.4.227 PI2FD: Packed Doubleword Integer to Single−Precision FP Convert 171
B.4.228 PF2IW: Packed Word Integer to Single−Precision FP Convert 171
B.4.229 PINSRW: Insert Word. 171
B.4.230 PMACHRIW: Packed Multiply and Accumulate with Rounding 171
B.4.231 PMADDWD: MMX Packed Multiply and Add 171
B.4.232 PMAGW: MMX Packed Magnitude 172
B.4.233 PMAXSW: Packed Signed Integer Word Maximum 172
B.4.234 PMAXUB: Packed Unsigned Integer Byte Maximum 172
B.4.235 PMINSW: Packed Signed Integer Word Minimum 172
B.4.236 PMINUB: Packed Unsigned Integer Byte Minimum. 172
B.4.237 PMOVMSKB: Move Byte Mask To Integer 172
B.4.238 PMULHRWC, PMULHRIW: Multiply Packed 16−bit Integers With Rounding, and Store High Word 173
B.4.239 PMULHRWA: Multiply Packed 16−bit Integers With Rounding, and Store High Word. 173
B.4.240 PMULHUW: Multiply Packed 16−bit Integers, and Store High Word 173
B.4.241 PMULHW, PMULLW: Multiply Packed 16−bit Integers, and Store 173
B.4.242 PMULUDQ: Multiply Packed Unsigned 32−bit Integers, and Store.. 174
B.4.243 PMVccZB: MMX Packed Conditional Move 174
B.4.244 POP: Pop Data from Stack 174
B.4.245 POPAx: Pop All General−Purpose Registers 175
B.4.246 POPFx: Pop Flags Register. 175
B.4.247 POR: MMX Bitwise OR 175
B.4.248 PREFETCH: Prefetch Data Into Caches 175
B.4.249 PREFETCHh: Prefetch Data Into Caches 176
B.4.250 PSADBW: Packed Sum of Absolute Differences 176
B.4.251 PSHUFD: Shuffle Packed Doublewords 176

B.4.252 PSHUFHW: Shuffle Packed High Words 176
B.4.253 PSHUFLW: Shuffle Packed Low Words 177
B.4.254 PSHUFW: Shuffle Packed Words 177
B.4.255 PSLLx: Packed Data Bit Shift Left Logical 177
B.4.256 PSRAx: Packed Data Bit Shift Right Arithmetic 177
B.4.257 PSRLx: Packed Data Bit Shift Right Logical 178
B.4.258 PSUBx: Subtract Packed Integers 178
B.4.259 PSUBSxx, PSUBUSx: Subtract Packed Integers With Saturation 179
B.4.260 PSUBSIW: MMX Packed Subtract with Saturation to Implied Destination 179
B.4.261 PSWAPD: Swap Packed Data 179
B.4.262 PUNPCKxxx: Unpack and Interleave Data 180
B.4.263 PUSH: Push Data on Stack 180
B.4.264 PUSHAx: Push All General−Purpose Registers 181
B.4.265 PUSHFx: Push Flags Register 181
B.4.266 PXOR: MMX Bitwise XOR. 182
B.4.267 RCL, RCR: Bitwise Rotate through Carry Bit 182
B.4.268 RCPPS: Packed Single−Precision FP Reciprocal 182
B.4.269 RCPSS: Scalar Single−Precision FP Reciprocal. 183
B.4.270 RDMSR: Read Model−Specific Registers 183
B.4.271 RDPMC: Read Performance−Monitoring Counters 183
B.4.272 RDSHR: Read SMM Header Pointer Register 183
B.4.273 RDTSC: Read Time−Stamp Counter 183
B.4.274 RET, RETF, RETN: Return from Procedure Call 183
B.4.275 ROL, ROR: Bitwise Rotate 183
B.4.276 RSDC: Restore Segment Register and Descriptor 184
B.4.277 RSLDT: Restore Segment Register and Descriptor 184
B.4.278 RSM: Resume from System−Management Mode 184
B.4.279 RSQRTPS: Packed Single−Precision FP Square Root Reciprocal 184
B.4.280 RSQRTSS: Scalar Single−Precision FP Square Root Reciprocal 184
B.4.281 RSTS: Restore TSR and Descriptor 185
B.4.282 SAHF: Store AH to Flags 185
B.4.283 SAL, SAR: Bitwise Arithmetic Shifts 185
B.4.284 SALC: Set AL from Carry Flag. 185
B.4.285 SBB: Subtract with Borrow. 186
B.4.286 SCASB, SCASW, SCASD: Scan String 186

B.4.287 SETcc: Set Register from Condition 186
B.4.288 SFENCE: Store Fence. 187
B.4.289 SGDT, SIDT, SLDT: Store Descriptor Table Pointers 187
B.4.290 SHL, SHR: Bitwise Logical Shifts 187
B.4.291 SHLD, SHRD: Bitwise Double−Precision Shifts 188
B.4.292 SHUFPD: Shuffle Packed Double−Precision FP Values 188
B.4.293 SHUFPS: Shuffle Packed Single−Precision FP Values 189
B.4.294 SMI: System Management Interrupt 189
B.4.295 SMINT, SMINTOLD: Software SMM Entry (CYRIX) 189
B.4.296 SMSW: Store Machine Status Word. 189
B.4.297 SQRTPD: Packed Double−Precision FP Square Root 189
B.4.298 SQRTPS: Packed Single−Precision FP Square Root 189
B.4.299 SQRTSD: Scalar Double−Precision FP Square Root 189
B.4.300 SQRTSS: Scalar Single−Precision FP Square Root 190
B.4.301 STC, STD, STI: Set Flags 190
B.4.302 STMXCSR: Store Streaming SIMD Extension Control/Status 190
B.4.303 STOSB, STOSW, STOSD: Store Byte to String 190
B.4.304 STR: Store Task Register 190
B.4.305 SUB: Subtract Integers 191
B.4.306 SUBPD: Packed Double−Precision FP Subtract 191
B.4.307 SUBPS: Packed Single−Precision FP Subtract 191
B.4.308 SUBSD: Scalar Single−FP Subtract 191
B.4.309 SUBSS: Scalar Single−FP Subtract 191
B.4.310 SVDC: Save Segment Register and Descriptor. 192
B.4.311 SVLDT: Save LDTR and Descriptor 192
B.4.312 SVTS: Save TSR and Descriptor 192
B.4.313 SYSCALL: Call Operating System 192
B.4.314 SYSENTER: Fast System Call 192
B.4.315 SYSEXIT: Fast Return From System Call. 193
B.4.316 SYSRET: Return From Operating System 193
B.4.317 TEST: Test Bits (notional bitwise AND) 194
B.4.318 UCOMISD: Unordered Scalar Double−Precision FP compare and set EFLAGS . . . 194
B.4.319 UCOMISS: Unordered Scalar Single−Precision FP compare and set EFLAGS . . . 194
B.4.320 UD0, UD1, UD2: Undefined Instruction 194
B.4.321 UMOV: User Move Data 195

B.4.322 UNPCKHPD: Unpack and Interleave High Packed Double−Precision FP Values . . 195
B.4.323 UNPCKHPS: Unpack and Interleave High Packed Single−Precision FP Values . . . 195
B.4.324 UNPCKLPD: Unpack and Interleave Low Packed Double−Precision FP Data 195
B.4.325 UNPCKLPS: Unpack and Interleave Low Packed Single−Precision FP Data 195
B.4.326 VERR, VERW: Verify Segment Readability/Writability 196
B.4.327 WAIT: Wait for Floating−Point Processor 196
B.4.328 WBINVD: Write Back and Invalidate Cache 196
B.4.329 WRMSR: Write Model−Specific Registers 196
B.4.330 WRSHR: Write SMM Header Pointer Register 196
B.4.331 XADD: Exchange and Add 196
B.4.332 XBTS: Extract Bit String 197
B.4.333 XCHG: Exchange 197
B.4.334 XLATB: Translate Byte in Lookup Table 197
B.4.335 XOR: Bitwise Exclusive OR 197
B.4.336 XORPD: Bitwise Logical XOR of Double−Precision FP Values 198
B.4.337 XORPS: Bitwise Logical XOR of Single−Precision FP Values 198

Chapter 1: Introduction

1.1 What Is NASM?

The Netwide Assembler, NASM, is an 80x86 assembler designed for portability and modularity. It supports a range of object file formats, including Linux and NetBSD/FreeBSD a.out, ELF, COFF, Microsoft 16–bit OBJ and Win32. It will also output plain binary files. Its syntax is designed to be simple and easy to understand, similar to Intel's but less complex. It supports Pentium, P6, MMX, 3DNow!, SSE and SSE2 opcodes, and has macro capability.

1.1.1 Why Yet Another Assembler?

The Netwide Assembler grew out of an idea on comp.lang.asm.x86 (or possibly alt.lang.asm – I forget which), which was essentially that there didn't seem to be a good *free* x86–series assembler around, and that maybe someone ought to write one.

- a86 is good, but not free, and in particular you don't get any 32–bit capability until you pay. It's DOS only, too.
- gas is free, and ports over DOS and Unix, but it's not very good, since it's designed to be a back end to gcc, which always feeds it correct code. So its error checking is minimal. Also, its syntax is horrible, from the point of view of anyone trying to actually *write* anything in it. Plus you can't write 16–bit code in it (properly).
- as86 is Minix– and Linux–specific, and (my version at least) doesn't seem to have much (or any) documentation.
- MASM isn't very good, and it's (was) expensive, and it runs only under DOS.
- TASM is better, but still strives for MASM compatibility, which means millions of directives and tons of red tape. And its syntax is essentially MASM's, with the contradictions and quirks that entails (although it sorts out some of those by means of Ideal mode). It's expensive too. And it's DOS–only.

So here, for your coding pleasure, is NASM. At present it's still in prototype stage – we don't promise that it can outperform any of these assemblers. But please, *please* send us bug reports, fixes, helpful information, and anything else you can get your hands on (and thanks to the many people who've done this already! You all know who you are), and we'll improve it out of all recognition. Again.

1.1.2 Licence Conditions

Please see the file COPYING, supplied as part of any NASM distribution archive, for the licence conditions under which you may use NASM. NASM is now under the so–called GNU Lesser General Public License, LGPL.

1.2 Contact Information

The current version of NASM (since about 0.98.08) are maintained by a team of developers, accessible through the nasm-devel mailing list (see below for the link). If you want to report a bug, please read section 10.2 first.

NASM has a WWW page at http://nasm.sourceforge.net. If it's not there, google for us!

The original authors are e–mailable as jules@dsf.org.uk and anakin@pobox.com. The latter is no longer involved in the development team.

New releases of NASM are uploaded to the official sites http://nasm.sourceforge.net and to ftp.kernel.org and ibiblio.org.

Announcements are posted to comp.lang.asm.x86, alt.lang.asm and comp.os.linux.announce

If you want information about NASM beta releases, and the current development status, please subscribe to the nasm–devel email list by registering at http://sourceforge.net/projects/nasm.

1.3 Installation

1.3.1 Installing NASM under MS–DOS or Windows

Once you've obtained the DOS archive for NASM, nasmXXX.zip (where XXX denotes the version number of NASM contained in the archive), unpack it into its own directory (for example c:\nasm).

The archive will contain four executable files: the NASM executable files nasm.exe and nasmw.exe, and the NDISASM executable files ndisasm.exe and ndisasmw.exe. In each case, the file whose name ends in w is a Win32 executable, designed to run under Windows 95 or Windows NT Intel, and the other one is a 16–bit DOS executable.

The only file NASM needs to run is its own executable, so copy (at least) one of nasm.exe and nasmw.exe to a directory on your PATH, or alternatively edit autoexec.bat to add the nasm directory to your PATH. (If you're only installing the Win32 version, you may wish to rename it to nasm.exe.)

That's it – NASM is installed. You don't need the nasm directory to be present to run NASM (unless you've added it to your PATH), so you can delete it if you need to save space; however, you may want to keep the documentation or test programs.

If you've downloaded the DOS source archive, nasmXXXs.zip, the nasm directory will also contain the full NASM source code, and a selection of Makefiles you can (hopefully) use to rebuild your copy of NASM from scratch.

Note that the source files insnsa.c, insnsd.c, insnsi.h and insnsn.c are automatically generated from the master instruction table insns.dat by a Perl script; the file macros.c is generated from standard.mac by another Perl script. Although the NASM source distribution includes these generated files, you will need to rebuild them (and hence, will need a Perl interpreter) if you change insns.dat, standard.mac or the documentation. It is possible future source distributions may not include these files at all. Ports of Perl for a variety of platforms, including DOS and Windows, are available from www.cpan.org.

1.3.2 Installing NASM under Unix

Once you've obtained the Unix source archive for NASM, nasm–X.XX.tar.gz (where X.XX denotes the version number of NASM contained in the archive), unpack it into a directory such as /usr/local/src. The archive, when unpacked, will create its own subdirectory nasm–X.XX.

NASM is an auto–configuring package: once you've unpacked it, cd to the directory it's been unpacked into and type ./configure. This shell script will find the best C compiler to use for building NASM and set up Makefiles accordingly.

Once NASM has auto–configured, you can type make to build the nasm and ndisasm binaries, and then make install to install them in /usr/local/bin and install the man pages nasm.1 and ndisasm.1 in /usr/local/man/man1. Alternatively, you can give options

such as `--prefix` to the configure script (see the file `INSTALL` for more details), or install the programs yourself.

NASM also comes with a set of utilities for handling the `RDOFF` custom object–file format, which are in the `rdoff` subdirectory of the NASM archive. You can build these with `make rdf` and install them with `make rdf_install`, if you want them.

If NASM fails to auto–configure, you may still be able to make it compile by using the fall–back Unix makefile `Makefile.unx`. Copy or rename that file to `Makefile` and try typing `make`. There is also a Makefile.unx file in the `rdoff` subdirectory.

Chapter 2: Running NASM

2.1 NASM Command–Line Syntax

To assemble a file, you issue a command of the form

```
nasm -f <format> <filename> [-o <output>]
```

For example,

```
nasm -f elf myfile.asm
```

will assemble `myfile.asm` into an `ELF` object file `myfile.o`. And

```
nasm -f bin myfile.asm -o myfile.com
```

will assemble `myfile.asm` into a raw binary file `myfile.com`.

To produce a listing file, with the hex codes output from NASM displayed on the left of the original sources, use the `-l` option to give a listing file name, for example:

```
nasm -f coff myfile.asm -l myfile.lst
```

To get further usage instructions from NASM, try typing

```
nasm -h
```

As `-hf`, this will also list the available output file formats, and what they are.

If you use Linux but aren't sure whether your system is `a.out` or `ELF`, type

```
file nasm
```

(in the directory in which you put the NASM binary when you installed it). If it says something like

```
nasm: ELF 32-bit LSB executable i386 (386 and up) Version 1
```

then your system is `ELF`, and you should use the option `-f elf` when you want NASM to produce Linux object files. If it says

```
nasm: Linux/i386 demand-paged executable (QMAGIC)
```

or something similar, your system is `a.out`, and you should use `-f aout` instead (Linux `a.out` systems have long been obsolete, and are rare these days.)

Like Unix compilers and assemblers, NASM is silent unless it goes wrong: you won't see any output at all, unless it gives error messages.

2.1.1 The `-o` Option: Specifying the Output File Name

NASM will normally choose the name of your output file for you; precisely how it does this is dependent on the object file format. For Microsoft object file formats (`obj` and `win32`), it will remove the `.asm` extension (or whatever extension you like to use – NASM doesn't care) from your source file name and substitute `.obj`. For Unix object file formats (`aout`, `coff`, `elf` and `as86`) it will substitute `.o`. For `rdf`, it will use `.rdf`, and for the `bin` format it will simply remove the extension, so that `myfile.asm` produces the output file `myfile`.

If the output file already exists, NASM will overwrite it, unless it has the same name as the input file, in which case it will give a warning and use `nasm.out` as the output file name instead.

For situations in which this behaviour is unacceptable, NASM provides the `-o` command-line option, which allows you to specify your desired output file name. You invoke `-o` by following it with the name you wish for the output file, either with or without an intervening space. For example:

```
nasm -f bin program.asm -o program.com
nasm -f bin driver.asm -odriver.sys
```

Note that this is a small o, and is different from a capital O , which is used to specify the number of optimisation passes required. See section 2.1.16.

2.1.2 The `-f` Option: Specifying the Output File Format

If you do not supply the `-f` option to NASM, it will choose an output file format for you itself. In the distribution versions of NASM, the default is always `bin`; if you've compiled your own copy of NASM, you can redefine `OF_DEFAULT` at compile time and choose what you want the default to be.

Like `-o`, the intervening space between `-f` and the output file format is optional; so `-f elf` and `-felf` are both valid.

A complete list of the available output file formats can be given by issuing the command `nasm -hf`.

2.1.3 The `-l` Option: Generating a Listing File

If you supply the `-l` option to NASM, followed (with the usual optional space) by a file name, NASM will generate a source-listing file for you, in which addresses and generated code are listed on the left, and the actual source code, with expansions of multi-line macros (except those which specifically request no expansion in source listings: see section 4.3.9) on the right. For example:

```
nasm -f elf myfile.asm -l myfile.lst
```

If a list file is selected, you may turn off listing for a section of your source with `[list -]`, and turn it back on with `[list +]`, (the default, obviously). There is no "user form" (without the brackets). This can be used to list only sections of interest, avoiding excessively long listings.

2.1.4 The `-M` Option: Generate Makefile Dependencies.

This option can be used to generate makefile dependencies on stdout. This can be redirected to a file for further processing. For example:

```
NASM -M myfile.asm > myfile.dep
```

2.1.5 The `-F` Option: Selecting a Debug Information Format

This option is used to select the format of the debug information emitted into the output file, to be used by a debugger (or *will* be). Use of this switch does *not* enable output of the selected debug info format. Use `-g`, see section 2.1.6, to enable output.

A complete list of the available debug file formats for an output format can be seen by issuing the command `nasm -f <format> -y`. (only "borland" in "-f obj", as of 0.98.35, but "watch this space") See: section 2.1.20.

This should not be confused with the "-f dbg" output format option which is not built into NASM by default. For information on how to enable it when building from the sources, see section 6.10

2.1.6 The `-g` Option: Enabling Debug Information.

This option can be used to generate debugging information in the specified format. See: section 2.1.5. Using `-g` without `-F` results in emitting debug info in the default format, if any, for the selected output format. If no debug information is currently implemented in the selected output format, `-g` is *silently ignored*.

2.1.7 The -X Option: Selecting an Error Reporting Format

This option can be used to select an error reporting format for any error messages that might be produced by NASM.

Currently, two error reporting formats may be selected. They are the -Xvc option and the -Xgnu option. The GNU format is the default and looks like this:

```
filename.asm:65: error: specific error message
```

where filename.asm is the name of the source file in which the error was detected, 65 is the source file line number on which the error was detected, error is the severity of the error (this could be warning), and specific error message is a more detailed text message which should help pinpoint the exact problem.

The other format, specified by -Xvc is the style used by Microsoft Visual C++ and some other programs. It looks like this:

```
filename.asm(65) : error: specific error message
```

where the only difference is that the line number is in parentheses instead of being delimited by colons.

See also the Visual C++ output format, section 6.3.

2.1.8 The -E Option: Send Errors to a File

Under MS-DOS it can be difficult (though there are ways) to redirect the standard-error output of a program to a file. Since NASM usually produces its warning and error messages on stderr, this can make it hard to capture the errors if (for example) you want to load them into an editor.

NASM therefore provides the -E option, taking a filename argument which causes errors to be sent to the specified files rather than standard error. Therefore you can redirect the errors into a file by typing

```
nasm -E myfile.err -f obj myfile.asm
```

2.1.9 The -s Option: Send Errors to stdout

The -s option redirects error messages to stdout rather than stderr, so it can be redirected under MS-DOS. To assemble the file myfile.asm and pipe its output to the more program, you can type:

```
nasm -s -f obj myfile.asm | more
```

See also the -E option, section 2.1.8.

2.1.10 The -i Option: Include File Search Directories

When NASM sees the %include or incbin directive in a source file (see section 4.6 or section 3.2.3), it will search for the given file not only in the current directory, but also in any directories specified on the command line by the use of the -i option. Therefore you can include files from a macro library, for example, by typing

```
nasm -ic:\\macrolib\\ -f obj myfile.asm
```

(As usual, a space between -i and the path name is allowed, and optional).

NASM, in the interests of complete source-code portability, does not understand the file naming conventions of the OS it is running on; the string you provide as an argument to the -i option will be prepended exactly as written to the name of the include file. Therefore the trailing backslash in the above example is necessary. Under Unix, a trailing forward slash is similarly necessary.

(You can use this to your advantage, if you're really perverse, by noting that the option `-ifoo` will cause `%include "bar.i"` to search for the file `foobar.i`...)

If you want to define a *standard* include search path, similar to `/usr/include` on Unix systems, you should place one or more `-i` directives in the `NASMENV` environment variable (see section 2.1.22).

For Makefile compatibility with many C compilers, this option can also be specified as `-I`.

2.1.11 The `-p` Option: Pre-Include a File

NASM allows you to specify files to be *pre-included* into your source file, by the use of the `-p` option. So running

```
nasm myfile.asm -p myinc.inc
```

is equivalent to running `nasm myfile.asm` and placing the directive `%include "myinc.inc"` at the start of the file.

For consistency with the `-I`, `-D` and `-U` options, this option can also be specified as `-P`.

2.1.12 The `-d` Option: Pre-Define a Macro

Just as the `-p` option gives an alternative to placing `%include` directives at the start of a source file, the `-d` option gives an alternative to placing a `%define` directive. You could code

```
nasm myfile.asm -dFOO=100
```

as an alternative to placing the directive

```
%define FOO 100
```

at the start of the file. You can miss off the macro value, as well: the option `-dFOO` is equivalent to coding `%define FOO`. This form of the directive may be useful for selecting assembly-time options which are then tested using `%ifdef`, for example `-dDEBUG`.

For Makefile compatibility with many C compilers, this option can also be specified as `-D`.

2.1.13 The `-u` Option: Undefine a Macro

The `-u` option undefines a macro that would otherwise have been pre-defined, either automatically or by a `-p` or `-d` option specified earlier on the command lines.

For example, the following command line:

```
nasm myfile.asm -dFOO=100 -uFOO
```

would result in `FOO` *not* being a predefined macro in the program. This is useful to override options specified at a different point in a Makefile.

For Makefile compatibility with many C compilers, this option can also be specified as `-U`.

2.1.14 The `-e` Option: Preprocess Only

NASM allows the preprocessor to be run on its own, up to a point. Using the `-e` option (which requires no arguments) will cause NASM to preprocess its input file, expand all the macro references, remove all the comments and preprocessor directives, and print the resulting file on standard output (or save it to a file, if the `-o` option is also used).

This option cannot be applied to programs which require the preprocessor to evaluate expressions which depend on the values of symbols: so code such as

```
%assign tablesize ($-tablestart)
```

will cause an error in preprocess–only mode.

2.1.15 The `-a` Option: Don't Preprocess At All

If NASM is being used as the back end to a compiler, it might be desirable to suppress preprocessing completely and assume the compiler has already done it, to save time and increase compilation speeds. The `-a` option, requiring no argument, instructs NASM to replace its powerful preprocessor with a stub preprocessor which does nothing.

2.1.16 The `-On` Option: Specifying Multipass Optimization.

NASM defaults to being a two pass assembler. This means that if you have a complex source file which needs more than 2 passes to assemble optimally, you have to enable extra passes.

Using the `-O` option, you can tell NASM to carry out multiple passes. The syntax is:

- `-O0` strict two–pass assembly, JMP and Jcc are handled more like v0.98, except that backward JMPs are short, if possible. Immediate operands take their long forms if a short form is not specified.
- `-O1` strict two–pass assembly, but forward branches are assembled with code guaranteed to reach; may produce larger code than –O0, but will produce successful assembly more often if branch offset sizes are not specified. Additionally, immediate operands which will fit in a signed byte are optimised, unless the long form is specified.
- `-On` multi–pass optimization, minimize branch offsets; also will minimize signed immediate bytes, overriding size specification unless the `strict` keyword has been used (see section 3.7). The number specifies the maximum number of passes. The more passes, the better the code, but the slower is the assembly.

Note that this is a capital O, and is different from a small o, which is used to specify the output format. See section 2.1.1.

2.1.17 The `-t` option: Enable TASM Compatibility Mode

NASM includes a limited form of compatibility with Borland's `TASM`. When NASM's `-t` option is used, the following changes are made:

- local labels may be prefixed with `@@` instead of `.`
- TASM–style response files beginning with `@` may be specified on the command line. This is different from the `-@resp` style that NASM natively supports.
- size override is supported within brackets. In TASM compatible mode, a size override inside square brackets changes the size of the operand, and not the address type of the operand as it does in NASM syntax. E.g. `mov eax,[DWORD val]` is valid syntax in TASM compatibility mode. Note that you lose the ability to override the default address type for the instruction.
- `%arg` preprocessor directive is supported which is similar to TASM's `ARG` directive.
- `%local` preprocessor directive
- `%stacksize` preprocessor directive
- unprefixed forms of some directives supported (`arg`, `elif`, `else`, `endif`, `if`, `ifdef`, `ifdifi`, `ifndef`, `include`, `local`)
- more...

For more information on the directives, see the section on TASM Compatiblity preprocessor directives in section 4.9.

2.1.18 The `-w` Option: Enable or Disable Assembly Warnings

NASM can observe many conditions during the course of assembly which are worth mentioning to the user, but not a sufficiently severe error to justify NASM refusing to generate an output file. These conditions are reported like errors, but come up with the word 'warning' before the message. Warnings do not prevent NASM from generating an output file and returning a success status to the operating system.

Some conditions are even less severe than that: they are only sometimes worth mentioning to the user. Therefore NASM supports the `-w` command-line option, which enables or disables certain classes of assembly warning. Such warning classes are described by a name, for example `orphan-labels`; you can enable warnings of this class by the command-line option `-w+orphan-labels` and disable it by `-w-orphan-labels`.

The suppressible warning classes are:

- `macro-params` covers warnings about multi-line macros being invoked with the wrong number of parameters. This warning class is enabled by default; see section 4.3.1 for an example of why you might want to disable it.
- `macro-selfref` warns if a macro references itself. This warning class is enabled by default.
- `orphan-labels` covers warnings about source lines which contain no instruction but define a label without a trailing colon. NASM does not warn about this somewhat obscure condition by default; see section 3.1 for an example of why you might want it to.
- `number-overflow` covers warnings about numeric constants which don't fit in 32 bits (for example, it's easy to type one too many Fs and produce `0x7fffffff` by mistake). This warning class is enabled by default.
- `gnu-elf-extensions` warns if 8-bit or 16-bit relocations are used in `-f elf` format. The GNU extensions allow this. This warning class is enabled by default.
- In addition, warning classes may be enabled or disabled across sections of source code with `[warning +warning-name]` or `[warning -warning-name]`. No "user form" (without the brackets) exists.

2.1.19 The `-v` Option: Display Version Info

Typing `NASM -v` will display the version of NASM which you are using, and the date on which it was compiled. This replaces the deprecated `-r`.

You will need the version number if you report a bug.

2.1.20 The `-y` Option: Display Available Debug Info Formats

Typing `nasm -f <option> -y` will display a list of the available debug info formats for the given output format. The default format is indicated by an asterisk. E.g. `nasm -f obj -y` yields `* borland`. (as of 0.98.35, the *only* debug info format implemented).

2.1.21 The `--prefix` and `--postfix` Options.

The `--prefix` and `--postfix` options prepend or append (respectively) the given argument to all `global` or `extern` variables. E.g. `--prefix_` will prepend the underscore to all global and external variables, as C sometimes (but not always) likes it.

2.1.22 The `NASMENV` Environment Variable

If you define an environment variable called `NASMENV`, the program will interpret it as a list of extra command-line options, which are processed before the real command line. You can use this to define standard search directories for include files, by putting `-i` options in the `NASMENV` variable.

The value of the variable is split up at white space, so that the value `-s -ic:\nasmlib` will be treated as two separate options. However, that means that the value `-dNAME="my name"` won't do what you might want, because it will be split at the space and the NASM command-line processing will get confused by the two nonsensical words `-dNAME="my` and `name"`.

To get round this, NASM provides a feature whereby, if you begin the `NASMENV` environment variable with some character that isn't a minus sign, then NASM will treat this character as the separator character for options. So setting the `NASMENV` variable to the value `!-s!-ic:\nasmlib` is equivalent to setting it to `-s -ic:\nasmlib`, but `!-dNAME="my name"` will work.

This environment variable was previously called `NASM`. This was changed with version 0.98.31.

2.2 Quick Start for MASM Users

If you're used to writing programs with MASM, or with TASM in MASM-compatible (non-Ideal) mode, or with `a86`, this section attempts to outline the major differences between MASM's syntax and NASM's. If you're not already used to MASM, it's probably worth skipping this section.

2.2.1 NASM Is Case-Sensitive

One simple difference is that NASM is case-sensitive. It makes a difference whether you call your label `foo`, `Foo` or `FOO`. If you're assembling to `DOS` or `OS/2 .OBJ` files, you can invoke the `UPPERCASE` directive (documented in section 6.2) to ensure that all symbols exported to other code modules are forced to be upper case; but even then, *within* a single module, NASM will distinguish between labels differing only in case.

2.2.2 NASM Requires Square Brackets For Memory References

NASM was designed with simplicity of syntax in mind. One of the design goals of NASM is that it should be possible, as far as is practical, for the user to look at a single line of NASM code and tell what opcode is generated by it. You can't do this in MASM: if you declare, for example,

```
foo     equ     1
bar     dw      2
```

then the two lines of code

```
        mov     ax,foo
        mov     ax,bar
```

generate completely different opcodes, despite having identical-looking syntaxes.

NASM avoids this undesirable situation by having a much simpler syntax for memory references. The rule is simply that any access to the *contents* of a memory location requires square brackets around the address, and any access to the *address* of a variable doesn't. So an instruction of the form `mov ax,foo` will *always* refer to a compile-time constant, whether it's an `EQU` or the address of a variable; and to access the *contents* of the variable `bar`, you must code `mov ax,[bar]`.

This also means that NASM has no need for MASM's `OFFSET` keyword, since the MASM code `mov ax,offset bar` means exactly the same thing as NASM's `mov ax,bar`. If you're trying to get large amounts of MASM code to assemble sensibly under NASM, you can always code `%idefine offset` to make the preprocessor treat the `OFFSET` keyword as a no-op.

This issue is even more confusing in `a86`, where declaring a label with a trailing colon defines it to be a 'label' as opposed to a 'variable' and causes `a86` to adopt NASM-style semantics; so in `a86`, `mov ax,var` has different behaviour depending on whether `var` was declared as `var: dw 0` (a label) or `var dw 0` (a word-size variable). NASM is very simple by comparison: *everything* is a label.

NASM, in the interests of simplicity, also does not support the hybrid syntaxes supported by MASM and its clones, such as `mov ax,table[bx]`, where a memory reference is denoted by one portion outside square brackets and another portion inside. The correct syntax for the above is `mov ax,[table+bx]`. Likewise, `mov ax,es:[di]` is wrong and `mov ax,[es:di]` is right.

2.2.3 NASM Doesn't Store Variable Types

NASM, by design, chooses not to remember the types of variables you declare. Whereas MASM will remember, on seeing `var dw 0`, that you declared `var` as a word–size variable, and will then be able to fill in the ambiguity in the size of the instruction `mov var,2`, NASM will deliberately remember nothing about the symbol `var` except where it begins, and so you must explicitly code `mov word [var],2`.

For this reason, NASM doesn't support the `LODS`, `MOVS`, `STOS`, `SCAS`, `CMPS`, `INS`, or `OUTS` instructions, but only supports the forms such as `LODSB`, `MOVSW`, and `SCASD`, which explicitly specify the size of the components of the strings being manipulated.

2.2.4 NASM Doesn't `ASSUME`

As part of NASM's drive for simplicity, it also does not support the `ASSUME` directive. NASM will not keep track of what values you choose to put in your segment registers, and will never *automatically* generate a segment override prefix.

2.2.5 NASM Doesn't Support Memory Models

NASM also does not have any directives to support different 16–bit memory models. The programmer has to keep track of which functions are supposed to be called with a far call and which with a near call, and is responsible for putting the correct form of `RET` instruction (`RETN` or `RETF`; NASM accepts `RET` itself as an alternate form for `RETN`); in addition, the programmer is responsible for coding CALL FAR instructions where necessary when calling *external* functions, and must also keep track of which external variable definitions are far and which are near.

2.2.6 Floating–Point Differences

NASM uses different names to refer to floating–point registers from MASM: where MASM would call them `ST(0)`, `ST(1)` and so on, and `a86` would call them simply `0`, `1` and so on, NASM chooses to call them `st0`, `st1` etc.

As of version 0.96, NASM now treats the instructions with 'nowait' forms in the same way as MASM–compatible assemblers. The idiosyncratic treatment employed by 0.95 and earlier was based on a misunderstanding by the authors.

2.2.7 Other Differences

For historical reasons, NASM uses the keyword `TWORD` where MASM and compatible assemblers use `TBYTE`.

NASM does not declare uninitialised storage in the same way as MASM: where a MASM programmer might use `stack db 64 dup (?)`, NASM requires `stack resb 64`, intended to be read as 'reserve 64 bytes'. For a limited amount of compatibility, since NASM treats `?` as a valid character in symbol names, you can code `? equ 0` and then writing `dw ?` will at least do something vaguely useful. `DUP` is still not a supported syntax, however.

In addition to all of this, macros and directives work completely differently to MASM. See chapter 4 and chapter 5 for further details.

Chapter 3: The NASM Language

3.1 Layout of a NASM Source Line

Like most assemblers, each NASM source line contains (unless it is a macro, a preprocessor directive or an assembler directive: see chapter 4 and chapter 5) some combination of the four fields

```
label:    instruction operands        ; comment
```

As usual, most of these fields are optional; the presence or absence of any combination of a label, an instruction and a comment is allowed. Of course, the operand field is either required or forbidden by the presence and nature of the instruction field.

NASM uses backslash (\) as the line continuation character; if a line ends with backslash, the next line is considered to be a part of the backslash–ended line.

NASM places no restrictions on white space within a line: labels may have white space before them, or instructions may have no space before them, or anything. The colon after a label is also optional. (Note that this means that if you intend to code `lodsb` alone on a line, and type `lodab` by accident, then that's still a valid source line which does nothing but define a label. Running NASM with the command–line option `-w+orphan-labels` will cause it to warn you if you define a label alone on a line without a trailing colon.)

Valid characters in labels are letters, numbers, _, `$`, `#`, `@`, `~`, `.`, and `?`. The only characters which may be used as the *first* character of an identifier are letters, `.` (with special meaning: see section 3.9), _ and `?`. An identifier may also be prefixed with a `$` to indicate that it is intended to be read as an identifier and not a reserved word; thus, if some other module you are linking with defines a symbol called `eax`, you can refer to `$eax` in NASM code to distinguish the symbol from the register. Maximum length of an identifier is 4095 characters.

The instruction field may contain any machine instruction: Pentium and P6 instructions, FPU instructions, MMX instructions and even undocumented instructions are all supported. The instruction may be prefixed by `LOCK`, `REP`, `REPE/REPZ` or `REPNE/REPNZ`, in the usual way. Explicit address–size and operand–size prefixes `A16`, `A32`, `O16` and `O32` are provided – one example of their use is given in chapter 9. You can also use the name of a segment register as an instruction prefix: coding `es mov [bx],ax` is equivalent to coding `mov [es:bx],ax`. We recommend the latter syntax, since it is consistent with other syntactic features of the language, but for instructions such as `LODSB`, which has no operands and yet can require a segment override, there is no clean syntactic way to proceed apart from `es lodsb`.

An instruction is not required to use a prefix: prefixes such as `CS`, `A32`, `LOCK` or `REPE` can appear on a line by themselves, and NASM will just generate the prefix bytes.

In addition to actual machine instructions, NASM also supports a number of pseudo–instructions, described in section 3.2.

Instruction operands may take a number of forms: they can be registers, described simply by the register name (e.g. `ax`, `bp`, `ebx`, `cr0`: NASM does not use the `gas`–style syntax in which register names must be prefixed by a `%` sign), or they can be effective addresses (see section 3.3), constants (section 3.4) or expressions (section 3.5).

For floating–point instructions, NASM accepts a wide range of syntaxes: you can use two–operand forms like MASM supports, or you can use NASM's native single–operand forms in most cases.

Details of all forms of each supported instruction are given in appendix B. For example, you can code:

```
fadd    st1             ; this sets st0 := st0 + st1
fadd    st0,st1         ; so does this

fadd    st1,st0         ; this sets st1 := st1 + st0
fadd    to st1          ; so does this
```

Almost any floating−point instruction that references memory must use one of the prefixes `DWORD`, `QWORD` or `TWORD` to indicate what size of memory operand it refers to.

3.2 Pseudo−Instructions

Pseudo−instructions are things which, though not real x86 machine instructions, are used in the instruction field anyway because that's the most convenient place to put them. The current pseudo−instructions are `DB`, `DW`, `DD`, `DQ` and `DT`, their uninitialised counterparts `RESB`, `RESW`, `RESD`, `RESQ` and `REST`, the `INCBIN` command, the `EQU` command, and the `TIMES` prefix.

3.2.1 `DB` and friends: Declaring Initialised Data

`DB`, `DW`, `DD`, `DQ` and `DT` are used, much as in MASM, to declare initialised data in the output file. They can be invoked in a wide range of ways:

```
db    0x55                ; just the byte 0x55
db    0x55,0x56,0x57      ; three bytes in succession
db    'a',0x55            ; character constants are OK
db    'hello',13,10,'$'   ; so are string constants
dw    0x1234              ; 0x34 0x12
dw    'a'                 ; 0x61 0x00 (it's just a number)
dw    'ab'                ; 0x61 0x62 (character constant)
dw    'abc'               ; 0x61 0x62 0x63 0x00 (string)
dd    0x12345678          ; 0x78 0x56 0x34 0x12
dd    1.234567e20         ; floating-point constant
dq    1.234567e20         ; double-precision float
dt    1.234567e20         ; extended-precision float
```

`DQ` and `DT` do not accept numeric constants or string constants as operands.

3.2.2 `RESB` and friends: Declaring Uninitialised Data

`RESB`, `RESW`, `RESD`, `RESQ` and `REST` are designed to be used in the BSS section of a module: they declare *uninitialised* storage space. Each takes a single operand, which is the number of bytes, words, doublewords or whatever to reserve. As stated in section 2.2.7, NASM does not support the MASM/TASM syntax of reserving uninitialised space by writing `DW ?` or similar things: this is what it does instead. The operand to a `RESB`−type pseudo−instruction is a *critical expression*: see section 3.8.

For example:

```
buffer:         resb    64              ; reserve 64 bytes
wordvar:        resw    1               ; reserve a word
realarray       resq    10              ; array of ten reals
```

3.2.3 `INCBIN`: Including External Binary Files

`INCBIN` is borrowed from the old Amiga assembler DevPac: it includes a binary file verbatim into the output file. This can be handy for (for example) including graphics and sound data directly into a game executable file. It can be called in one of these three ways:

```
incbin  "file.dat"              ; include the whole file
incbin  "file.dat",1024         ; skip the first 1024 bytes
incbin  "file.dat",1024,512     ; skip the first 1024, and
                                ; actually include at most 512
```

3.2.4 EQU: Defining Constants

EQU defines a symbol to a given constant value: when EQU is used, the source line must contain a label. The action of EQU is to define the given label name to the value of its (only) operand. This definition is absolute, and cannot change later. So, for example,

```
message         db      'hello, world'
msglen          equ     $-message
```

defines msglen to be the constant 12. msglen may not then be redefined later. This is not a preprocessor definition either: the value of msglen is evaluated *once*, using the value of $ (see section 3.5 for an explanation of $) at the point of definition, rather than being evaluated wherever it is referenced and using the value of $ at the point of reference. Note that the operand to an EQU is also a critical expression (section 3.8).

3.2.5 TIMES: Repeating Instructions or Data

The TIMES prefix causes the instruction to be assembled multiple times. This is partly present as NASM's equivalent of the DUP syntax supported by MASM–compatible assemblers, in that you can code

```
zerobuf:        times 64 db 0
```

or similar things; but TIMES is more versatile than that. The argument to TIMES is not just a numeric constant, but a numeric *expression*, so you can do things like

```
buffer: db      'hello, world'
        times 64-$+buffer db ' '
```

which will store exactly enough spaces to make the total length of buffer up to 64. Finally, TIMES can be applied to ordinary instructions, so you can code trivial unrolled loops in it:

```
        times 100 movsb
```

Note that there is no effective difference between times 100 resb 1 and resb 100, except that the latter will be assembled about 100 times faster due to the internal structure of the assembler.

The operand to TIMES, like that of EQU and those of RESB and friends, is a critical expression (section 3.8).

Note also that TIMES can't be applied to macros: the reason for this is that TIMES is processed after the macro phase, which allows the argument to TIMES to contain expressions such as 64-$+buffer as above. To repeat more than one line of code, or a complex macro, use the preprocessor %rep directive.

3.3 Effective Addresses

An effective address is any operand to an instruction which references memory. Effective addresses, in NASM, have a very simple syntax: they consist of an expression evaluating to the desired address, enclosed in square brackets. For example:

```
wordvar dw      123
        mov     ax,[wordvar]
        mov     ax,[wordvar+1]
        mov     ax,[es:wordvar+bx]
```

Anything not conforming to this simple system is not a valid memory reference in NASM, for example `es:wordvar[bx]`.

More complicated effective addresses, such as those involving more than one register, work in exactly the same way:

```
mov     eax,[ebx*2+ecx+offset]
mov     ax,[bp+di+8]
```

NASM is capable of doing algebra on these effective addresses, so that things which don't necessarily *look* legal are perfectly all right:

```
mov     eax,[ebx*5]             ; assembles as [ebx*4+ebx]
mov     eax,[label1*2-label2]   ; ie [label1+(label1-label2)]
```

Some forms of effective address have more than one assembled form; in most such cases NASM will generate the smallest form it can. For example, there are distinct assembled forms for the 32–bit effective addresses `[eax*2+0]` and `[eax+eax]`, and NASM will generally generate the latter on the grounds that the former requires four bytes to store a zero offset.

NASM has a hinting mechanism which will cause `[eax+ebx]` and `[ebx+eax]` to generate different opcodes; this is occasionally useful because `[esi+ebp]` and `[ebp+esi]` have different default segment registers.

However, you can force NASM to generate an effective address in a particular form by the use of the keywords `BYTE`, `WORD`, `DWORD` and `NOSPLIT`. If you need `[eax+3]` to be assembled using a double–word offset field instead of the one byte NASM will normally generate, you can code `[dword eax+3]`. Similarly, you can force NASM to use a byte offset for a small value which it hasn't seen on the first pass (see section 3.8 for an example of such a code fragment) by using `[byte eax+offset]`. As special cases, `[byte eax]` will code `[eax+0]` with a byte offset of zero, and `[dword eax]` will code it with a double–word offset of zero. The normal form, `[eax]`, will be coded with no offset field.

The form described in the previous paragraph is also useful if you are trying to access data in a 32–bit segment from within 16 bit code. For more information on this see the section on mixed–size addressing (section 9.2). In particular, if you need to access data with a known offset that is larger than will fit in a 16–bit value, if you don't specify that it is a dword offset, nasm will cause the high word of the offset to be lost.

Similarly, NASM will split `[eax*2]` into `[eax+eax]` because that allows the offset field to be absent and space to be saved; in fact, it will also split `[eax*2+offset]` into `[eax+eax+offset]`. You can combat this behaviour by the use of the `NOSPLIT` keyword: `[nosplit eax*2]` will force `[eax*2+0]` to be generated literally.

3.4 Constants

NASM understands four different types of constant: numeric, character, string and floating–point.

3.4.1 Numeric Constants

A numeric constant is simply a number. NASM allows you to specify numbers in a variety of number bases, in a variety of ways: you can suffix `H`, `Q` or `O`, and `B` for hex, octal and binary, or you can prefix `0x` for hex in the style of C, or you can prefix `$` for hex in the style of Borland Pascal. Note, though, that the `$` prefix does double duty as a prefix on identifiers (see section 3.1), so a hex number prefixed with a `$` sign must have a digit after the `$` rather than a letter.

Some examples:

```
mov     ax,100          ; decimal
mov     ax,0a2h         ; hex
```

```
mov     ax,$0a2         ; hex again: the 0 is required
mov     ax,0xa2         ; hex yet again
mov     ax,777q         ; octal
mov     ax,777o         ; octal again
mov     ax,10010011b    ; binary
```

3.4.2 Character Constants

A character constant consists of up to four characters enclosed in either single or double quotes. The type of quote makes no difference to NASM, except of course that surrounding the constant with single quotes allows double quotes to appear within it and vice versa.

A character constant with more than one character will be arranged with little–endian order in mind: if you code

```
mov eax,'abcd'
```

then the constant generated is not `0x61626364`, but `0x64636261`, so that if you were then to store the value into memory, it would read `abcd` rather than `dcba`. This is also the sense of character constants understood by the Pentium's `CPUID` instruction (see section B.4.34).

3.4.3 String Constants

String constants are only acceptable to some pseudo–instructions, namely the `DB` family and `INCBIN`.

A string constant looks like a character constant, only longer. It is treated as a concatenation of maximum–size character constants for the conditions. So the following are equivalent:

```
db    'hello'               ; string constant
db    'h','e','l','l','o'   ; equivalent character constants
```

And the following are also equivalent:

```
dd    'ninechars'           ; doubleword string constant
dd    'nine','char','s'     ; becomes three doublewords
db    'ninechars',0,0,0     ; and really looks like this
```

Note that when used as an operand to `db`, a constant like `'ab'` is treated as a string constant despite being short enough to be a character constant, because otherwise `db 'ab'` would have the same effect as `db 'a'`, which would be silly. Similarly, three–character or four–character constants are treated as strings when they are operands to `dw`.

3.4.4 Floating–Point Constants

Floating–point constants are acceptable only as arguments to `DD`, `DQ` and `DT`. They are expressed in the traditional form: digits, then a period, then optionally more digits, then optionally an `E` followed by an exponent. The period is mandatory, so that NASM can distinguish between `dd 1`, which declares an integer constant, and `dd 1.0` which declares a floating–point constant.

Some examples:

```
dd    1.2                     ; an easy one
dq    1.e10                   ; 10,000,000,000
dq    1.e+10                  ; synonymous with 1.e10
dq    1.e-10                  ; 0.000 000 000 1
dt    3.14159265358979323846  ; pi
```

NASM cannot do compile–time arithmetic on floating–point constants. This is because NASM is designed to be portable – although it always generates code to run on x86 processors, the assembler itself can run on any system with an ANSI C compiler. Therefore, the assembler cannot guarantee

the presence of a floating−point unit capable of handling the Intel number formats, and so for NASM to be able to do floating arithmetic it would have to include its own complete set of floating−point routines, which would significantly increase the size of the assembler for very little benefit.

3.5 Expressions

Expressions in NASM are similar in syntax to those in C.

NASM does not guarantee the size of the integers used to evaluate expressions at compile time: since NASM can compile and run on 64−bit systems quite happily, don't assume that expressions are evaluated in 32−bit registers and so try to make deliberate use of integer overflow. It might not always work. The only thing NASM will guarantee is what's guaranteed by ANSI C: you always have *at least* 32 bits to work in.

NASM supports two special tokens in expressions, allowing calculations to involve the current assembly position: the `$` and `$$` tokens. `$` evaluates to the assembly position at the beginning of the line containing the expression; so you can code an infinite loop using `JMP $`. `$$` evaluates to the beginning of the current section; so you can tell how far into the section you are by using `($−$$)`.

The arithmetic operators provided by NASM are listed here, in increasing order of precedence.

3.5.1 |: Bitwise OR Operator

The | operator gives a bitwise OR, exactly as performed by the `OR` machine instruction. Bitwise OR is the lowest−priority arithmetic operator supported by NASM.

3.5.2 ^: Bitwise XOR Operator

^ provides the bitwise XOR operation.

3.5.3 &: Bitwise AND Operator

`&` provides the bitwise AND operation.

3.5.4 << and >>: Bit Shift Operators

<< gives a bit−shift to the left, just as it does in C. So `5<<3` evaluates to 5 times 8, or 40. >> gives a bit−shift to the right; in NASM, such a shift is *always* unsigned, so that the bits shifted in from the left−hand end are filled with zero rather than a sign−extension of the previous highest bit.

3.5.5 + and −: Addition and Subtraction Operators

The + and − operators do perfectly ordinary addition and subtraction.

3.5.6 *, /, //, % and %%: Multiplication and Division

`*` is the multiplication operator. `/` and `//` are both division operators: `/` is unsigned division and `//` is signed division. Similarly, `%` and `%%` provide unsigned and signed modulo operators respectively.

NASM, like ANSI C, provides no guarantees about the sensible operation of the signed modulo operator.

Since the `%` character is used extensively by the macro preprocessor, you should ensure that both the signed and unsigned modulo operators are followed by white space wherever they appear.

3.5.7 Unary Operators: +, −, ~ and SEG

The highest−priority operators in NASM's expression grammar are those which only apply to one argument. − negates its operand, + does nothing (it's provided for symmetry with −), ~ computes the one's complement of its operand, and `SEG` provides the segment address of its operand (explained in more detail in section 3.6).

3.6 SEG and WRT

When writing large 16–bit programs, which must be split into multiple segments, it is often necessary to be able to refer to the segment part of the address of a symbol. NASM supports the SEG operator to perform this function.

The SEG operator returns the *preferred* segment base of a symbol, defined as the segment base relative to which the offset of the symbol makes sense. So the code

```
mov     ax,seg symbol
mov     es,ax
mov     bx,symbol
```

will load ES:BX with a valid pointer to the symbol symbol.

Things can be more complex than this: since 16–bit segments and groups may overlap, you might occasionally want to refer to some symbol using a different segment base from the preferred one. NASM lets you do this, by the use of the WRT (With Reference To) keyword. So you can do things like

```
mov     ax,weird_seg        ; weird_seg is a segment base
mov     es,ax
mov     bx,symbol wrt weird_seg
```

to load ES:BX with a different, but functionally equivalent, pointer to the symbol symbol.

NASM supports far (inter–segment) calls and jumps by means of the syntax call segment:offset, where segment and offset both represent immediate values. So to call a far procedure, you could code either of

```
call    (seg procedure):procedure
call    weird_seg:(procedure wrt weird_seg)
```

(The parentheses are included for clarity, to show the intended parsing of the above instructions. They are not necessary in practice.)

NASM supports the syntax call far procedure as a synonym for the first of the above usages. JMP works identically to CALL in these examples.

To declare a far pointer to a data item in a data segment, you must code

```
dw      symbol, seg symbol
```

NASM supports no convenient synonym for this, though you can always invent one using the macro processor.

3.7 STRICT: Inhibiting Optimization

When assembling with the optimizer set to level 2 or higher (see section 2.1.16), NASM will use size specifiers (BYTE, WORD, DWORD, QWORD, or TWORD), but will give them the smallest possible size. The keyword STRICT can be used to inhibit optimization and force a particular operand to be emitted in the specified size. For example, with the optimizer on, and in BITS 16 mode,

```
push dword 33
```

is encoded in three bytes 66 6A 21, whereas

```
push strict dword 33
```

is encoded in six bytes, with a full dword immediate operand 66 68 21 00 00 00.

With the optimizer off, the same code (six bytes) is generated whether the `STRICT` keyword was used or not.

3.8 Critical Expressions

A limitation of NASM is that it is a two−pass assembler; unlike TASM and others, it will always do exactly two assembly passes. Therefore it is unable to cope with source files that are complex enough to require three or more passes.

The first pass is used to determine the size of all the assembled code and data, so that the second pass, when generating all the code, knows all the symbol addresses the code refers to. So one thing NASM can't handle is code whose size depends on the value of a symbol declared after the code in question. For example,

```
        times (label-$) db 0
label:  db      'Where am I?'
```

The argument to `TIMES` in this case could equally legally evaluate to anything at all; NASM will reject this example because it cannot tell the size of the `TIMES` line when it first sees it. It will just as firmly reject the slightly paradoxical code

```
        times (label-$+1) db 0
label:  db      'NOW where am I?'
```

in which *any* value for the `TIMES` argument is by definition wrong!

NASM rejects these examples by means of a concept called a *critical expression*, which is defined to be an expression whose value is required to be computable in the first pass, and which must therefore depend only on symbols defined before it. The argument to the `TIMES` prefix is a critical expression; for the same reason, the arguments to the `RESB` family of pseudo−instructions are also critical expressions.

Critical expressions can crop up in other contexts as well: consider the following code.

```
                mov     ax,symbol1
symbol1         equ     symbol2
symbol2:
```

On the first pass, NASM cannot determine the value of `symbol1`, because `symbol1` is defined to be equal to `symbol2` which NASM hasn't seen yet. On the second pass, therefore, when it encounters the line `mov ax,symbol1`, it is unable to generate the code for it because it still doesn't know the value of `symbol1`. On the next line, it would see the `EQU` again and be able to determine the value of `symbol1`, but by then it would be too late.

NASM avoids this problem by defining the right−hand side of an `EQU` statement to be a critical expression, so the definition of `symbol1` would be rejected in the first pass.

There is a related issue involving forward references: consider this code fragment.

```
        mov     eax,[ebx+offset]
offset  equ     10
```

NASM, on pass one, must calculate the size of the instruction `mov eax,[ebx+offset]` without knowing the value of `offset`. It has no way of knowing that `offset` is small enough to fit into a one−byte offset field and that it could therefore get away with generating a shorter form of the effective−address encoding; for all it knows, in pass one, `offset` could be a symbol in the code segment, and it might need the full four−byte form. So it is forced to compute the size of the instruction to accommodate a four−byte address part. In pass two, having made this decision, it is now forced to honour it and keep the instruction large, so the code generated in this case is not as

small as it could have been. This problem can be solved by defining `offset` before using it, or by forcing byte size in the effective address by coding `[byte ebx+offset]`.

Note that use of the `-On` switch (with n>=2) makes some of the above no longer true (see section 2.1.16).

3.9 Local Labels

NASM gives special treatment to symbols beginning with a period. A label beginning with a single period is treated as a *local* label, which means that it is associated with the previous non–local label. So, for example:

```
label1  ; some code

.loop
        ; some more code

        jne     .loop
        ret

label2  ; some code

.loop
        ; some more code

        jne     .loop
        ret
```

In the above code fragment, each `JNE` instruction jumps to the line immediately before it, because the two definitions of `.loop` are kept separate by virtue of each being associated with the previous non–local label.

This form of local label handling is borrowed from the old Amiga assembler DevPac; however, NASM goes one step further, in allowing access to local labels from other parts of the code. This is achieved by means of *defining* a local label in terms of the previous non–local label: the first definition of `.loop` above is really defining a symbol called `label1.loop`, and the second defines a symbol called `label2.loop`. So, if you really needed to, you could write

```
label3  ; some more code
        ; and some more

        jmp label1.loop
```

Sometimes it is useful – in a macro, for instance – to be able to define a label which can be referenced from anywhere but which doesn't interfere with the normal local–label mechanism. Such a label can't be non–local because it would interfere with subsequent definitions of, and references to, local labels; and it can't be local because the macro that defined it wouldn't know the label's full name. NASM therefore introduces a third type of label, which is probably only useful in macro definitions: if a label begins with the special prefix `..@`, then it does nothing to the local label mechanism. So you could code

```
label1:                         ; a non-local label
.local:                         ; this is really label1.local
..@foo:                          ; this is a special symbol
label2:                         ; another non-local label
.local:                         ; this is really label2.local
```

```
        jmp     ..@@foo         ; this will jump three lines up
```

NASM has the capacity to define other special symbols beginning with a double period: for example, `..start` is used to specify the entry point in the `obj` output format (see section 6.2.6).

Chapter 4: The NASM Preprocessor

NASM contains a powerful macro processor, which supports conditional assembly, multi−level file inclusion, two forms of macro (single−line and multi−line), and a 'context stack' mechanism for extra macro power. Preprocessor directives all begin with a `%` sign.

The preprocessor collapses all lines which end with a backslash (\) character into a single line. Thus:

```
%define THIS_VERY_LONG_MACRO_NAME_IS_DEFINED_TO \\
        THIS_VALUE
```

will work like a single−line macro without the backslash−newline sequence.

4.1 Single−Line Macros

4.1.1 The Normal Way: `%define`

Single−line macros are defined using the `%define` preprocessor directive. The definitions work in a similar way to C; so you can do things like

```
%define ctrl    0x1F &
%define param(a,b) ((a)+(a)*(b))

        mov     byte [param(2,ebx)], ctrl 'D'
```

which will expand to

```
        mov     byte [(2)+(2)*(ebx)], 0x1F & 'D'
```

When the expansion of a single−line macro contains tokens which invoke another macro, the expansion is performed at invocation time, not at definition time. Thus the code

```
%define a(x)    1+b(x)
%define b(x)    2*x

        mov     ax,a(8)
```

will evaluate in the expected way to `mov ax,1+2*8`, even though the macro `b` wasn't defined at the time of definition of `a`.

Macros defined with `%define` are case sensitive: after `%define foo bar`, only `foo` will expand to `bar`: `Foo` or `FOO` will not. By using `%idefine` instead of `%define` (the 'i' stands for 'insensitive') you can define all the case variants of a macro at once, so that `%idefine foo bar` would cause `foo`, `Foo`, `FOO`, `fOO` and so on all to expand to `bar`.

There is a mechanism which detects when a macro call has occurred as a result of a previous expansion of the same macro, to guard against circular references and infinite loops. If this happens, the preprocessor will only expand the first occurrence of the macro. Hence, if you code

```
%define a(x)    1+a(x)

        mov     ax,a(3)
```

the macro `a(3)` will expand once, becoming `1+a(3)`, and will then expand no further. This behaviour can be useful: see section 8.1 for an example of its use.

You can overload single−line macros: if you write

```
%define foo(x)    1+x
%define foo(x,y) 1+x*y
```

the preprocessor will be able to handle both types of macro call, by counting the parameters you pass; so `foo(3)` will become `1+3` whereas `foo(ebx,2)` will become `1+ebx*2`. However, if you define

```
%define foo bar
```

then no other definition of `foo` will be accepted: a macro with no parameters prohibits the definition of the same name as a macro *with* parameters, and vice versa.

This doesn't prevent single−line macros being *redefined*: you can perfectly well define a macro with

```
%define foo bar
```

and then re−define it later in the same source file with

```
%define foo baz
```

Then everywhere the macro `foo` is invoked, it will be expanded according to the most recent definition. This is particularly useful when defining single−line macros with `%assign` (see section 4.1.5).

You can pre−define single−line macros using the '−d' option on the NASM command line: see section 2.1.12.

4.1.2 Enhancing %define: %xdefine

To have a reference to an embedded single−line macro resolved at the time that it is embedded, as opposed to when the calling macro is expanded, you need a different mechanism to the one offered by `%define`. The solution is to use `%xdefine`, or it's case−insensitive counterpart `%xidefine`.

Suppose you have the following code:

```
%define  isTrue  1
%define  isFalse isTrue
%define  isTrue  0

val1:    db      isFalse

%define  isTrue  1

val2:    db      isFalse
```

In this case, `val1` is equal to 0, and `val2` is equal to 1. This is because, when a single−line macro is defined using `%define`, it is expanded only when it is called. As `isFalse` expands to `isTrue`, the expansion will be the current value of `isTrue`. The first time it is called that is 0, and the second time it is 1.

If you wanted `isFalse` to expand to the value assigned to the embedded macro `isTrue` at the time that `isFalse` was defined, you need to change the above code to use `%xdefine`.

```
%xdefine isTrue  1
%xdefine isFalse isTrue
%xdefine isTrue  0

val1:    db      isFalse
```

```
%xdefine isTrue  1

val2:    db      isFalse
```

Now, each time that `isFalse` is called, it expands to 1, as that is what the embedded macro `isTrue` expanded to at the time that `isFalse` was defined.

4.1.3 Concatenating Single Line Macro Tokens: %+

Individual tokens in single line macros can be concatenated, to produce longer tokens for later processing. This can be useful if there are several similar macros that perform similar functions.

As an example, consider the following:

```
%define BDASTART 400h                ; Start of BIOS data area

struc   tBIOSDA                      ; its structure
        .COM1addr       RESW    1
        .COM2addr       RESW    1
        ; ..and so on
endstruc
```

Now, if we need to access the elements of tBIOSDA in different places, we can end up with:

```
        mov     ax,BDASTART + tBIOSDA.COM1addr
        mov     bx,BDASTART + tBIOSDA.COM2addr
```

This will become pretty ugly (and tedious) if used in many places, and can be reduced in size significantly by using the following macro:

```
; Macro to access BIOS variables by their names (from tBDA):

%define BDA(x)  BDASTART + tBIOSDA. %+ x
```

Now the above code can be written as:

```
        mov     ax,BDA(COM1addr)
        mov     bx,BDA(COM2addr)
```

Using this feature, we can simplify references to a lot of macros (and, in turn, reduce typing errors).

4.1.4 Undefining macros: %undef

Single–line macros can be removed with the `%undef` command. For example, the following sequence:

```
%define foo bar
%undef  foo

        mov     eax, foo
```

will expand to the instruction `mov eax, foo`, since after `%undef` the macro `foo` is no longer defined.

Macros that would otherwise be pre–defined can be undefined on the command–line using the '–u' option on the NASM command line: see section 2.1.13.

4.1.5 Preprocessor Variables: %assign

An alternative way to define single–line macros is by means of the `%assign` command (and its case–insensitive counterpart `%iassign`, which differs from `%assign` in exactly the same way that `%idefine` differs from `%define`).

`%assign` is used to define single–line macros which take no parameters and have a numeric value. This value can be specified in the form of an expression, and it will be evaluated once, when the `%assign` directive is processed.

Like `%define`, macros defined using `%assign` can be re–defined later, so you can do things like

```
%assign i i+1
```

to increment the numeric value of a macro.

`%assign` is useful for controlling the termination of `%rep` preprocessor loops: see section 4.5 for an example of this. Another use for `%assign` is given in section 7.4 and section 8.1.

The expression passed to `%assign` is a critical expression (see section 3.8), and must also evaluate to a pure number (rather than a relocatable reference such as a code or data address, or anything involving a register).

4.2 String Handling in Macros: `%strlen` and `%substr`

It's often useful to be able to handle strings in macros. NASM supports two simple string handling macro operators from which more complex operations can be constructed.

4.2.1 String Length: `%strlen`

The `%strlen` macro is like `%assign` macro in that it creates (or redefines) a numeric value to a macro. The difference is that with `%strlen`, the numeric value is the length of a string. An example of the use of this would be:

```
%strlen charcnt 'my string'
```

In this example, `charcnt` would receive the value 8, just as if an `%assign` had been used. In this example, `'my string'` was a literal string but it could also have been a single–line macro that expands to a string, as in the following example:

```
%define sometext 'my string'
%strlen charcnt sometext
```

As in the first case, this would result in `charcnt` being assigned the value of 8.

4.2.2 Sub–strings: `%substr`

Individual letters in strings can be extracted using `%substr`. An example of its use is probably more useful than the description:

```
%substr mychar  'xyz' 1         ; equivalent to %define mychar 'x'
%substr mychar  'xyz' 2         ; equivalent to %define mychar 'y'
%substr mychar  'xyz' 3         ; equivalent to %define mychar 'z'
```

In this example, mychar gets the value of 'y'. As with `%strlen` (see section 4.2.1), the first parameter is the single–line macro to be created and the second is the string. The third parameter specifies which character is to be selected. Note that the first index is 1, not 0 and the last index is equal to the value that `%strlen` would assign given the same string. Index values out of range result in an empty string.

4.3 Multi–Line Macros: `%macro`

Multi–line macros are much more like the type of macro seen in MASM and TASM: a multi–line macro definition in NASM looks something like this.

```
%macro  prologue 1

        push    ebp
```

```
        mov     ebp,esp
        sub     esp,%1

%endmacro
```

This defines a C–like function prologue as a macro: so you would invoke the macro with a call such as

```
myfunc:   prologue 12
```

which would expand to the three lines of code

```
myfunc: push    ebp
        mov     ebp,esp
        sub     esp,12
```

The number `1` after the macro name in the `%macro` line defines the number of parameters the macro `prologue` expects to receive. The use of `%1` inside the macro definition refers to the first parameter to the macro call. With a macro taking more than one parameter, subsequent parameters would be referred to as `%2`, `%3` and so on.

Multi–line macros, like single–line macros, are case–sensitive, unless you define them using the alternative directive `%imacro`.

If you need to pass a comma as *part* of a parameter to a multi–line macro, you can do that by enclosing the entire parameter in braces. So you could code things like

```
%macro  silly 2

    %2: db      %1

%endmacro

        silly 'a', letter_a             ; letter_a:  db 'a'
        silly 'ab', string_ab           ; string_ab: db 'ab'
        silly @\{13,10@\}, crlf             ; crlf:      db 13,10
```

4.3.1 Overloading Multi–Line Macros

As with single–line macros, multi–line macros can be overloaded by defining the same macro name several times with different numbers of parameters. This time, no exception is made for macros with no parameters at all. So you could define

```
%macro  prologue 0

        push    ebp
        mov     ebp,esp

%endmacro
```

to define an alternative form of the function prologue which allocates no local stack space.

Sometimes, however, you might want to 'overload' a machine instruction; for example, you might want to define

```
%macro  push 2

        push    %1
        push    %2
```

```
%endmacro
```

so that you could code

```
        push    ebx             ; this line is not a macro call
        push    eax,ecx         ; but this one is
```

Ordinarily, NASM will give a warning for the first of the above two lines, since `push` is now defined to be a macro, and is being invoked with a number of parameters for which no definition has been given. The correct code will still be generated, but the assembler will give a warning. This warning can be disabled by the use of the `-w-macro-params` command-line option (see section 2.1.18).

4.3.2 Macro-Local Labels

NASM allows you to define labels within a multi-line macro definition in such a way as to make them local to the macro call: so calling the same macro multiple times will use a different label each time. You do this by prefixing `%%` to the label name. So you can invent an instruction which executes a `RET` if the `Z` flag is set by doing this:

```
%macro  retz 0

        jnz     %%skip
        ret
    %%skip:

%endmacro
```

You can call this macro as many times as you want, and every time you call it NASM will make up a different 'real' name to substitute for the label `%%skip`. The names NASM invents are of the form `..@2345.skip`, where the number 2345 changes with every macro call. The `..@` prefix prevents macro-local labels from interfering with the local label mechanism, as described in section 3.9. You should avoid defining your own labels in this form (the `..@` prefix, then a number, then another period) in case they interfere with macro-local labels.

4.3.3 Greedy Macro Parameters

Occasionally it is useful to define a macro which lumps its entire command line into one parameter definition, possibly after extracting one or two smaller parameters from the front. An example might be a macro to write a text string to a file in MS-DOS, where you might want to be able to write

```
        writefile [filehandle],"hello, world",13,10
```

NASM allows you to define the last parameter of a macro to be *greedy*, meaning that if you invoke the macro with more parameters than it expects, all the spare parameters get lumped into the last defined one along with the separating commas. So if you code:

```
%macro  writefile 2+

        jmp     %%endstr
  %%str:        db      %2
  %%endstr:
        mov     dx,%%str
        mov     cx,%%endstr-%%str
        mov     bx,%1
        mov     ah,0x40
        int     0x21
```

```
%endmacro
```

then the example call to `writefile` above will work as expected: the text before the first comma, `[filehandle]`, is used as the first macro parameter and expanded when `%1` is referred to, and all the subsequent text is lumped into `%2` and placed after the `db`.

The greedy nature of the macro is indicated to NASM by the use of the + sign after the parameter count on the `%macro` line.

If you define a greedy macro, you are effectively telling NASM how it should expand the macro given *any* number of parameters from the actual number specified up to infinity; in this case, for example, NASM now knows what to do when it sees a call to `writefile` with 2, 3, 4 or more parameters. NASM will take this into account when overloading macros, and will not allow you to define another form of `writefile` taking 4 parameters (for example).

Of course, the above macro could have been implemented as a non−greedy macro, in which case the call to it would have had to look like

```
writefile [filehandle], @\{"hello, world",13,10@\}
```

NASM provides both mechanisms for putting commas in macro parameters, and you choose which one you prefer for each macro definition.

See section 5.2.1 for a better way to write the above macro.

4.3.4 Default Macro Parameters

NASM also allows you to define a multi−line macro with a *range* of allowable parameter counts. If you do this, you can specify defaults for omitted parameters. So, for example:

```
%macro  die 0-1 "Painful program death has occurred."

        writefile 2,%1
        mov     ax,0x4c01
        int     0x21

%endmacro
```

This macro (which makes use of the `writefile` macro defined in section 4.3.3) can be called with an explicit error message, which it will display on the error output stream before exiting, or it can be called with no parameters, in which case it will use the default error message supplied in the macro definition.

In general, you supply a minimum and maximum number of parameters for a macro of this type; the minimum number of parameters are then required in the macro call, and then you provide defaults for the optional ones. So if a macro definition began with the line

```
%macro foobar 1-3 eax,[ebx+2]
```

then it could be called with between one and three parameters, and `%1` would always be taken from the macro call. `%2`, if not specified by the macro call, would default to `eax`, and `%3` if not specified would default to `[ebx+2]`.

You may omit parameter defaults from the macro definition, in which case the parameter default is taken to be blank. This can be useful for macros which can take a variable number of parameters, since the `%0` token (see section 4.3.5) allows you to determine how many parameters were really passed to the macro call.

This defaulting mechanism can be combined with the greedy–parameter mechanism; so the `die` macro above could be made more powerful, and more useful, by changing the first line of the definition to

```
%macro die 0-1+ "Painful program death has occurred.",13,10
```

The maximum parameter count can be infinite, denoted by `*`. In this case, of course, it is impossible to provide a *full* set of default parameters. Examples of this usage are shown in section 4.3.6.

4.3.5 `%0`: Macro Parameter Counter

For a macro which can take a variable number of parameters, the parameter reference `%0` will return a numeric constant giving the number of parameters passed to the macro. This can be used as an argument to `%rep` (see section 4.5) in order to iterate through all the parameters of a macro. Examples are given in section 4.3.6.

4.3.6 `%rotate`: Rotating Macro Parameters

Unix shell programmers will be familiar with the `shift` shell command, which allows the arguments passed to a shell script (referenced as `$1`, `$2` and so on) to be moved left by one place, so that the argument previously referenced as `$2` becomes available as `$1`, and the argument previously referenced as `$1` is no longer available at all.

NASM provides a similar mechanism, in the form of `%rotate`. As its name suggests, it differs from the Unix `shift` in that no parameters are lost: parameters rotated off the left end of the argument list reappear on the right, and vice versa.

`%rotate` is invoked with a single numeric argument (which may be an expression). The macro parameters are rotated to the left by that many places. If the argument to `%rotate` is negative, the macro parameters are rotated to the right.

So a pair of macros to save and restore a set of registers might work as follows:

```
%macro  multipush 1-*

  %rep  %0
        push    %1
  %rotate 1
  %endrep

%endmacro
```

This macro invokes the `PUSH` instruction on each of its arguments in turn, from left to right. It begins by pushing its first argument, `%1`, then invokes `%rotate` to move all the arguments one place to the left, so that the original second argument is now available as `%1`. Repeating this procedure as many times as there were arguments (achieved by supplying `%0` as the argument to `%rep`) causes each argument in turn to be pushed.

Note also the use of `*` as the maximum parameter count, indicating that there is no upper limit on the number of parameters you may supply to the `multipush` macro.

It would be convenient, when using this macro, to have a `POP` equivalent, which *didn't* require the arguments to be given in reverse order. Ideally, you would write the `multipush` macro call, then cut–and–paste the line to where the pop needed to be done, and change the name of the called macro to `multipop`, and the macro would take care of popping the registers in the opposite order from the one in which they were pushed.

This can be done by the following definition:

```
%macro  multipop 1-*

  %rep %0
  %rotate -1
        pop     %1
  %endrep

%endmacro
```

This macro begins by rotating its arguments one place to the *right*, so that the original *last* argument appears as `%1`. This is then popped, and the arguments are rotated right again, so the second–to–last argument becomes `%1`. Thus the arguments are iterated through in reverse order.

4.3.7 Concatenating Macro Parameters

NASM can concatenate macro parameters on to other text surrounding them. This allows you to declare a family of symbols, for example, in a macro definition. If, for example, you wanted to generate a table of key codes along with offsets into the table, you could code something like

```
%macro keytab_entry 2

    keypos%1    equ     $-keytab
                db      %2

%endmacro

keytab:
          keytab_entry F1,128+1
          keytab_entry F2,128+2
          keytab_entry Return,13
```

which would expand to

```
keytab:
keyposF1        equ     $-keytab
                db     128+1
keyposF2        equ     $-keytab
                db      128+2
keyposReturn    equ     $-keytab
                db      13
```

You can just as easily concatenate text on to the other end of a macro parameter, by writing `%1foo`.

If you need to append a *digit* to a macro parameter, for example defining labels `foo1` and `foo2` when passed the parameter `foo`, you can't code `%11` because that would be taken as the eleventh macro parameter. Instead, you must code `%{1}1`, which will separate the first `1` (giving the number of the macro parameter) from the second (literal text to be concatenated to the parameter).

This concatenation can also be applied to other preprocessor in–line objects, such as macro–local labels (section 4.3.2) and context–local labels (section 4.7.2). In all cases, ambiguities in syntax can be resolved by enclosing everything after the `%` sign and before the literal text in braces: so `%{%foo}bar` concatenates the text `bar` to the end of the real name of the macro–local label `%%foo`. (This is unnecessary, since the form NASM uses for the real names of macro–local labels means that the two usages `%{%foo}bar` and `%%foobar` would both expand to the same thing anyway; nevertheless, the capability is there.)

4.3.8 Condition Codes as Macro Parameters

NASM can give special treatment to a macro parameter which contains a condition code. For a start, you can refer to the macro parameter `%1` by means of the alternative syntax `%+1`, which informs NASM that this macro parameter is supposed to contain a condition code, and will cause the preprocessor to report an error message if the macro is called with a parameter which is *not* a valid condition code.

Far more usefully, though, you can refer to the macro parameter by means of `%-1`, which NASM will expand as the *inverse* condition code. So the `retz` macro defined in section 4.3.2 can be replaced by a general conditional-return macro like this:

```
%macro  retc 1

        j%-1    %%skip
        ret
  %%skip:

%endmacro
```

This macro can now be invoked using calls like `retc ne`, which will cause the conditional-jump instruction in the macro expansion to come out as `JE`, or `retc po` which will make the jump a `JPE`.

The `%+1` macro-parameter reference is quite happy to interpret the arguments `CXZ` and `ECXZ` as valid condition codes; however, `%-1` will report an error if passed either of these, because no inverse condition code exists.

4.3.9 Disabling Listing Expansion

When NASM is generating a listing file from your program, it will generally expand multi-line macros by means of writing the macro call and then listing each line of the expansion. This allows you to see which instructions in the macro expansion are generating what code; however, for some macros this clutters the listing up unnecessarily.

NASM therefore provides the `.nolist` qualifier, which you can include in a macro definition to inhibit the expansion of the macro in the listing file. The `.nolist` qualifier comes directly after the number of parameters, like this:

```
%macro foo 1.nolist
```

Or like this:

```
%macro bar 1-5+.nolist a,b,c,d,e,f,g,h
```

4.4 Conditional Assembly

Similarly to the C preprocessor, NASM allows sections of a source file to be assembled only if certain conditions are met. The general syntax of this feature looks like this:

```
%if<condition>
    ; some code which only appears if <condition> is met
%elif<condition2>
    ; only appears if <condition> is not met but <condition2> is
%else
    ; this appears if neither <condition> nor <condition2> was met
%endif
```

The `%else` clause is optional, as is the `%elif` clause. You can have more than one `%elif` clause as well.

4.4.1 `%ifdef`: Testing Single–Line Macro Existence

Beginning a conditional–assembly block with the line `%ifdef MACRO` will assemble the subsequent code if, and only if, a single–line macro called `MACRO` is defined. If not, then the `%elif` and `%else` blocks (if any) will be processed instead.

For example, when debugging a program, you might want to write code such as

```
          ; perform some function
%ifdef DEBUG
          writefile 2,"Function performed successfully",13,10
%endif
          ; go and do something else
```

Then you could use the command–line option `-dDEBUG` to create a version of the program which produced debugging messages, and remove the option to generate the final release version of the program.

You can test for a macro *not* being defined by using `%ifndef` instead of `%ifdef`. You can also test for macro definitions in `%elif` blocks by using `%elifdef` and `%elifndef`.

4.4.2 `ifmacro`: Testing Multi–Line Macro Existence

The `%ifmacro` directive operates in the same way as the `%ifdef` directive, except that it checks for the existence of a multi–line macro.

For example, you may be working with a large project and not have control over the macros in a library. You may want to create a macro with one name if it doesn't already exist, and another name if one with that name does exist.

The `%ifmacro` is considered true if defining a macro with the given name and number of arguments would cause a definitions conflict. For example:

```
%ifmacro MyMacro 1-3

     %error "MyMacro 1-3" causes a conflict with an existing macro.

%else

     %macro MyMacro 1-3

             ; insert code to define the macro

     %endmacro

%endif
```

This will create the macro "MyMacro 1–3" if no macro already exists which would conflict with it, and emits a warning if there would be a definition conflict.

You can test for the macro not existing by using the `%ifnmacro` instead of `%ifmacro`. Additional tests can be performed in `%elif` blocks by using `%elifmacro` and `%elifnmacro`.

4.4.3 `%ifctx`: Testing the Context Stack

The conditional–assembly construct `%ifctx ctxname` will cause the subsequent code to be assembled if and only if the top context on the preprocessor's context stack has the name `ctxname`. As with `%ifdef`, the inverse and `%elif` forms `%ifnctx`, `%elifctx` and `%elifnctx` are also supported.

For more details of the context stack, see section 4.7. For a sample use of `%ifctx`, see section 4.7.5.

4.4.4 `%if`: Testing Arbitrary Numeric Expressions

The conditional–assembly construct `%if expr` will cause the subsequent code to be assembled if and only if the value of the numeric expression `expr` is non–zero. An example of the use of this feature is in deciding when to break out of a `%rep` preprocessor loop: see section 4.5 for a detailed example.

The expression given to `%if`, and its counterpart `%elif`, is a critical expression (see section 3.8).

`%if` extends the normal NASM expression syntax, by providing a set of relational operators which are not normally available in expressions. The operators =, <, >, <=, >= and <> test equality, less–than, greater–than, less–or–equal, greater–or–equal and not–equal respectively. The C–like forms == and `!=` are supported as alternative forms of = and <>. In addition, low–priority logical operators `&&`, `^^` and `||` are provided, supplying logical AND, logical XOR and logical OR. These work like the C logical operators (although C has no logical XOR), in that they always return either 0 or 1, and treat any non–zero input as 1 (so that `^^`, for example, returns 1 if exactly one of its inputs is zero, and 0 otherwise). The relational operators also return 1 for true and 0 for false.

4.4.5 `%ifidn` and `%ifidni`: Testing Exact Text Identity

The construct `%ifidn text1,text2` will cause the subsequent code to be assembled if and only if `text1` and `text2`, after expanding single–line macros, are identical pieces of text. Differences in white space are not counted.

`%ifidni` is similar to `%ifidn`, but is case–insensitive.

For example, the following macro pushes a register or number on the stack, and allows you to treat `IP` as a real register:

```
%macro  pushparam 1

  %ifidni %1,ip
        call    %%label
  %%label:
  %else
        push    %1
  %endif

%endmacro
```

Like most other `%if` constructs, `%ifidn` has a counterpart `%elifidn`, and negative forms `%ifnidn` and `%elifnidn`. Similarly, `%ifidni` has counterparts `%elifidni`, `%ifnidni` and `%elifnidni`.

4.4.6 `%ifid`, `%ifnum`, `%ifstr`: Testing Token Types

Some macros will want to perform different tasks depending on whether they are passed a number, a string, or an identifier. For example, a string output macro might want to be able to cope with being passed either a string constant or a pointer to an existing string.

The conditional assembly construct `%ifid`, taking one parameter (which may be blank), assembles the subsequent code if and only if the first token in the parameter exists and is an identifier. `%ifnum` works similarly, but tests for the token being a numeric constant; `%ifstr` tests for it being a string.

For example, the `writefile` macro defined in section 4.3.3 can be extended to take advantage of `%ifstr` in the following fashion:

```
%macro writefile 2-3+

  %ifstr %2
        jmp     %%endstr
    %if %0 = 3
      %%str:    db      %2,%3
    %else
      %%str:    db      %2
    %endif
      %%endstr: mov     dx,%%str
                mov     cx,%%endstr-%%str
  %else
                mov     dx,%2
                mov     cx,%3
  %endif
                mov     bx,%1
                mov     ah,0x40
                int     0x21

%endmacro
```

Then the `writefile` macro can cope with being called in either of the following two ways:

```
        writefile [file], strpointer, length
        writefile [file], "hello", 13, 10
```

In the first, `strpointer` is used as the address of an already-declared string, and `length` is used as its length; in the second, a string is given to the macro, which therefore declares it itself and works out the address and length for itself.

Note the use of `%if` inside the `%ifstr`: this is to detect whether the macro was passed two arguments (so the string would be a single string constant, and `db %2` would be adequate) or more (in which case, all but the first two would be lumped together into `%3`, and `db %2,%3` would be required).

The usual `%elifXXX`, `%ifnXXX` and `%elifnXXX` versions exist for each of `%ifid`, `%ifnum` and `%ifstr`.

4.4.7 `%error`: Reporting User-Defined Errors

The preprocessor directive `%error` will cause NASM to report an error if it occurs in assembled code. So if other users are going to try to assemble your source files, you can ensure that they define the right macros by means of code like this:

```
%ifdef SOME_MACRO
    ; do some setup
%elifdef SOME_OTHER_MACRO
    ; do some different setup
%else
    %error Neither SOME_MACRO nor SOME_OTHER_MACRO was defined.
%endif
```

Then any user who fails to understand the way your code is supposed to be assembled will be quickly warned of their mistake, rather than having to wait until the program crashes on being run and then not knowing what went wrong.

4.5 Preprocessor Loops: %rep

NASM's TIMES prefix, though useful, cannot be used to invoke a multi-line macro multiple times, because it is processed by NASM after macros have already been expanded. Therefore NASM provides another form of loop, this time at the preprocessor level: %rep.

The directives %rep and %endrep (%rep takes a numeric argument, which can be an expression; %endrep takes no arguments) can be used to enclose a chunk of code, which is then replicated as many times as specified by the preprocessor:

```
%assign i 0
%rep    64
        inc     word [table+2*i]
%assign i i+1
%endrep
```

This will generate a sequence of 64 INC instructions, incrementing every word of memory from [table] to [table+126].

For more complex termination conditions, or to break out of a repeat loop part way along, you can use the %exitrep directive to terminate the loop, like this:

```
fibonacci:
%assign i 0
%assign j 1
%rep 100
%if j > 65535
    %exitrep
%endif
        dw j
%assign k j+i
%assign i j
%assign j k
%endrep

fib_number equ ($-fibonacci)/2
```

This produces a list of all the Fibonacci numbers that will fit in 16 bits. Note that a maximum repeat count must still be given to %rep. This is to prevent the possibility of NASM getting into an infinite loop in the preprocessor, which (on multitasking or multi-user systems) would typically cause all the system memory to be gradually used up and other applications to start crashing.

4.6 Including Other Files

Using, once again, a very similar syntax to the C preprocessor, NASM's preprocessor lets you include other source files into your code. This is done by the use of the %include directive:

```
%include "macros.mac"
```

will include the contents of the file macros.mac into the source file containing the %include directive.

Include files are searched for in the current directory (the directory you're in when you run NASM, as opposed to the location of the NASM executable or the location of the source file), plus any directories specified on the NASM command line using the -i option.

The standard C idiom for preventing a file being included more than once is just as applicable in NASM: if the file macros.mac has the form

```
%ifndef MACROS_MAC
    %define MACROS_MAC
    ; now define some macros
%endif
```

then including the file more than once will not cause errors, because the second time the file is included nothing will happen because the macro `MACROS_MAC` will already be defined.

You can force a file to be included even if there is no `%include` directive that explicitly includes it, by using the `-p` option on the NASM command line (see section 2.1.11).

4.7 The Context Stack

Having labels that are local to a macro definition is sometimes not quite powerful enough: sometimes you want to be able to share labels between several macro calls. An example might be a `REPEAT` ... `UNTIL` loop, in which the expansion of the `REPEAT` macro would need to be able to refer to a label which the `UNTIL` macro had defined. However, for such a macro you would also want to be able to nest these loops.

NASM provides this level of power by means of a *context stack*. The preprocessor maintains a stack of *contexts*, each of which is characterised by a name. You add a new context to the stack using the `%push` directive, and remove one using `%pop`. You can define labels that are local to a particular context on the stack.

4.7.1 %push and %pop: Creating and Removing Contexts

The `%push` directive is used to create a new context and place it on the top of the context stack. `%push` requires one argument, which is the name of the context. For example:

```
%push     foobar
```

This pushes a new context called `foobar` on the stack. You can have several contexts on the stack with the same name: they can still be distinguished.

The directive `%pop`, requiring no arguments, removes the top context from the context stack and destroys it, along with any labels associated with it.

4.7.2 Context-Local Labels

Just as the usage `%%foo` defines a label which is local to the particular macro call in which it is used, the usage `%$foo` is used to define a label which is local to the context on the top of the context stack. So the `REPEAT` and `UNTIL` example given above could be implemented by means of:

```
%macro repeat 0

    %push   repeat
    %$begin:

%endmacro

%macro until 1

        j%-1    %$begin
    %pop

%endmacro
```

and invoked by means of, for example,

```
        mov     cx,string
        repeat
        add     cx,3
        scasb
        until   e
```

which would scan every fourth byte of a string in search of the byte in `AL`.

If you need to define, or access, labels local to the context *below* the top one on the stack, you can use `%$$foo`, or `%$$$foo` for the context below that, and so on.

4.7.3 Context–Local Single–Line Macros

NASM also allows you to define single–line macros which are local to a particular context, in just the same way:

```
%define %$localmac 3
```

will define the single–line macro `%$localmac` to be local to the top context on the stack. Of course, after a subsequent `%push`, it can then still be accessed by the name `%$$localmac`.

4.7.4 `%repl`: Renaming a Context

If you need to change the name of the top context on the stack (in order, for example, to have it respond differently to `%ifctx`), you can execute a `%pop` followed by a `%push`; but this will have the side effect of destroying all context–local labels and macros associated with the context that was just popped.

NASM provides the directive `%repl`, which *replaces* a context with a different name, without touching the associated macros and labels. So you could replace the destructive code

```
%pop
%push   newname
```

with the non–destructive version `%repl newname`.

4.7.5 Example Use of the Context Stack: Block IFs

This example makes use of almost all the context–stack features, including the conditional–assembly construct `%ifctx`, to implement a block IF statement as a set of macros.

```
%macro if 1

    %push if
    j%-1  %$ifnot

%endmacro

%macro else 0

  %ifctx if
        %repl   else
        jmp     %$ifend
        %$ifnot:
  %else
        %error  "expected `if' before `else'"
  %endif

%endmacro
```

```
%macro endif 0

  %ifctx if
        %$ifnot:
        %pop
  %elifctx      else
        %$ifend:
        %pop
  %else
        %error  "expected `if' or `else' before `endif'"
  %endif

%endmacro
```

This code is more robust than the `REPEAT` and `UNTIL` macros given in section 4.7.2, because it uses conditional assembly to check that the macros are issued in the right order (for example, not calling `endif` before `if`) and issues a `%error` if they're not.

In addition, the `endif` macro has to be able to cope with the two distinct cases of either directly following an `if`, or following an `else`. It achieves this, again, by using conditional assembly to do different things depending on whether the context on top of the stack is `if` or `else`.

The `else` macro has to preserve the context on the stack, in order to have the `%$ifnot` referred to by the `if` macro be the same as the one defined by the `endif` macro, but has to change the context's name so that `endif` will know there was an intervening `else`. It does this by the use of `%repl`.

A sample usage of these macros might look like:

```
        cmp     ax,bx

        if ae
                cmp     bx,cx

                if ae
                        mov     ax,cx
                else
                        mov     ax,bx
                endif

        else
                cmp     ax,cx

                if ae
                        mov     ax,cx
                endif

        endif
```

The block-`IF` macros handle nesting quite happily, by means of pushing another context, describing the inner `if`, on top of the one describing the outer `if`; thus `else` and `endif` always refer to the last unmatched `if` or `else`.

4.8 Standard Macros

NASM defines a set of standard macros, which are already defined when it starts to process any source file. If you really need a program to be assembled with no pre–defined macros, you can use the `%clear` directive to empty the preprocessor of everything but context–local preprocessor variables and single–line macros.

Most user–level assembler directives (see chapter 5) are implemented as macros which invoke primitive directives; these are described in chapter 5. The rest of the standard macro set is described here.

4.8.1 `__NASM_MAJOR__`, `__NASM_MINOR__`, `__NASM_SUBMINOR__` and `____NASM_PATCHLEVEL__`: NASM Version

The single–line macros `__NASM_MAJOR__`, `__NASM_MINOR__`, `__NASM_SUBMINOR__` and `____NASM_PATCHLEVEL__` expand to the major, minor, subminor and patch level parts of the version number of NASM being used. So, under NASM 0.98.32p1 for example, `__NASM_MAJOR__` would be defined to be 0, `__NASM_MINOR__` would be defined as 98, `__NASM_SUBMINOR__` would be defined to 32, and `____NASM_PATCHLEVEL__` would be defined as 1.

4.8.2 `__NASM_VERSION_ID__`: NASM Version ID

The single–line macro `__NASM_VERSION_ID__` expands to a dword integer representing the full version number of the version of nasm being used. The value is the equivalent to `__NASM_MAJOR__`, `__NASM_MINOR__`, `__NASM_SUBMINOR__` and `____NASM_PATCHLEVEL__` concatenated to produce a single doubleword. Hence, for 0.98.32p1, the returned number would be equivalent to:

```
dd      0x00622001
```

or

```
db      1,32,98,0
```

Note that the above lines are generate exactly the same code, the second line is used just to give an indication of the order that the separate values will be present in memory.

4.8.3 `__NASM_VER__`: NASM Version string

The single–line macro `__NASM_VER__` expands to a string which defines the version number of nasm being used. So, under NASM 0.98.32 for example,

```
db      __NASM_VER__
```

would expand to

```
db      "0.98.32"
```

4.8.4 `__FILE__` and `__LINE__`: File Name and Line Number

Like the C preprocessor, NASM allows the user to find out the file name and line number containing the current instruction. The macro `__FILE__` expands to a string constant giving the name of the current input file (which may change through the course of assembly if `%include` directives are used), and `__LINE__` expands to a numeric constant giving the current line number in the input file.

These macros could be used, for example, to communicate debugging information to a macro, since invoking `__LINE__` inside a macro definition (either single–line or multi–line) will return the line number of the macro *call*, rather than *definition*. So to determine where in a piece of code a crash is

occurring, for example, one could write a routine `stillhere`, which is passed a line number in `EAX` and outputs something like ‘line 155: still here’. You could then write a macro

```
%macro  notdeadyet 0

        push    eax
        mov     eax,__LINE__
        call    stillhere
        pop     eax

%endmacro
```

and then pepper your code with calls to `notdeadyet` until you find the crash point.

4.8.5 STRUC and ENDSTRUC: Declaring Structure Data Types

The core of NASM contains no intrinsic means of defining data structures; instead, the preprocessor is sufficiently powerful that data structures can be implemented as a set of macros. The macros `STRUC` and `ENDSTRUC` are used to define a structure data type.

`STRUC` takes one parameter, which is the name of the data type. This name is defined as a symbol with the value zero, and also has the suffix `_size` appended to it and is then defined as an `EQU` giving the size of the structure. Once `STRUC` has been issued, you are defining the structure, and should define fields using the `RESB` family of pseudo−instructions, and then invoke `ENDSTRUC` to finish the definition.

For example, to define a structure called `mytype` containing a longword, a word, a byte and a string of bytes, you might code

```
struc   mytype

  mt_long:      resd    1
  mt_word:      resw    1
  mt_byte:      resb    1
  mt_str:       resb    32

endstruc
```

The above code defines six symbols: `mt_long` as 0 (the offset from the beginning of a `mytype` structure to the longword field), `mt_word` as 4, `mt_byte` as 6, `mt_str` as 7, `mytype_size` as 39, and `mytype` itself as zero.

The reason why the structure type name is defined at zero is a side effect of allowing structures to work with the local label mechanism: if your structure members tend to have the same names in more than one structure, you can define the above structure like this:

```
struc mytype

  .long:        resd    1
  .word:        resw    1
  .byte:        resb    1
  .str:         resb    32

endstruc
```

This defines the offsets to the structure fields as `mytype.long`, `mytype.word`, `mytype.byte` and `mytype.str`.

NASM, since it has no *intrinsic* structure support, does not support any form of period notation to refer to the elements of a structure once you have one (except the above local–label notation), so code such as `mov ax,[mystruc.mt_word]` is not valid. `mt_word` is a constant just like any other constant, so the correct syntax is `mov ax,[mystruc+mt_word]` or `mov ax,[mystruc+mytype.word]`.

4.8.6 `ISTRUC`, `AT` and `IEND`: Declaring Instances of Structures

Having defined a structure type, the next thing you typically want to do is to declare instances of that structure in your data segment. NASM provides an easy way to do this in the `ISTRUC` mechanism. To declare a structure of type `mytype` in a program, you code something like this:

```
mystruc:
    istruc mytype

        at mt_long, dd      123456
        at mt_word, dw      1024
        at mt_byte, db      'x'
        at mt_str,  db      'hello, world', 13, 10, 0

    iend
```

The function of the `AT` macro is to make use of the `TIMES` prefix to advance the assembly position to the correct point for the specified structure field, and then to declare the specified data. Therefore the structure fields must be declared in the same order as they were specified in the structure definition.

If the data to go in a structure field requires more than one source line to specify, the remaining source lines can easily come after the `AT` line. For example:

```
        at mt_str,  db      123,134,145,156,167,178,189
                    db      190,100,0
```

Depending on personal taste, you can also omit the code part of the `AT` line completely, and start the structure field on the next line:

```
        at mt_str
                db      'hello, world'
                db      13,10,0
```

4.8.7 `ALIGN` and `ALIGNB`: Data Alignment

The `ALIGN` and `ALIGNB` macros provides a convenient way to align code or data on a word, longword, paragraph or other boundary. (Some assemblers call this directive `EVEN`.) The syntax of the `ALIGN` and `ALIGNB` macros is

```
        align   4               ; align on 4-byte boundary
        align   16              ; align on 16-byte boundary
        align   8,db 0          ; pad with 0s rather than NOPs
        align   4,resb 1        ; align to 4 in the BSS
        alignb  4               ; equivalent to previous line
```

Both macros require their first argument to be a power of two; they both compute the number of additional bytes required to bring the length of the current section up to a multiple of that power of two, and then apply the `TIMES` prefix to their second argument to perform the alignment.

If the second argument is not specified, the default for `ALIGN` is `NOP`, and the default for `ALIGNB` is `RESB 1`. So if the second argument is specified, the two macros are equivalent. Normally, you

can just use ALIGN in code and data sections and ALIGNB in BSS sections, and never need the second argument except for special purposes.

ALIGN and ALIGNB, being simple macros, perform no error checking: they cannot warn you if their first argument fails to be a power of two, or if their second argument generates more than one byte of code. In each of these cases they will silently do the wrong thing.

ALIGNB (or ALIGN with a second argument of RESB 1) can be used within structure definitions:

```
struc mytype2

  mt_byte:
        resb 1
        alignb 2
  mt_word:
        resw 1
        alignb 4
  mt_long:
        resd 1
  mt_str:
        resb 32

endstruc
```

This will ensure that the structure members are sensibly aligned relative to the base of the structure.

A final caveat: ALIGN and ALIGNB work relative to the beginning of the *section*, not the beginning of the address space in the final executable. Aligning to a 16–byte boundary when the section you're in is only guaranteed to be aligned to a 4–byte boundary, for example, is a waste of effort. Again, NASM does not check that the section's alignment characteristics are sensible for the use of ALIGN or ALIGNB.

4.9 TASM Compatible Preprocessor Directives

The following preprocessor directives may only be used when TASM compatibility is turned on using the -t command line switch (This switch is described in section 2.1.17.)

- %arg (see section 4.9.1)
- %stacksize (see section 4.9.2)
- %local (see section 4.9.3)

4.9.1 %arg Directive

The %arg directive is used to simplify the handling of parameters passed on the stack. Stack based parameter passing is used by many high level languages, including C, C++ and Pascal.

While NASM comes with macros which attempt to duplicate this functionality (see section 7.4.5), the syntax is not particularly convenient to use and is not TASM compatible. Here is an example which shows the use of %arg without any external macros:

```
some_function:

    %push     mycontext         ; save the current context
    %stacksize large            ; tell NASM to use bp
    %arg      i:word, j_ptr:word

        mov     ax,[i]
```

```
        mov     bx,[j_ptr]
        add     ax,[bx]
        ret

    %pop                        ; restore original context
```

This is similar to the procedure defined in section 7.4.5 and adds the value in i to the value pointed to by j_ptr and returns the sum in the ax register. See section 4.7.1 for an explanation of `push` and `pop` and the use of context stacks.

4.9.2 `%stacksize` Directive

The `%stacksize` directive is used in conjunction with the `%arg` (see section 4.9.1) and the `%local` (see section 4.9.3) directives. It tells NASM the default size to use for subsequent `%arg` and `%local` directives. The `%stacksize` directive takes one required argument which is one of `flat`, `large` or `small`.

```
%stacksize flat
```

This form causes NASM to use stack–based parameter addressing relative to `ebp` and it assumes that a near form of call was used to get to this label (i.e. that `eip` is on the stack).

```
%stacksize large
```

This form uses `bp` to do stack–based parameter addressing and assumes that a far form of call was used to get to this address (i.e. that `ip` and `cs` are on the stack).

```
%stacksize small
```

This form also uses `bp` to address stack parameters, but it is different from `large` because it also assumes that the old value of bp is pushed onto the stack (i.e. it expects an `ENTER` instruction). In other words, it expects that `bp`, `ip` and `cs` are on the top of the stack, underneath any local space which may have been allocated by `ENTER`. This form is probably most useful when used in combination with the `%local` directive (see section 4.9.3).

4.9.3 `%local` Directive

The `%local` directive is used to simplify the use of local temporary stack variables allocated in a stack frame. Automatic local variables in C are an example of this kind of variable. The `%local` directive is most useful when used with the `%stacksize` (see section 4.9.2 and is also compatible with the `%arg` directive (see section 4.9.1). It allows simplified reference to variables on the stack which have been allocated typically by using the `ENTER` instruction (see section B.4.65 for a description of that instruction). An example of its use is the following:

```
silly_swap:

    %push mycontext             ; save the current context
    %stacksize small            ; tell NASM to use bp
    %assign %$localsize 0       ; see text for explanation
    %local old_ax:word, old_dx:word

        enter   %$localsize,0   ; see text for explanation
        mov     [old_ax],ax     ; swap ax & bx
        mov     [old_dx],dx     ; and swap dx & cx
        mov     ax,bx
        mov     dx,cx
        mov     bx,[old_ax]
        mov     cx,[old_dx]
        leave                   ; restore old bp
```

```
        ret                             ;

    %pop                                ; restore original context
```

The `%$localsize` variable is used internally by the `%local` directive and *must* be defined within the current context before the `%local` directive may be used. Failure to do so will result in one expression syntax error for each `%local` variable declared. It then may be used in the construction of an appropriately sized ENTER instruction as shown in the example.

4.10 Other Preprocessor Directives

NASM also has preprocessor directives which allow access to information from external sources. Currently they include:

The following preprocessor directive is supported to allow NASM to correctly handle output of the cpp C language preprocessor.

- `%line` enables NAsM to correctly handle the output of the cpp C language preprocessor (see section 4.10.1).
- `%!` enables NASM to read in the value of an environment variable, which can then be used in your program (see section 4.10.2).

4.10.1 `%line` Directive

The `%line` directive is used to notify NASM that the input line corresponds to a specific line number in another file. Typically this other file would be an original source file, with the current NASM input being the output of a pre−processor. The `%line` directive allows NASM to output messages which indicate the line number of the original source file, instead of the file that is being read by NASM.

This preprocessor directive is not generally of use to programmers, by may be of interest to preprocessor authors. The usage of the `%line` preprocessor directive is as follows:

```
%line nnn[+mmm] [filename]
```

In this directive, `nnn` indentifies the line of the original source file which this line corresponds to. `mmm` is an optional parameter which specifies a line increment value; each line of the input file read in is considered to correspond to `mmm` lines of the original source file. Finally, `filename` is an optional parameter which specifies the file name of the original source file.

After reading a `%line` preprocessor directive, NASM will report all file name and line numbers relative to the values specified therein.

4.10.2 `%!<env>`: Read an environment variable.

The `%!<env>` directive makes it possible to read the value of an environment variable at assembly time. This could, for example, be used to store the contents of an environment variable into a string, which could be used at some other point in your code.

For example, suppose that you have an environment variable `FOO`, and you want the contents of `FOO` to be embedded in your program. You could do that as follows:

```
%define FOO     %!FOO
%define quote   '

tmpstr  db      quote FOO quote
```

At the time of writing, this will generate an "unterminated string" warning at the time of defining "quote", and it will add a space before and after the string that is read in. I was unable to find a simple workaround (although a workaround can be created using a multi−line macro), so I believe

that you will need to either learn how to create more complex macros, or allow for the extra spaces if you make use of this feature in that way.

Chapter 5: Assembler Directives

NASM, though it attempts to avoid the bureaucracy of assemblers like MASM and TASM, is nevertheless forced to support a *few* directives. These are described in this chapter.

NASM's directives come in two types: *user−level* directives and *primitive* directives. Typically, each directive has a user−level form and a primitive form. In almost all cases, we recommend that users use the user−level forms of the directives, which are implemented as macros which call the primitive forms.

Primitive directives are enclosed in square brackets; user−level directives are not.

In addition to the universal directives described in this chapter, each object file format can optionally supply extra directives in order to control particular features of that file format. These *format−specific* directives are documented along with the formats that implement them, in chapter 6.

5.1 `BITS`: Specifying Target Processor Mode

The `BITS` directive specifies whether NASM should generate code designed to run on a processor operating in 16−bit mode, or code designed to run on a processor operating in 32−bit mode. The syntax is `BITS 16` or `BITS 32`.

In most cases, you should not need to use `BITS` explicitly. The `aout`, `coff`, `elf` and `win32` object formats, which are designed for use in 32−bit operating systems, all cause NASM to select 32−bit mode by default. The `obj` object format allows you to specify each segment you define as either `USE16` or `USE32`, and NASM will set its operating mode accordingly, so the use of the `BITS` directive is once again unnecessary.

The most likely reason for using the `BITS` directive is to write 32−bit code in a flat binary file; this is because the `bin` output format defaults to 16−bit mode in anticipation of it being used most frequently to write DOS `.COM` programs, DOS `.SYS` device drivers and boot loader software.

You do *not* need to specify `BITS 32` merely in order to use 32−bit instructions in a 16−bit DOS program; if you do, the assembler will generate incorrect code because it will be writing code targeted at a 32−bit platform, to be run on a 16−bit one.

When NASM is in `BITS 16` state, instructions which use 32−bit data are prefixed with an 0x66 byte, and those referring to 32−bit addresses have an 0x67 prefix. In `BITS 32` state, the reverse is true: 32−bit instructions require no prefixes, whereas instructions using 16−bit data need an 0x66 and those working on 16−bit addresses need an 0x67.

The `BITS` directive has an exactly equivalent primitive form, `[BITS 16]` and `[BITS 32]`. The user−level form is a macro which has no function other than to call the primitive form.

Note that the space is neccessary, `BITS32` will *not* work!

5.1.1 `USE16` & `USE32`: Aliases for BITS

The '`USE16`' and '`USE32`' directives can be used in place of '`BITS 16`' and '`BITS 32`', for compatibility with other assemblers.

5.2 `SECTION` or `SEGMENT`: Changing and Defining Sections

The `SECTION` directive (`SEGMENT` is an exactly equivalent synonym) changes which section of the output file the code you write will be assembled into. In some object file formats, the number

and names of sections are fixed; in others, the user may make up as many as they wish. Hence SECTION may sometimes give an error message, or may define a new section, if you try to switch to a section that does not (yet) exist.

The Unix object formats, and the bin object format (but see section 6.1.3, all support the standardised section names .text, .data and .bss for the code, data and uninitialised–data sections. The obj format, by contrast, does not recognise these section names as being special, and indeed will strip off the leading period of any section name that has one.

5.2.1 The __SECT__ Macro

The SECTION directive is unusual in that its user–level form functions differently from its primitive form. The primitive form, [SECTION xyz], simply switches the current target section to the one given. The user–level form, SECTION xyz, however, first defines the single–line macro __SECT__ to be the primitive [SECTION] directive which it is about to issue, and then issues it. So the user–level directive

```
        SECTION .text
```

expands to the two lines

```
%define __SECT__        [SECTION .text]
        [SECTION .text]
```

Users may find it useful to make use of this in their own macros. For example, the writefile macro defined in section 4.3.3 can be usefully rewritten in the following more sophisticated form:

```
%macro  writefile 2+

        [section .data]

  %%str:        db      %2
  %%endstr:

        __SECT__

        mov     dx,%%str
        mov     cx,%%endstr-%%str
        mov     bx,%1
        mov     ah,0x40
        int     0x21

%endmacro
```

This form of the macro, once passed a string to output, first switches temporarily to the data section of the file, using the primitive form of the SECTION directive so as not to modify __SECT__. It then declares its string in the data section, and then invokes __SECT__ to switch back to *whichever* section the user was previously working in. It thus avoids the need, in the previous version of the macro, to include a JMP instruction to jump over the data, and also does not fail if, in a complicated OBJ format module, the user could potentially be assembling the code in any of several separate code sections.

5.3 ABSOLUTE: Defining Absolute Labels

The ABSOLUTE directive can be thought of as an alternative form of SECTION: it causes the subsequent code to be directed at no physical section, but at the hypothetical section starting at the given absolute address. The only instructions you can use in this mode are the RESB family.

ABSOLUTE is used as follows:

```
absolute 0x1A

    kbuf_chr    resw    1
    kbuf_free   resw    1
    kbuf        resw    16
```

This example describes a section of the PC BIOS data area, at segment address 0x40: the above code defines kbuf_chr to be 0x1A, kbuf_free to be 0x1C, and kbuf to be 0x1E.

The user–level form of ABSOLUTE, like that of SECTION, redefines the __SECT__ macro when it is invoked.

STRUC and ENDSTRUC are defined as macros which use ABSOLUTE (and also __SECT__).

ABSOLUTE doesn't have to take an absolute constant as an argument: it can take an expression (actually, a critical expression: see section 3.8) and it can be a value in a segment. For example, a TSR can re–use its setup code as run–time BSS like this:

```
        org     100h                ; it's a .COM program

        jmp     setup               ; setup code comes last

        ; the resident part of the TSR goes here
setup:
        ; now write the code that installs the TSR here

absolute setup

runtimevar1     resw    1
runtimevar2     resd    20

tsr_end:
```

This defines some variables 'on top of' the setup code, so that after the setup has finished running, the space it took up can be re–used as data storage for the running TSR. The symbol 'tsr_end' can be used to calculate the total size of the part of the TSR that needs to be made resident.

5.4 EXTERN: Importing Symbols from Other Modules

EXTERN is similar to the MASM directive EXTRN and the C keyword extern: it is used to declare a symbol which is not defined anywhere in the module being assembled, but is assumed to be defined in some other module and needs to be referred to by this one. Not every object–file format can support external variables: the bin format cannot.

The EXTERN directive takes as many arguments as you like. Each argument is the name of a symbol:

```
extern  _printf
extern  _sscanf,_fscanf
```

Some object–file formats provide extra features to the EXTERN directive. In all cases, the extra features are used by suffixing a colon to the symbol name followed by object–format specific text. For example, the obj format allows you to declare that the default segment base of an external should be the group dgroup by means of the directive

```
extern  _variable:wrt dgroup
```

The primitive form of EXTERN differs from the user–level form only in that it can take only one argument at a time: the support for multiple arguments is implemented at the preprocessor level.

You can declare the same variable as EXTERN more than once: NASM will quietly ignore the second and later redeclarations. You can't declare a variable as EXTERN as well as something else, though.

5.5 GLOBAL: Exporting Symbols to Other Modules

GLOBAL is the other end of EXTERN: if one module declares a symbol as EXTERN and refers to it, then in order to prevent linker errors, some other module must actually *define* the symbol and declare it as GLOBAL. Some assemblers use the name PUBLIC for this purpose.

The GLOBAL directive applying to a symbol must appear *before* the definition of the symbol.

GLOBAL uses the same syntax as EXTERN, except that it must refer to symbols which *are* defined in the same module as the GLOBAL directive. For example:

```
global _main
_main:
        ; some code
```

GLOBAL, like EXTERN, allows object formats to define private extensions by means of a colon. The elf object format, for example, lets you specify whether global data items are functions or data:

```
global  hashlookup:function, hashtable:data
```

Like EXTERN, the primitive form of GLOBAL differs from the user–level form only in that it can take only one argument at a time.

5.6 COMMON: Defining Common Data Areas

The COMMON directive is used to declare *common variables*. A common variable is much like a global variable declared in the uninitialised data section, so that

```
common  intvar  4
```

is similar in function to

```
global  intvar
section .bss

intvar  resd    1
```

The difference is that if more than one module defines the same common variable, then at link time those variables will be *merged*, and references to intvar in all modules will point at the same piece of memory.

Like GLOBAL and EXTERN, COMMON supports object–format specific extensions. For example, the obj format allows common variables to be NEAR or FAR, and the elf format allows you to specify the alignment requirements of a common variable:

```
common  commvar  4:near  ; works in OBJ
common  intarray 100:4   ; works in ELF: 4 byte aligned
```

Once again, like EXTERN and GLOBAL, the primitive form of COMMON differs from the user–level form only in that it can take only one argument at a time.

5.7 CPU: Defining CPU Dependencies

The CPU directive restricts assembly to those instructions which are available on the specified CPU.

Options are:

- `CPU 8086` Assemble only 8086 instruction set
- `CPU 186` Assemble instructions up to the 80186 instruction set
- `CPU 286` Assemble instructions up to the 286 instruction set
- `CPU 386` Assemble instructions up to the 386 instruction set
- `CPU 486` 486 instruction set
- `CPU 586` Pentium instruction set
- `CPU PENTIUM` Same as 586
- `CPU 686` P6 instruction set
- `CPU PPRO` Same as 686
- `CPU P2` Same as 686
- `CPU P3` Pentium III (Katmai) instruction sets
- `CPU KATMAI` Same as P3
- `CPU P4` Pentium 4 (Willamette) instruction set
- `CPU WILLAMETTE` Same as P4
- `CPU PRESCOTT` Prescott instruction set
- `CPU IA64` IA64 CPU (in x86 mode) instruction set

All options are case insensitive. All instructions will be selected only if they apply to the selected CPU or lower. By default, all instructions are available.

Chapter 6: Output Formats

NASM is a portable assembler, designed to be able to compile on any ANSI C–supporting platform and produce output to run on a variety of Intel x86 operating systems. For this reason, it has a large number of available output formats, selected using the `-f` option on the NASM command line. Each of these formats, along with its extensions to the base NASM syntax, is detailed in this chapter.

As stated in section 2.1.1, NASM chooses a default name for your output file based on the input file name and the chosen output format. This will be generated by removing the extension (`.asm`, `.s`, or whatever you like to use) from the input file name, and substituting an extension defined by the output format. The extensions are given with each format below.

6.1 `bin`: Flat–Form Binary Output

The `bin` format does not produce object files: it generates nothing in the output file except the code you wrote. Such 'pure binary' files are used by MS–DOS: `.COM` executables and `.SYS` device drivers are pure binary files. Pure binary output is also useful for operating system and boot loader development.

The `bin` format supports multiple section names. For details of how nasm handles sections in the `bin` format, see section 6.1.3.

Using the `bin` format puts NASM by default into 16–bit mode (see section 5.1). In order to use `bin` to write 32–bit code such as an OS kernel, you need to explicitly issue the `BITS 32` directive.

`bin` has no default output file name extension: instead, it leaves your file name as it is once the original extension has been removed. Thus, the default is for NASM to assemble `binprog.asm` into a binary file called `binprog`.

6.1.1 `ORG`: Binary File Program Origin

The `bin` format provides an additional directive to the list given in chapter 5: `ORG`. The function of the `ORG` directive is to specify the origin address which NASM will assume the program begins at when it is loaded into memory.

For example, the following code will generate the longword `0x00000104`:

```
        org     0x100
        dd      label
label:
```

Unlike the `ORG` directive provided by MASM–compatible assemblers, which allows you to jump around in the object file and overwrite code you have already generated, NASM's `ORG` does exactly what the directive says: *origin*. Its sole function is to specify one offset which is added to all internal address references within the section; it does not permit any of the trickery that MASM's version does. See section 10.1.3 for further comments.

6.1.2 `bin` Extensions to the `SECTION` Directive

The `bin` output format extends the `SECTION` (or `SEGMENT`) directive to allow you to specify the alignment requirements of segments. This is done by appending the `ALIGN` qualifier to the end of the section–definition line. For example,

```
section .data   align=16
```

switches to the section `.data` and also specifies that it must be aligned on a 16–byte boundary.

The parameter to `ALIGN` specifies how many low bits of the section start address must be forced to zero. The alignment value given may be any power of two.

6.1.3 `Multisection` support for the BIN format.

The `bin` format allows the use of multiple sections, of arbitrary names, besides the "known" `.text`, `.data`, and `.bss` names.

- Sections may be designated `progbits` or `nobits`. Default is `progbits` (except `.bss`, which defaults to `nobits`, of course).
- Sections can be aligned at a specified boundary following the previous section with `align=`, or at an arbitrary byte–granular position with `start=`.
- Sections can be given a virtual start address, which will be used for the calculation of all memory references within that section with `vstart=`.
- Sections can be ordered using `follows=<section>` or `vfollows=<section>` as an alternative to specifying an explicit start address.
- Arguments to `org`, `start`, `vstart`, and `align=` are critical expressions. See section 3.8. E.g. `align=(1 << ALIGN_SHIFT) – ALIGN_SHIFT` must be defined before it is used here.
- Any code which comes before an explicit `SECTION` directive is directed by default into the `.text` section.
- If an `ORG` statement is not given, `ORG 0` is used by default.
- The `.bss` section will be placed after the last `progbits` section, unless `start=`, `vstart=`, `follows=`, or `vfollows=` has been specified.
- All sections are aligned on dword boundaries, unless a different alignment has been specified.
- Sections may not overlap.
- Nasm creates the `section.<secname>.start` for each section, which may be used in your code.

6.1.4 Map files

Map files can be generated in `–f bin` format by means of the `[map]` option. Map types of `all` (default), `brief`, `sections`, `segments`, or `symbols` may be specified. Output may be directed to `stdout` (default), `stderr`, or a specified file. E.g. `[map symbols myfile.map]`. No "user form" exists, the square brackets must be used.

6.2 `obj`: Microsoft OMF Object Files

The `obj` file format (NASM calls it `obj` rather than `omf` for historical reasons) is the one produced by MASM and TASM, which is typically fed to 16–bit DOS linkers to produce `.EXE` files. It is also the format used by OS/2.

`obj` provides a default output file–name extension of `.obj`.

`obj` is not exclusively a 16–bit format, though: NASM has full support for the 32–bit extensions to the format. In particular, 32–bit `obj` format files are used by Borland's Win32 compilers, instead of using Microsoft's newer `win32` object file format.

The `obj` format does not define any special segment names: you can call your segments anything you like. Typical names for segments in `obj` format files are `CODE`, `DATA` and `BSS`.

If your source file contains code before specifying an explicit SEGMENT directive, then NASM will invent its own segment called __NASMDEFSEG for you.

When you define a segment in an obj file, NASM defines the segment name as a symbol as well, so that you can access the segment address of the segment. So, for example:

```
segment data

dvar:   dw      1234

segment code

function:
        mov     ax,data         ; get segment address of data
        mov     ds,ax           ; and move it into DS
        inc     word [dvar]     ; now this reference will work
        ret
```

The obj format also enables the use of the SEG and WRT operators, so that you can write code which does things like

```
extern  foo

      mov   ax,seg foo            ; get preferred segment of foo
      mov   ds,ax
      mov   ax,data               ; a different segment
      mov   es,ax
      mov   ax,[ds:foo]           ; this accesses 'foo'
      mov   [es:foo wrt data],bx  ; so does this
```

6.2.1 obj Extensions to the SEGMENT Directive

The obj output format extends the SEGMENT (or SECTION) directive to allow you to specify various properties of the segment you are defining. This is done by appending extra qualifiers to the end of the segment–definition line. For example,

```
segment code private align=16
```

defines the segment code, but also declares it to be a private segment, and requires that the portion of it described in this code module must be aligned on a 16–byte boundary.

The available qualifiers are:

- PRIVATE, PUBLIC, COMMON and STACK specify the combination characteristics of the segment. PRIVATE segments do not get combined with any others by the linker; PUBLIC and STACK segments get concatenated together at link time; and COMMON segments all get overlaid on top of each other rather than stuck end–to–end.
- ALIGN is used, as shown above, to specify how many low bits of the segment start address must be forced to zero. The alignment value given may be any power of two from 1 to 4096; in reality, the only values supported are 1, 2, 4, 16, 256 and 4096, so if 8 is specified it will be rounded up to 16, and 32, 64 and 128 will all be rounded up to 256, and so on. Note that alignment to 4096–byte boundaries is a PharLap extension to the format and may not be supported by all linkers.
- CLASS can be used to specify the segment class; this feature indicates to the linker that segments of the same class should be placed near each other in the output file. The class name can be any word, e.g. CLASS=CODE.

- `OVERLAY`, like `CLASS`, is specified with an arbitrary word as an argument, and provides overlay information to an overlay–capable linker.
- Segments can be declared as `USE16` or `USE32`, which has the effect of recording the choice in the object file and also ensuring that NASM's default assembly mode when assembling in that segment is 16–bit or 32–bit respectively.
- When writing OS/2 object files, you should declare 32–bit segments as `FLAT`, which causes the default segment base for anything in the segment to be the special group `FLAT`, and also defines the group if it is not already defined.
- The `obj` file format also allows segments to be declared as having a pre–defined absolute segment address, although no linkers are currently known to make sensible use of this feature; nevertheless, NASM allows you to declare a segment such as `SEGMENT SCREEN ABSOLUTE=0xB800` if you need to. The `ABSOLUTE` and `ALIGN` keywords are mutually exclusive.

NASM's default segment attributes are `PUBLIC`, `ALIGN=1`, no class, no overlay, and `USE16`.

6.2.2 GROUP: Defining Groups of Segments

The `obj` format also allows segments to be grouped, so that a single segment register can be used to refer to all the segments in a group. NASM therefore supplies the `GROUP` directive, whereby you can code

```
segment data

        ; some data

segment bss

        ; some uninitialised data

group dgroup data bss
```

which will define a group called `dgroup` to contain the segments `data` and `bss`. Like `SEGMENT`, `GROUP` causes the group name to be defined as a symbol, so that you can refer to a variable `var` in the `data` segment as `var wrt data` or as `var wrt dgroup`, depending on which segment value is currently in your segment register.

If you just refer to `var`, however, and `var` is declared in a segment which is part of a group, then NASM will default to giving you the offset of `var` from the beginning of the *group*, not the *segment*. Therefore `SEG var`, also, will return the group base rather than the segment base.

NASM will allow a segment to be part of more than one group, but will generate a warning if you do this. Variables declared in a segment which is part of more than one group will default to being relative to the first group that was defined to contain the segment.

A group does not have to contain any segments; you can still make `WRT` references to a group which does not contain the variable you are referring to. OS/2, for example, defines the special group `FLAT` with no segments in it.

6.2.3 UPPERCASE: Disabling Case Sensitivity in Output

Although NASM itself is case sensitive, some OMF linkers are not; therefore it can be useful for NASM to output single–case object files. The `UPPERCASE` format–specific directive causes all segment, group and symbol names that are written to the object file to be forced to upper case just before being written. Within a source file, NASM is still case–sensitive; but the object file can be written entirely in upper case if desired.

UPPERCASE is used alone on a line; it requires no parameters.

6.2.4 IMPORT: Importing DLL Symbols

The IMPORT format−specific directive defines a symbol to be imported from a DLL, for use if you are writing a DLL's import library in NASM. You still need to declare the symbol as EXTERN as well as using the IMPORT directive.

The IMPORT directive takes two required parameters, separated by white space, which are (respectively) the name of the symbol you wish to import and the name of the library you wish to import it from. For example:

```
import  WSAStartup wsock32.dll
```

A third optional parameter gives the name by which the symbol is known in the library you are importing it from, in case this is not the same as the name you wish the symbol to be known by to your code once you have imported it. For example:

```
import  asyncsel wsock32.dll WSAAsyncSelect
```

6.2.5 EXPORT: Exporting DLL Symbols

The EXPORT format−specific directive defines a global symbol to be exported as a DLL symbol, for use if you are writing a DLL in NASM. You still need to declare the symbol as GLOBAL as well as using the EXPORT directive.

EXPORT takes one required parameter, which is the name of the symbol you wish to export, as it was defined in your source file. An optional second parameter (separated by white space from the first) gives the *external* name of the symbol: the name by which you wish the symbol to be known to programs using the DLL. If this name is the same as the internal name, you may leave the second parameter off.

Further parameters can be given to define attributes of the exported symbol. These parameters, like the second, are separated by white space. If further parameters are given, the external name must also be specified, even if it is the same as the internal name. The available attributes are:

- resident indicates that the exported name is to be kept resident by the system loader. This is an optimisation for frequently used symbols imported by name.
- nodata indicates that the exported symbol is a function which does not make use of any initialised data.
- parm=NNN, where NNN is an integer, sets the number of parameter words for the case in which the symbol is a call gate between 32−bit and 16−bit segments.
- An attribute which is just a number indicates that the symbol should be exported with an identifying number (ordinal), and gives the desired number.

For example:

```
export  myfunc
export  myfunc TheRealMoreFormalLookingFunctionName
export  myfunc myfunc 1234  ; export by ordinal
export  myfunc myfunc resident parm=23 nodata
```

6.2.6 ..start: Defining the Program Entry Point

OMF linkers require exactly one of the object files being linked to define the program entry point, where execution will begin when the program is run. If the object file that defines the entry point is assembled using NASM, you specify the entry point by declaring the special symbol ..start at the point where you wish execution to begin.

6.2.7 `obj` Extensions to the `EXTERN` Directive

If you declare an external symbol with the directive

```
extern  foo
```

then references such as `mov ax,foo` will give you the offset of `foo` from its preferred segment base (as specified in whichever module `foo` is actually defined in). So to access the contents of `foo` you will usually need to do something like

```
mov     ax,seg foo      ; get preferred segment base
mov     es,ax           ; move it into ES
mov     ax,[es:foo]     ; and use offset 'foo' from it
```

This is a little unwieldy, particularly if you know that an external is going to be accessible from a given segment or group, say `dgroup`. So if `DS` already contained `dgroup`, you could simply code

```
mov     ax,[foo wrt dgroup]
```

However, having to type this every time you want to access `foo` can be a pain; so NASM allows you to declare `foo` in the alternative form

```
extern  foo:wrt dgroup
```

This form causes NASM to pretend that the preferred segment base of `foo` is in fact `dgroup`; so the expression `seg foo` will now return `dgroup`, and the expression `foo` is equivalent to `foo wrt dgroup`.

This default-`WRT` mechanism can be used to make externals appear to be relative to any group or segment in your program. It can also be applied to common variables: see section 6.2.8.

6.2.8 `obj` Extensions to the `COMMON` Directive

The `obj` format allows common variables to be either near or far; NASM allows you to specify which your variables should be by the use of the syntax

```
common  nearvar 2:near   ; 'nearvar' is a near common
common  farvar  10:far   ; and 'farvar' is far
```

Far common variables may be greater in size than 64Kb, and so the OMF specification says that they are declared as a number of *elements* of a given size. So a 10-byte far common variable could be declared as ten one-byte elements, five two-byte elements, two five-byte elements or one ten-byte element.

Some `OMF` linkers require the element size, as well as the variable size, to match when resolving common variables declared in more than one module. Therefore NASM must allow you to specify the element size on your far common variables. This is done by the following syntax:

```
common  c_5by2  10:far 5        ; two five-byte elements
common  c_2by5  10:far 2        ; five two-byte elements
```

If no element size is specified, the default is 1. Also, the `FAR` keyword is not required when an element size is specified, since only far commons may have element sizes at all. So the above declarations could equivalently be

```
common  c_5by2  10:5            ; two five-byte elements
common  c_2by5  10:2            ; five two-byte elements
```

In addition to these extensions, the `COMMON` directive in `obj` also supports default-`WRT` specification like `EXTERN` does (explained in section 6.2.7). So you can also declare things like

```
common  foo     10:wrt dgroup
common  bar     16:far 2:wrt data
common  baz     24:wrt data:6
```

6.3 win32: Microsoft Win32 Object Files

The win32 output format generates Microsoft Win32 object files, suitable for passing to Microsoft linkers such as Visual C++. Note that Borland Win32 compilers do not use this format, but use obj instead (see section 6.2).

win32 provides a default output file–name extension of .obj.

Note that although Microsoft say that Win32 object files follow the COFF (Common Object File Format) standard, the object files produced by Microsoft Win32 compilers are not compatible with COFF linkers such as DJGPP's, and vice versa. This is due to a difference of opinion over the precise semantics of PC–relative relocations. To produce COFF files suitable for DJGPP, use NASM's coff output format; conversely, the coff format does not produce object files that Win32 linkers can generate correct output from.

6.3.1 win32 Extensions to the SECTION Directive

Like the obj format, win32 allows you to specify additional information on the SECTION directive line, to control the type and properties of sections you declare. Section types and properties are generated automatically by NASM for the standard section names .text, .data and .bss, but may still be overridden by these qualifiers.

The available qualifiers are:

- code, or equivalently text, defines the section to be a code section. This marks the section as readable and executable, but not writable, and also indicates to the linker that the type of the section is code.
- data and bss define the section to be a data section, analogously to code. Data sections are marked as readable and writable, but not executable. data declares an initialised data section, whereas bss declares an uninitialised data section.
- rdata declares an initialised data section that is readable but not writable. Microsoft compilers use this section to place constants in it.
- info defines the section to be an informational section, which is not included in the executable file by the linker, but may (for example) pass information *to* the linker. For example, declaring an info–type section called .drectve causes the linker to interpret the contents of the section as command–line options.
- align=, used with a trailing number as in obj, gives the alignment requirements of the section. The maximum you may specify is 64: the Win32 object file format contains no means to request a greater section alignment than this. If alignment is not explicitly specified, the defaults are 16–byte alignment for code sections, 8–byte alignment for rdata sections and 4–byte alignment for data (and BSS) sections. Informational sections get a default alignment of 1 byte (no alignment), though the value does not matter.

The defaults assumed by NASM if you do not specify the above qualifiers are:

```
section .text    code  align=16
section .data    data  align=4
section .rdata   rdata align=8
section .bss     bss   align=4
```

Any other section name is treated by default like .text.

6.4 coff: Common Object File Format

The coff output type produces COFF object files suitable for linking with the DJGPP linker.

coff provides a default output file−name extension of .o.

The coff format supports the same extensions to the SECTION directive as win32 does, except that the align qualifier and the info section type are not supported.

6.5 elf: Executable and Linkable Format Object Files

The elf output format generates ELF32 (Executable and Linkable Format) object files, as used by Linux as well as Unix System V, including Solaris x86, UnixWare and SCO Unix. elf provides a default output file−name extension of .o.

6.5.1 elf Extensions to the SECTION Directive

Like the obj format, elf allows you to specify additional information on the SECTION directive line, to control the type and properties of sections you declare. Section types and properties are generated automatically by NASM for the standard section names .text, .data and .bss, but may still be overridden by these qualifiers.

The available qualifiers are:

- alloc defines the section to be one which is loaded into memory when the program is run. noalloc defines it to be one which is not, such as an informational or comment section.
- exec defines the section to be one which should have execute permission when the program is run. noexec defines it as one which should not.
- write defines the section to be one which should be writable when the program is run. nowrite defines it as one which should not.
- progbits defines the section to be one with explicit contents stored in the object file: an ordinary code or data section, for example, nobits defines the section to be one with no explicit contents given, such as a BSS section.
- align=, used with a trailing number as in obj, gives the alignment requirements of the section.

The defaults assumed by NASM if you do not specify the above qualifiers are:

```
section .text    progbits  alloc  exec    nowrite  align=16
section .rodata  progbits  alloc  noexec  nowrite  align=4
section .data    progbits  alloc  noexec  write    align=4
section .bss     nobits    alloc  noexec  write    align=4
section other    progbits  alloc  noexec  nowrite  align=1
```

(Any section name other than .text, .rodata, .data and .bss is treated by default like other in the above code.)

6.5.2 Position−Independent Code: elf Special Symbols and WRT

The ELF specification contains enough features to allow position−independent code (PIC) to be written, which makes ELF shared libraries very flexible. However, it also means NASM has to be able to generate a variety of strange relocation types in ELF object files, if it is to be an assembler which can write PIC.

Since ELF does not support segment−base references, the WRT operator is not used for its normal purpose; therefore NASM's elf output format makes use of WRT for a different purpose, namely the PIC−specific relocation types.

elf defines five special symbols which you can use as the right-hand side of the WRT operator to obtain PIC relocation types. They are ..gotpc, ..gotoff, ..got, ..plt and ..sym. Their functions are summarised here:

- Referring to the symbol marking the global offset table base using wrt ..gotpc will end up giving the distance from the beginning of the current section to the global offset table. (_GLOBAL_OFFSET_TABLE_ is the standard symbol name used to refer to the GOT.) So you would then need to add $$ to the result to get the real address of the GOT.
- Referring to a location in one of your own sections using wrt ..gotoff will give the distance from the beginning of the GOT to the specified location, so that adding on the address of the GOT would give the real address of the location you wanted.
- Referring to an external or global symbol using wrt ..got causes the linker to build an entry *in* the GOT containing the address of the symbol, and the reference gives the distance from the beginning of the GOT to the entry; so you can add on the address of the GOT, load from the resulting address, and end up with the address of the symbol.
- Referring to a procedure name using wrt ..plt causes the linker to build a procedure linkage table entry for the symbol, and the reference gives the address of the PLT entry. You can only use this in contexts which would generate a PC-relative relocation normally (i.e. as the destination for CALL or JMP), since ELF contains no relocation type to refer to PLT entries absolutely.
- Referring to a symbol name using wrt ..sym causes NASM to write an ordinary relocation, but instead of making the relocation relative to the start of the section and then adding on the offset to the symbol, it will write a relocation record aimed directly at the symbol in question. The distinction is a necessary one due to a peculiarity of the dynamic linker.

A fuller explanation of how to use these relocation types to write shared libraries entirely in NASM is given in section 8.2.

6.5.3 elf Extensions to the GLOBAL Directive

ELF object files can contain more information about a global symbol than just its address: they can contain the size of the symbol and its type as well. These are not merely debugger conveniences, but are actually necessary when the program being written is a shared library. NASM therefore supports some extensions to the GLOBAL directive, allowing you to specify these features.

You can specify whether a global variable is a function or a data object by suffixing the name with a colon and the word function or data. (object is a synonym for data.) For example:

```
global   hashlookup:function, hashtable:data
```

exports the global symbol hashlookup as a function and hashtable as a data object.

You can also specify the size of the data associated with the symbol, as a numeric expression (which may involve labels, and even forward references) after the type specifier. Like this:

```
global  hashtable:data (hashtable.end - hashtable)

hashtable:
        db this,that,theother  ; some data here
.end:
```

This makes NASM automatically calculate the length of the table and place that information into the ELF symbol table.

Declaring the type and size of global symbols is necessary when writing shared library code. For more information, see section 8.2.4.

6.5.4 `elf` Extensions to the `COMMON` Directive

`ELF` also allows you to specify alignment requirements on common variables. This is done by putting a number (which must be a power of two) after the name and size of the common variable, separated (as usual) by a colon. For example, an array of doublewords would benefit from 4–byte alignment:

```
common  dwordarray 128:4
```

This declares the total size of the array to be 128 bytes, and requires that it be aligned on a 4–byte boundary.

6.5.5 16–bit code and ELF

The `ELF32` specification doesn't provide relocations for 8– and 16–bit values, but the GNU `ld` linker adds these as an extension. NASM can generate GNU–compatible relocations, to allow 16–bit code to be linked as ELF using GNU `ld`. If NASM is used with the `–w+gnu–elf–extensions` option, a warning is issued when one of these relocations is generated.

6.6 `aout`: Linux `a.out` Object Files

The `aout` format generates `a.out` object files, in the form used by early Linux systems (current Linux systems use ELF, see section 6.5.) These differ from other `a.out` object files in that the magic number in the first four bytes of the file is different; also, some implementations of `a.out`, for example NetBSD's, support position–independent code, which Linux's implementation does not.

`a.out` provides a default output file–name extension of `.o`.

`a.out` is a very simple object format. It supports no special directives, no special symbols, no use of `SEG` or `WRT`, and no extensions to any standard directives. It supports only the three standard section names `.text`, `.data` and `.bss`.

6.7 `aoutb`: NetBSD/FreeBSD/OpenBSD `a.out` Object Files

The `aoutb` format generates `a.out` object files, in the form used by the various free `BSD Unix` clones, `NetBSD`, `FreeBSD` and `OpenBSD`. For simple object files, this object format is exactly the same as `aout` except for the magic number in the first four bytes of the file. However, the `aoutb` format supports position–independent code in the same way as the `elf` format, so you can use it to write `BSD` shared libraries.

`aoutb` provides a default output file–name extension of `.o`.

`aoutb` supports no special directives, no special symbols, and only the three standard section names `.text`, `.data` and `.bss`. However, it also supports the same use of `WRT` as `elf` does, to provide position–independent code relocation types. See section 6.5.2 for full documentation of this feature.

`aoutb` also supports the same extensions to the `GLOBAL` directive as `elf` does: see section 6.5.3 for documentation of this.

6.8 `as86`: Minix/Linux `as86` Object Files

The Minix/Linux 16–bit assembler `as86` has its own non–standard object file format. Although its companion linker `ld86` produces something close to ordinary `a.out` binaries as output, the object file format used to communicate between `as86` and `ld86` is not itself `a.out`.

NASM supports this format, just in case it is useful, as `as86`. `as86` provides a default output file–name extension of `.o`.

as86 is a very simple object format (from the NASM user's point of view). It supports no special directives, no special symbols, no use of SEG or WRT, and no extensions to any standard directives. It supports only the three standard section names .text, .data and .bss.

6.9 rdf: Relocatable Dynamic Object File Format

The rdf output format produces RDOFF object files. RDOFF (Relocatable Dynamic Object File Format) is a home−grown object−file format, designed alongside NASM itself and reflecting in its file format the internal structure of the assembler.

RDOFF is not used by any well−known operating systems. Those writing their own systems, however, may well wish to use RDOFF as their object format, on the grounds that it is designed primarily for simplicity and contains very little file−header bureaucracy.

The Unix NASM archive, and the DOS archive which includes sources, both contain an rdoff subdirectory holding a set of RDOFF utilities: an RDF linker, an RDF static−library manager, an RDF file dump utility, and a program which will load and execute an RDF executable under Linux.

rdf supports only the standard section names .text, .data and .bss.

6.9.1 Requiring a Library: The LIBRARY Directive

RDOFF contains a mechanism for an object file to demand a given library to be linked to the module, either at load time or run time. This is done by the LIBRARY directive, which takes one argument which is the name of the module:

```
library  mylib.rdl
```

6.9.2 Specifying a Module Name: The MODULE Directive

Special RDOFF header record is used to store the name of the module. It can be used, for example, by run−time loader to perform dynamic linking. MODULE directive takes one argument which is the name of current module:

```
module  mymodname
```

Note that when you statically link modules and tell linker to strip the symbols from output file, all module names will be stripped too. To avoid it, you should start module names with $, like:

```
module  $kernel.core
```

6.9.3 rdf Extensions to the GLOBAL directive

RDOFF global symbols can contain additional information needed by the static linker. You can mark a global symbol as exported, thus telling the linker do not strip it from target executable or library file. Like in ELF, you can also specify whether an exported symbol is a procedure (function) or data object.

Suffixing the name with a colon and the word export you make the symbol exported:

```
global  sys_open:export
```

To specify that exported symbol is a procedure (function), you add the word proc or function after declaration:

```
global  sys_open:export proc
```

Similarly, to specify exported data object, add the word data or object to the directive:

```
global  kernel_ticks:export data
```

6.9.4 `rdf` Extensions to the `EXTERN` directive

By default the `EXTERN` directive in `RDOFF` declares a "pure external" symbol (i.e. the static linker will complain if such a symbol is not resolved). To declare an "imported" symbol, which must be resolved later during a dynamic linking phase, `RDOFF` offers an additional `import` modifier. As in `GLOBAL`, you can also specify whether an imported symbol is a procedure (function) or data object. For example:

```
library $libc
extern  _open:import
extern  _printf:import proc
extern  _errno:import data
```

Here the directive `LIBRARY` is also included, which gives the dynamic linker a hint as to where to find requested symbols.

6.10 `dbg`: Debugging Format

The `dbg` output format is not built into NASM in the default configuration. If you are building your own NASM executable from the sources, you can define `OF_DBG` in `outform.h` or on the compiler command line, and obtain the `dbg` output format.

The `dbg` format does not output an object file as such; instead, it outputs a text file which contains a complete list of all the transactions between the main body of NASM and the output–format back end module. It is primarily intended to aid people who want to write their own output drivers, so that they can get a clearer idea of the various requests the main program makes of the output driver, and in what order they happen.

For simple files, one can easily use the `dbg` format like this:

```
nasm -f dbg filename.asm
```

which will generate a diagnostic file called `filename.dbg`. However, this will not work well on files which were designed for a different object format, because each object format defines its own macros (usually user–level forms of directives), and those macros will not be defined in the `dbg` format. Therefore it can be useful to run NASM twice, in order to do the preprocessing with the native object format selected:

```
nasm -e -f rdf -o rdfprog.i rdfprog.asm
nasm -a -f dbg rdfprog.i
```

This preprocesses `rdfprog.asm` into `rdfprog.i`, keeping the `rdf` object format selected in order to make sure RDF special directives are converted into primitive form correctly. Then the preprocessed source is fed through the `dbg` format to generate the final diagnostic output.

This workaround will still typically not work for programs intended for `obj` format, because the `obj SEGMENT` and `GROUP` directives have side effects of defining the segment and group names as symbols; `dbg` will not do this, so the program will not assemble. You will have to work around that by defining the symbols yourself (using `EXTERN`, for example) if you really need to get a `dbg` trace of an `obj`–specific source file.

`dbg` accepts any section name and any directives at all, and logs them all to its output file.

Chapter 7: Writing 16–bit Code (DOS, Windows 3/3.1)

This chapter attempts to cover some of the common issues encountered when writing 16–bit code to run under MS–DOS or Windows 3.x. It covers how to link programs to produce .EXE or .COM files, how to write .SYS device drivers, and how to interface assembly language code with 16–bit C compilers and with Borland Pascal.

7.1 Producing .EXE Files

Any large program written under DOS needs to be built as a .EXE file: only .EXE files have the necessary internal structure required to span more than one 64K segment. Windows programs, also, have to be built as .EXE files, since Windows does not support the .COM format.

In general, you generate .EXE files by using the obj output format to produce one or more .OBJ files, and then linking them together using a linker. However, NASM also supports the direct generation of simple DOS .EXE files using the bin output format (by using DB and DW to construct the .EXE file header), and a macro package is supplied to do this. Thanks to Yann Guidon for contributing the code for this.

NASM may also support .EXE natively as another output format in future releases.

7.1.1 Using the obj Format To Generate .EXE Files

This section describes the usual method of generating .EXE files by linking .OBJ files together.

Most 16–bit programming language packages come with a suitable linker; if you have none of these, there is a free linker called VAL, available in LZH archive format from x2ftp.oulu.fi. An LZH archiver can be found at ftp.simtel.net. There is another 'free' linker (though this one doesn't come with sources) called FREELINK, available from www.pcorner.com. A third, djlink, written by DJ Delorie, is available at www.delorie.com. A fourth linker, ALINK, written by Anthony A.J. Williams, is available at alink.sourceforge.net.

When linking several .OBJ files into a .EXE file, you should ensure that exactly one of them has a start point defined (using the ..start special symbol defined by the obj format: see section 6.2.6). If no module defines a start point, the linker will not know what value to give the entry–point field in the output file header; if more than one defines a start point, the linker will not know *which* value to use.

An example of a NASM source file which can be assembled to a .OBJ file and linked on its own to a .EXE is given here. It demonstrates the basic principles of defining a stack, initialising the segment registers, and declaring a start point. This file is also provided in the test subdirectory of the NASM archives, under the name objexe.asm.

```
segment code

..start:
        mov     ax,data
        mov     ds,ax
        mov     ax,stack
        mov     ss,ax
        mov     sp,stacktop
```

This initial piece of code sets up DS to point to the data segment, and initialises SS and SP to point to the top of the provided stack. Notice that interrupts are implicitly disabled for one instruction after a move into SS, precisely for this situation, so that there's no chance of an interrupt occurring between the loads of SS and SP and not having a stack to execute on.

Note also that the special symbol ..start is defined at the beginning of this code, which means that will be the entry point into the resulting executable file.

```
        mov     dx,hello
        mov     ah,9
        int     0x21
```

The above is the main program: load DS:DX with a pointer to the greeting message (hello is implicitly relative to the segment data, which was loaded into DS in the setup code, so the full pointer is valid), and call the DOS print−string function.

```
        mov     ax,0x4c00
        int     0x21
```

This terminates the program using another DOS system call.

```
segment data

hello:  db      'hello, world', 13, 10, '$'
```

The data segment contains the string we want to display.

```
segment stack stack
        resb 64
stacktop:
```

The above code declares a stack segment containing 64 bytes of uninitialised stack space, and points stacktop at the top of it. The directive segment stack stack defines a segment *called* stack, and also of *type* STACK. The latter is not necessary to the correct running of the program, but linkers are likely to issue warnings or errors if your program has no segment of type STACK.

The above file, when assembled into a .OBJ file, will link on its own to a valid .EXE file, which when run will print 'hello, world' and then exit.

7.1.2 Using the bin Format To Generate .EXE Files

The .EXE file format is simple enough that it's possible to build a .EXE file by writing a pure−binary program and sticking a 32−byte header on the front. This header is simple enough that it can be generated using DB and DW commands by NASM itself, so that you can use the bin output format to directly generate .EXE files.

Included in the NASM archives, in the misc subdirectory, is a file exebin.mac of macros. It defines three macros: EXE_begin, EXE_stack and EXE_end.

To produce a .EXE file using this method, you should start by using %include to load the exebin.mac macro package into your source file. You should then issue the EXE_begin macro call (which takes no arguments) to generate the file header data. Then write code as normal for the bin format – you can use all three standard sections .text, .data and .bss. At the end of the file you should call the EXE_end macro (again, no arguments), which defines some symbols to mark section sizes, and these symbols are referred to in the header code generated by EXE_begin.

In this model, the code you end up writing starts at 0x100, just like a .COM file – in fact, if you strip off the 32−byte header from the resulting .EXE file, you will have a valid .COM program. All the segment bases are the same, so you are limited to a 64K program, again just like a .COM file.

Note that an ORG directive is issued by the EXE_begin macro, so you should not explicitly issue one of your own.

You can't directly refer to your segment base value, unfortunately, since this would require a relocation in the header, and things would get a lot more complicated. So you should get your segment base by copying it out of CS instead.

On entry to your .EXE file, SS:SP are already set up to point to the top of a 2Kb stack. You can adjust the default stack size of 2Kb by calling the EXE_stack macro. For example, to change the stack size of your program to 64 bytes, you would call EXE_stack 64.

A sample program which generates a .EXE file in this way is given in the test subdirectory of the NASM archive, as binexe.asm.

7.2 Producing .COM Files

While large DOS programs must be written as .EXE files, small ones are often better written as .COM files. .COM files are pure binary, and therefore most easily produced using the bin output format.

7.2.1 Using the bin Format To Generate .COM Files

.COM files expect to be loaded at offset 100h into their segment (though the segment may change). Execution then begins at 100h, i.e. right at the start of the program. So to write a .COM program, you would create a source file looking like

```
        org 100h

section .text

start:
        ; put your code here

section .data

        ; put data items here

section .bss

        ; put uninitialised data here
```

The bin format puts the .text section first in the file, so you can declare data or BSS items before beginning to write code if you want to and the code will still end up at the front of the file where it belongs.

The BSS (uninitialised data) section does not take up space in the .COM file itself: instead, addresses of BSS items are resolved to point at space beyond the end of the file, on the grounds that this will be free memory when the program is run. Therefore you should not rely on your BSS being initialised to all zeros when you run.

To assemble the above program, you should use a command line like

```
nasm myprog.asm -fbin -o myprog.com
```

The bin format would produce a file called myprog if no explicit output file name were specified, so you have to override it and give the desired file name.

7.2.2 Using the `obj` Format To Generate `.COM` Files

If you are writing a `.COM` program as more than one module, you may wish to assemble several `.OBJ` files and link them together into a `.COM` program. You can do this, provided you have a linker capable of outputting `.COM` files directly (TLINK does this), or alternatively a converter program such as `EXE2BIN` to transform the `.EXE` file output from the linker into a `.COM` file.

If you do this, you need to take care of several things:

- The first object file containing code should start its code segment with a line like `RESB 100h`. This is to ensure that the code begins at offset `100h` relative to the beginning of the code segment, so that the linker or converter program does not have to adjust address references within the file when generating the `.COM` file. Other assemblers use an `ORG` directive for this purpose, but `ORG` in NASM is a format−specific directive to the `bin` output format, and does not mean the same thing as it does in MASM−compatible assemblers.
- You don't need to define a stack segment.
- All your segments should be in the same group, so that every time your code or data references a symbol offset, all offsets are relative to the same segment base. This is because, when a `.COM` file is loaded, all the segment registers contain the same value.

7.3 Producing `.SYS` Files

MS−DOS device drivers – `.SYS` files – are pure binary files, similar to `.COM` files, except that they start at origin zero rather than `100h`. Therefore, if you are writing a device driver using the `bin` format, you do not need the `ORG` directive, since the default origin for `bin` is zero. Similarly, if you are using `obj`, you do not need the `RESB 100h` at the start of your code segment.

`.SYS` files start with a header structure, containing pointers to the various routines inside the driver which do the work. This structure should be defined at the start of the code segment, even though it is not actually code.

For more information on the format of `.SYS` files, and the data which has to go in the header structure, a list of books is given in the Frequently Asked Questions list for the newsgroup `comp.os.msdos.programmer`.

7.4 Interfacing to 16−bit C Programs

This section covers the basics of writing assembly routines that call, or are called from, C programs. To do this, you would typically write an assembly module as a `.OBJ` file, and link it with your C modules to produce a mixed−language program.

7.4.1 External Symbol Names

C compilers have the convention that the names of all global symbols (functions or data) they define are formed by prefixing an underscore to the name as it appears in the C program. So, for example, the function a C programmer thinks of as `printf` appears to an assembly language programmer as `_printf`. This means that in your assembly programs, you can define symbols without a leading underscore, and not have to worry about name clashes with C symbols.

If you find the underscores inconvenient, you can define macros to replace the `GLOBAL` and `EXTERN` directives as follows:

```
%macro  cglobal 1

  global  _%1
  %define %1 _%1
```

```
%endmacro

%macro  cextern 1

  extern  _%1
  %define %1 _%1

%endmacro
```

(These forms of the macros only take one argument at a time; a `%rep` construct could solve this.)

If you then declare an external like this:

```
cextern printf
```

then the macro will expand it as

```
extern  _printf
%define printf _printf
```

Thereafter, you can reference `printf` as if it was a symbol, and the preprocessor will put the leading underscore on where necessary.

The `cglobal` macro works similarly. You must use `cglobal` before defining the symbol in question, but you would have had to do that anyway if you used `GLOBAL`.

Also see section 2.1.21.

7.4.2 Memory Models

NASM contains no mechanism to support the various C memory models directly; you have to keep track yourself of which one you are writing for. This means you have to keep track of the following things:

- In models using a single code segment (tiny, small and compact), functions are near. This means that function pointers, when stored in data segments or pushed on the stack as function arguments, are 16 bits long and contain only an offset field (the `CS` register never changes its value, and always gives the segment part of the full function address), and that functions are called using ordinary near `CALL` instructions and return using `RETN` (which, in NASM, is synonymous with `RET` anyway). This means both that you should write your own routines to return with `RETN`, and that you should call external C routines with near `CALL` instructions.
- In models using more than one code segment (medium, large and huge), functions are far. This means that function pointers are 32 bits long (consisting of a 16−bit offset followed by a 16−bit segment), and that functions are called using `CALL FAR` (or `CALL seg:offset`) and return using `RETF`. Again, you should therefore write your own routines to return with `RETF` and use `CALL FAR` to call external routines.
- In models using a single data segment (tiny, small and medium), data pointers are 16 bits long, containing only an offset field (the `DS` register doesn't change its value, and always gives the segment part of the full data item address).
- In models using more than one data segment (compact, large and huge), data pointers are 32 bits long, consisting of a 16−bit offset followed by a 16−bit segment. You should still be careful not to modify `DS` in your routines without restoring it afterwards, but `ES` is free for you to use to access the contents of 32−bit data pointers you are passed.
- The huge memory model allows single data items to exceed 64K in size. In all other memory models, you can access the whole of a data item just by doing arithmetic on the offset field of the

pointer you are given, whether a segment field is present or not; in huge model, you have to be more careful of your pointer arithmetic.

- In most memory models, there is a *default* data segment, whose segment address is kept in `DS` throughout the program. This data segment is typically the same segment as the stack, kept in `SS`, so that functions' local variables (which are stored on the stack) and global data items can both be accessed easily without changing `DS`. Particularly large data items are typically stored in other segments. However, some memory models (though not the standard ones, usually) allow the assumption that `SS` and `DS` hold the same value to be removed. Be careful about functions' local variables in this latter case.

In models with a single code segment, the segment is called `_TEXT`, so your code segment must also go by this name in order to be linked into the same place as the main code segment. In models with a single data segment, or with a default data segment, it is called `_DATA`.

7.4.3 Function Definitions and Function Calls

The C calling convention in 16–bit programs is as follows. In the following description, the words *caller* and *callee* are used to denote the function doing the calling and the function which gets called.

- The caller pushes the function's parameters on the stack, one after another, in reverse order (right to left, so that the first argument specified to the function is pushed last).
- The caller then executes a `CALL` instruction to pass control to the callee. This `CALL` is either near or far depending on the memory model.
- The callee receives control, and typically (although this is not actually necessary, in functions which do not need to access their parameters) starts by saving the value of `SP` in `BP` so as to be able to use `BP` as a base pointer to find its parameters on the stack. However, the caller was probably doing this too, so part of the calling convention states that `BP` must be preserved by any C function. Hence the callee, if it is going to set up `BP` as a *frame pointer*, must push the previous value first.
- The callee may then access its parameters relative to `BP`. The word at `[BP]` holds the previous value of `BP` as it was pushed; the next word, at `[BP+2]`, holds the offset part of the return address, pushed implicitly by `CALL`. In a small–model (near) function, the parameters start after that, at `[BP+4]`; in a large–model (far) function, the segment part of the return address lives at `[BP+4]`, and the parameters begin at `[BP+6]`. The leftmost parameter of the function, since it was pushed last, is accessible at this offset from `BP`; the others follow, at successively greater offsets. Thus, in a function such as `printf` which takes a variable number of parameters, the pushing of the parameters in reverse order means that the function knows where to find its first parameter, which tells it the number and type of the remaining ones.
- The callee may also wish to decrease `SP` further, so as to allocate space on the stack for local variables, which will then be accessible at negative offsets from `BP`.
- The callee, if it wishes to return a value to the caller, should leave the value in `AL`, `AX` or `DX:AX` depending on the size of the value. Floating–point results are sometimes (depending on the compiler) returned in `ST0`.
- Once the callee has finished processing, it restores `SP` from `BP` if it had allocated local stack space, then pops the previous value of `BP`, and returns via `RETN` or `RETF` depending on memory model.
- When the caller regains control from the callee, the function parameters are still on the stack, so it typically adds an immediate constant to `SP` to remove them (instead of executing a number of slow `POP` instructions). Thus, if a function is accidentally called with the wrong number of parameters due to a prototype mismatch, the stack will still be returned to a sensible state since the caller, which *knows* how many parameters it pushed, does the removing.

It is instructive to compare this calling convention with that for Pascal programs (described in section 7.5.1). Pascal has a simpler convention, since no functions have variable numbers of parameters. Therefore the callee knows how many parameters it should have been passed, and is able to deallocate them from the stack itself by passing an immediate argument to the `RET` or `RETF` instruction, so the caller does not have to do it. Also, the parameters are pushed in left–to–right order, not right–to–left, which means that a compiler can give better guarantees about sequence points without performance suffering.

Thus, you would define a function in C style in the following way. The following example is for small model:

```
global  _myfunc

_myfunc:
        push    bp
        mov     bp,sp
        sub     sp,0x40         ; 64 bytes of local stack space
        mov     bx,[bp+4]       ; first parameter to function

        ; some more code

        mov     sp,bp           ; undo "sub sp,0x40" above
        pop     bp
        ret
```

For a large–model function, you would replace `RET` by `RETF`, and look for the first parameter at `[BP+6]` instead of `[BP+4]`. Of course, if one of the parameters is a pointer, then the offsets of *subsequent* parameters will change depending on the memory model as well: far pointers take up four bytes on the stack when passed as a parameter, whereas near pointers take up two.

At the other end of the process, to call a C function from your assembly code, you would do something like this:

```
extern  _printf

      ; and then, further down...

      push    word [myint]        ; one of my integer variables
      push    word mystring       ; pointer into my data segment
      call    _printf
      add     sp,byte 4           ; 'byte' saves space

      ; then those data items...

segment _DATA

myint         dw    1234
mystring      db    'This number -> %d <- should be 1234',10,0
```

This piece of code is the small–model assembly equivalent of the C code

```
    int myint = 1234;
    printf("This number -> %d <- should be 1234\\n", myint);
```

In large model, the function–call code might look more like this. In this example, it is assumed that `DS` already holds the segment base of the segment `_DATA`. If not, you would have to initialise it first.

```
push    word [myint]
push    word seg mystring   ; Now push the segment, and...
push    word mystring       ; ... offset of "mystring"
call    far _printf
add    sp,byte 6
```

The integer value still takes up one word on the stack, since large model does not affect the size of the `int` data type. The first argument (pushed last) to `printf`, however, is a data pointer, and therefore has to contain a segment and offset part. The segment should be stored second in memory, and therefore must be pushed first. (Of course, `PUSH DS` would have been a shorter instruction than `PUSH WORD SEG mystring`, if `DS` was set up as the above example assumed.) Then the actual call becomes a far call, since functions expect far calls in large model; and `SP` has to be increased by 6 rather than 4 afterwards to make up for the extra word of parameters.

7.4.4 Accessing Data Items

To get at the contents of C variables, or to declare variables which C can access, you need only declare the names as `GLOBAL` or `EXTERN`. (Again, the names require leading underscores, as stated in section 7.4.1.) Thus, a C variable declared as `int i` can be accessed from assembler as

```
extern _i

        mov ax,[_i]
```

And to declare your own integer variable which C programs can access as `extern int j`, you do this (making sure you are assembling in the `_DATA` segment, if necessary):

```
global  _j

_j      dw      0
```

To access a C array, you need to know the size of the components of the array. For example, `int` variables are two bytes long, so if a C program declares an array as `int a[10]`, you can access `a[3]` by coding `mov ax,[_a+6]`. (The byte offset 6 is obtained by multiplying the desired array index, 3, by the size of the array element, 2.) The sizes of the C base types in 16–bit compilers are: 1 for `char`, 2 for `short` and `int`, 4 for `long` and `float`, and 8 for `double`.

To access a C data structure, you need to know the offset from the base of the structure to the field you are interested in. You can either do this by converting the C structure definition into a NASM structure definition (using `STRUC`), or by calculating the one offset and using just that.

To do either of these, you should read your C compiler's manual to find out how it organises data structures. NASM gives no special alignment to structure members in its own `STRUC` macro, so you have to specify alignment yourself if the C compiler generates it. Typically, you might find that a structure like

```
struct @\{
    char c;
    int i;
@\} foo;
```

might be four bytes long rather than three, since the `int` field would be aligned to a two–byte boundary. However, this sort of feature tends to be a configurable option in the C compiler, either using command–line options or `#pragma` lines, so you have to find out how your own compiler does it.

7.4.5 `c16.mac`: Helper Macros for the 16–bit C Interface

Included in the NASM archives, in the `misc` directory, is a file `c16.mac` of macros. It defines three macros: `proc`, `arg` and `endproc`. These are intended to be used for C–style procedure definitions, and they automate a lot of the work involved in keeping track of the calling convention.

(An alternative, TASM compatible form of `arg` is also now built into NASM's preprocessor. See section 4.9 for details.)

An example of an assembly function using the macro set is given here:

```
proc    _nearproc

%$i     arg
%$j     arg
        mov     ax,[bp + %$i]
        mov     bx,[bp + %$j]
        add     ax,[bx]

endproc
```

This defines `_nearproc` to be a procedure taking two arguments, the first (`i`) an integer and the second (`j`) a pointer to an integer. It returns `i + *j`.

Note that the `arg` macro has an `EQU` as the first line of its expansion, and since the label before the macro call gets prepended to the first line of the expanded macro, the `EQU` works, defining `%$i` to be an offset from `BP`. A context–local variable is used, local to the context pushed by the `proc` macro and popped by the `endproc` macro, so that the same argument name can be used in later procedures. Of course, you don't *have* to do that.

The macro set produces code for near functions (tiny, small and compact–model code) by default. You can have it generate far functions (medium, large and huge–model code) by means of coding `%define FARCODE`. This changes the kind of return instruction generated by `endproc`, and also changes the starting point for the argument offsets. The macro set contains no intrinsic dependency on whether data pointers are far or not.

`arg` can take an optional parameter, giving the size of the argument. If no size is given, 2 is assumed, since it is likely that many function parameters will be of type `int`.

The large–model equivalent of the above function would look like this:

```
%define FARCODE

proc    _farproc

%$i     arg
%$j     arg     4
        mov     ax,[bp + %$i]
        mov     bx,[bp + %$j]
        mov     es,[bp + %$j + 2]
        add     ax,[bx]

endproc
```

This makes use of the argument to the `arg` macro to define a parameter of size 4, because `j` is now a far pointer. When we load from `j`, we must load a segment and an offset.

7.5 Interfacing to Borland Pascal Programs

Interfacing to Borland Pascal programs is similar in concept to interfacing to 16–bit C programs. The differences are:

- The leading underscore required for interfacing to C programs is not required for Pascal.
- The memory model is always large: functions are far, data pointers are far, and no data item can be more than 64K long. (Actually, some functions are near, but only those functions that are local to a Pascal unit and never called from outside it. All assembly functions that Pascal calls, and all Pascal functions that assembly routines are able to call, are far.) However, all static data declared in a Pascal program goes into the default data segment, which is the one whose segment address will be in `DS` when control is passed to your assembly code. The only things that do not live in the default data segment are local variables (they live in the stack segment) and dynamically allocated variables. All data *pointers*, however, are far.
- The function calling convention is different – described below.
- Some data types, such as strings, are stored differently.
- There are restrictions on the segment names you are allowed to use – Borland Pascal will ignore code or data declared in a segment it doesn't like the name of. The restrictions are described below.

7.5.1 The Pascal Calling Convention

The 16–bit Pascal calling convention is as follows. In the following description, the words *caller* and *callee* are used to denote the function doing the calling and the function which gets called.

- The caller pushes the function's parameters on the stack, one after another, in normal order (left to right, so that the first argument specified to the function is pushed first).
- The caller then executes a far `CALL` instruction to pass control to the callee.
- The callee receives control, and typically (although this is not actually necessary, in functions which do not need to access their parameters) starts by saving the value of `SP` in `BP` so as to be able to use `BP` as a base pointer to find its parameters on the stack. However, the caller was probably doing this too, so part of the calling convention states that `BP` must be preserved by any function. Hence the callee, if it is going to set up `BP` as a frame pointer, must push the previous value first.
- The callee may then access its parameters relative to `BP`. The word at `[BP]` holds the previous value of `BP` as it was pushed. The next word, at `[BP+2]`, holds the offset part of the return address, and the next one at `[BP+4]` the segment part. The parameters begin at `[BP+6]`. The rightmost parameter of the function, since it was pushed last, is accessible at this offset from `BP`; the others follow, at successively greater offsets.
- The callee may also wish to decrease `SP` further, so as to allocate space on the stack for local variables, which will then be accessible at negative offsets from `BP`.
- The callee, if it wishes to return a value to the caller, should leave the value in `AL`, `AX` or `DX:AX` depending on the size of the value. Floating–point results are returned in `ST0`. Results of type `Real` (Borland's own custom floating–point data type, not handled directly by the FPU) are returned in `DX:BX:AX`. To return a result of type `String`, the caller pushes a pointer to a temporary string before pushing the parameters, and the callee places the returned string value at that location. The pointer is not a parameter, and should not be removed from the stack by the `RETF` instruction.
- Once the callee has finished processing, it restores `SP` from `BP` if it had allocated local stack space, then pops the previous value of `BP`, and returns via `RETF`. It uses the form of `RETF` with

an immediate parameter, giving the number of bytes taken up by the parameters on the stack. This causes the parameters to be removed from the stack as a side effect of the return instruction.

- When the caller regains control from the callee, the function parameters have already been removed from the stack, so it needs to do nothing further.

Thus, you would define a function in Pascal style, taking two `Integer`–type parameters, in the following way:

```
global  myfunc

myfunc: push    bp
        mov     bp,sp
        sub     sp,0x40         ; 64 bytes of local stack space
        mov     bx,[bp+8]       ; first parameter to function
        mov     bx,[bp+6]       ; second parameter to function

        ; some more code

        mov     sp,bp           ; undo "sub sp,0x40" above
        pop     bp
        retf    4               ; total size of params is 4
```

At the other end of the process, to call a Pascal function from your assembly code, you would do something like this:

```
extern  SomeFunc

       ; and then, further down...

       push   word seg mystring   ; Now push the segment, and...
       push   word mystring       ; ... offset of "mystring"
       push   word [myint]        ; one of my variables
       call   far SomeFunc
```

This is equivalent to the Pascal code

```
procedure SomeFunc(String: PChar; Int: Integer);
    SomeFunc(@@mystring, myint);
```

7.5.2 Borland Pascal Segment Name Restrictions

Since Borland Pascal's internal unit file format is completely different from `OBJ`, it only makes a very sketchy job of actually reading and understanding the various information contained in a real `OBJ` file when it links that in. Therefore an object file intended to be linked to a Pascal program must obey a number of restrictions:

- Procedures and functions must be in a segment whose name is either `CODE`, `CSEG`, or something ending in `_TEXT`.
- Initialised data must be in a segment whose name is either `CONST` or something ending in `_DATA`.
- Uninitialised data must be in a segment whose name is either `DATA`, `DSEG`, or something ending in `_BSS`.
- Any other segments in the object file are completely ignored. `GROUP` directives and segment attributes are also ignored.

7.5.3 Using `c16.mac` With Pascal Programs

The `c16.mac` macro package, described in section 7.4.5, can also be used to simplify writing functions to be called from Pascal programs, if you code `%define PASCAL`. This definition ensures that functions are far (it implies `FARCODE`), and also causes procedure return instructions to be generated with an operand.

Defining `PASCAL` does not change the code which calculates the argument offsets; you must declare your function's arguments in reverse order. For example:

```
%define PASCAL

proc    _pascalproc

%$j     arg 4
%$i     arg
        mov     ax,[bp + %$i]
        mov     bx,[bp + %$j]
        mov     es,[bp + %$j + 2]
        add     ax,[bx]

endproc
```

This defines the same routine, conceptually, as the example in section 7.4.5: it defines a function taking two arguments, an integer and a pointer to an integer, which returns the sum of the integer and the contents of the pointer. The only difference between this code and the large–model C version is that `PASCAL` is defined instead of `FARCODE`, and that the arguments are declared in reverse order.

Chapter 8: Writing 32−bit Code (Unix, Win32, DJGPP)

This chapter attempts to cover some of the common issues involved when writing 32−bit code, to run under Win32 or Unix, or to be linked with C code generated by a Unix−style C compiler such as DJGPP. It covers how to write assembly code to interface with 32−bit C routines, and how to write position−independent code for shared libraries.

Almost all 32−bit code, and in particular all code running under `Win32`, `DJGPP` or any of the PC Unix variants, runs in *flat* memory model. This means that the segment registers and paging have already been set up to give you the same 32−bit 4Gb address space no matter what segment you work relative to, and that you should ignore all segment registers completely. When writing flat−model application code, you never need to use a segment override or modify any segment register, and the code−section addresses you pass to `CALL` and `JMP` live in the same address space as the data−section addresses you access your variables by and the stack−section addresses you access local variables and procedure parameters by. Every address is 32 bits long and contains only an offset part.

8.1 Interfacing to 32−bit C Programs

A lot of the discussion in section 7.4, about interfacing to 16−bit C programs, still applies when working in 32 bits. The absence of memory models or segmentation worries simplifies things a lot.

8.1.1 External Symbol Names

Most 32−bit C compilers share the convention used by 16−bit compilers, that the names of all global symbols (functions or data) they define are formed by prefixing an underscore to the name as it appears in the C program. However, not all of them do: the `ELF` specification states that C symbols do *not* have a leading underscore on their assembly−language names.

The older Linux `a.out` C compiler, all `Win32` compilers, `DJGPP`, and `NetBSD` and `FreeBSD`, all use the leading underscore; for these compilers, the macros `cextern` and `cglobal`, as given in section 7.4.1, will still work. For `ELF`, though, the leading underscore should not be used.

See also section 2.1.21.

8.1.2 Function Definitions and Function Calls

The C calling conventionThe C calling convention in 32−bit programs is as follows. In the following description, the words *caller* and *callee* are used to denote the function doing the calling and the function which gets called.

- The caller pushes the function's parameters on the stack, one after another, in reverse order (right to left, so that the first argument specified to the function is pushed last).
- The caller then executes a near `CALL` instruction to pass control to the callee.
- The callee receives control, and typically (although this is not actually necessary, in functions which do not need to access their parameters) starts by saving the value of `ESP` in `EBP` so as to be able to use `EBP` as a base pointer to find its parameters on the stack. However, the caller was probably doing this too, so part of the calling convention states that `EBP` must be preserved by any C function. Hence the callee, if it is going to set up `EBP` as a frame pointer, must push the previous value first.

- The callee may then access its parameters relative to EBP. The doubleword at [EBP] holds the previous value of EBP as it was pushed; the next doubleword, at [EBP+4], holds the return address, pushed implicitly by CALL. The parameters start after that, at [EBP+8]. The leftmost parameter of the function, since it was pushed last, is accessible at this offset from EBP; the others follow, at successively greater offsets. Thus, in a function such as printf which takes a variable number of parameters, the pushing of the parameters in reverse order means that the function knows where to find its first parameter, which tells it the number and type of the remaining ones.
- The callee may also wish to decrease ESP further, so as to allocate space on the stack for local variables, which will then be accessible at negative offsets from EBP.
- The callee, if it wishes to return a value to the caller, should leave the value in AL, AX or EAX depending on the size of the value. Floating−point results are typically returned in ST0.
- Once the callee has finished processing, it restores ESP from EBP if it had allocated local stack space, then pops the previous value of EBP, and returns via RET (equivalently, RETN).
- When the caller regains control from the callee, the function parameters are still on the stack, so it typically adds an immediate constant to ESP to remove them (instead of executing a number of slow POP instructions). Thus, if a function is accidentally called with the wrong number of parameters due to a prototype mismatch, the stack will still be returned to a sensible state since the caller, which *knows* how many parameters it pushed, does the removing.

There is an alternative calling convention used by Win32 programs for Windows API calls, and also for functions called *by* the Windows API such as window procedures: they follow what Microsoft calls the __stdcall convention. This is slightly closer to the Pascal convention, in that the callee clears the stack by passing a parameter to the RET instruction. However, the parameters are still pushed in right−to−left order.

Thus, you would define a function in C style in the following way:

```
global  _myfunc

_myfunc:
        push    ebp
        mov     ebp,esp
        sub     esp,0x40        ; 64 bytes of local stack space
        mov     ebx,[ebp+8]     ; first parameter to function

        ; some more code

        leave                   ; mov esp,ebp / pop ebp
        ret
```

At the other end of the process, to call a C function from your assembly code, you would do something like this:

```
extern  _printf

        ; and then, further down...

        push    dword [myint]   ; one of my integer variables
        push    dword mystring  ; pointer into my data segment
        call    _printf
        add     esp,byte 8      ; 'byte' saves space
```

```
        ; then those data items...

segment _DATA

myint       dd   1234
mystring    db   'This number -> %d <- should be 1234',10,0
```

This piece of code is the assembly equivalent of the C code

```
int myint = 1234;
printf("This number -> %d <- should be 1234\\n", myint);
```

8.1.3 Accessing Data Items

To get at the contents of C variables, or to declare variables which C can access, you need only declare the names as `GLOBAL` or `EXTERN`. (Again, the names require leading underscores, as stated in section 8.1.1.) Thus, a C variable declared as `int i` can be accessed from assembler as

```
          extern _i
          mov eax,[_i]
```

And to declare your own integer variable which C programs can access as `extern int j`, you do this (making sure you are assembling in the `_DATA` segment, if necessary):

```
          global _j
_j        dd 0
```

To access a C array, you need to know the size of the components of the array. For example, `int` variables are four bytes long, so if a C program declares an array as `int a[10]`, you can access `a[3]` by coding `mov ax,[_a+12]`. (The byte offset 12 is obtained by multiplying the desired array index, 3, by the size of the array element, 4.) The sizes of the C base types in 32–bit compilers are: 1 for `char`, 2 for `short`, 4 for `int`, `long` and `float`, and 8 for `double`. Pointers, being 32–bit addresses, are also 4 bytes long.

To access a C data structure, you need to know the offset from the base of the structure to the field you are interested in. You can either do this by converting the C structure definition into a NASM structure definition (using `STRUC`), or by calculating the one offset and using just that.

To do either of these, you should read your C compiler's manual to find out how it organises data structures. NASM gives no special alignment to structure members in its own `STRUC` macro, so you have to specify alignment yourself if the C compiler generates it. Typically, you might find that a structure like

```
struct @\{
    char c;
    int i;
@\} foo;
```

might be eight bytes long rather than five, since the `int` field would be aligned to a four–byte boundary. However, this sort of feature is sometimes a configurable option in the C compiler, either using command–line options or `#pragma` lines, so you have to find out how your own compiler does it.

8.1.4 `c32.mac`: Helper Macros for the 32–bit C Interface

Included in the NASM archives, in the `misc` directory, is a file `c32.mac` of macros. It defines three macros: `proc`, `arg` and `endproc`. These are intended to be used for C–style procedure definitions, and they automate a lot of the work involved in keeping track of the calling convention.

An example of an assembly function using the macro set is given here:

```
proc    _proc32

%$i     arg
%$j     arg
        mov     eax,[ebp + %$i]
        mov     ebx,[ebp + %$j]
        add     eax,[ebx]

endproc
```

This defines `_proc32` to be a procedure taking two arguments, the first (`i`) an integer and the second (`j`) a pointer to an integer. It returns `i + *j`.

Note that the `arg` macro has an `EQU` as the first line of its expansion, and since the label before the macro call gets prepended to the first line of the expanded macro, the `EQU` works, defining `%$i` to be an offset from `BP`. A context–local variable is used, local to the context pushed by the `proc` macro and popped by the `endproc` macro, so that the same argument name can be used in later procedures. Of course, you don't *have* to do that.

`arg` can take an optional parameter, giving the size of the argument. If no size is given, 4 is assumed, since it is likely that many function parameters will be of type `int` or pointers.

8.2 Writing NetBSD/FreeBSD/OpenBSD and Linux/ELF Shared Libraries

`ELF` replaced the older `a.out` object file format under Linux because it contains support for position–independent code (PIC), which makes writing shared libraries much easier. NASM supports the `ELF` position–independent code features, so you can write Linux `ELF` shared libraries in NASM.

NetBSD, and its close cousins FreeBSD and OpenBSD, take a different approach by hacking PIC support into the `a.out` format. NASM supports this as the `aoutb` output format, so you can write BSD shared libraries in NASM too.

The operating system loads a PIC shared library by memory–mapping the library file at an arbitrarily chosen point in the address space of the running process. The contents of the library's code section must therefore not depend on where it is loaded in memory.

Therefore, you cannot get at your variables by writing code like this:

```
        mov     eax,[myvar]             ; WRONG
```

Instead, the linker provides an area of memory called the *global offset table*, or GOT; the GOT is situated at a constant distance from your library's code, so if you can find out where your library is loaded (which is typically done using a `CALL` and `POP` combination), you can obtain the address of the GOT, and you can then load the addresses of your variables out of linker–generated entries in the GOT.

The *data* section of a PIC shared library does not have these restrictions: since the data section is writable, it has to be copied into memory anyway rather than just paged in from the library file, so as long as it's being copied it can be relocated too. So you can put ordinary types of relocation in the data section without too much worry (but see section 8.2.4 for a caveat).

8.2.1 Obtaining the Address of the GOT

Each code module in your shared library should define the GOT as an external symbol:

```
extern  _GLOBAL_OFFSET_TABLE_   ; in ELF
extern  __GLOBAL_OFFSET_TABLE_  ; in BSD a.out
```

At the beginning of any function in your shared library which plans to access your data or BSS sections, you must first calculate the address of the GOT. This is typically done by writing the function in this form:

```
func:   push    ebp
        mov     ebp,esp
        push    ebx
        call    .get_GOT
.get_GOT:
        pop     ebx
        add     ebx,_GLOBAL_OFFSET_TABLE_+$$-.get_GOT wrt ..gotpc

        ; the function body comes here

        mov     ebx,[ebp-4]
        mov     esp,ebp
        pop     ebp
        ret
```

(For BSD, again, the symbol `_GLOBAL_OFFSET_TABLE` requires a second leading underscore.)

The first two lines of this function are simply the standard C prologue to set up a stack frame, and the last three lines are standard C function epilogue. The third line, and the fourth to last line, save and restore the `EBX` register, because PIC shared libraries use this register to store the address of the GOT.

The interesting bit is the `CALL` instruction and the following two lines. The `CALL` and `POP` combination obtains the address of the label `.get_GOT`, without having to know in advance where the program was loaded (since the `CALL` instruction is encoded relative to the current position). The `ADD` instruction makes use of one of the special PIC relocation types: GOTPC relocation. With the `WRT ..gotpc` qualifier specified, the symbol referenced (here `_GLOBAL_OFFSET_TABLE_`, the special symbol assigned to the GOT) is given as an offset from the beginning of the section. (Actually, `ELF` encodes it as the offset from the operand field of the `ADD` instruction, but NASM simplifies this deliberately, so you do things the same way for both `ELF` and `BSD`.) So the instruction then *adds* the beginning of the section, to get the real address of the GOT, and subtracts the value of `.get_GOT` which it knows is in `EBX`. Therefore, by the time that instruction has finished, `EBX` contains the address of the GOT.

If you didn't follow that, don't worry: it's never necessary to obtain the address of the GOT by any other means, so you can put those three instructions into a macro and safely ignore them:

```
%macro  get_GOT 0

        call    %%getgot
  %%getgot:
        pop     ebx
        add     ebx,_GLOBAL_OFFSET_TABLE_+$$-%%getgot wrt ..gotpc

%endmacro
```

8.2.2 Finding Your Local Data Items

Having got the GOT, you can then use it to obtain the addresses of your data items. Most variables will reside in the sections you have declared; they can be accessed using the `..gotoff` special `WRT` type. The way this works is like this:

```
        lea     eax,[ebx+myvar wrt ..gotoff]
```

The expression `myvar wrt ..gotoff` is calculated, when the shared library is linked, to be the offset to the local variable `myvar` from the beginning of the GOT. Therefore, adding it to `EBX` as above will place the real address of `myvar` in `EAX`.

If you declare variables as `GLOBAL` without specifying a size for them, they are shared between code modules in the library, but do not get exported from the library to the program that loaded it. They will still be in your ordinary data and BSS sections, so you can access them in the same way as local variables, using the above `..gotoff` mechanism.

Note that due to a peculiarity of the way BSD `a.out` format handles this relocation type, there must be at least one non–local symbol in the same section as the address you're trying to access.

8.2.3 Finding External and Common Data Items

If your library needs to get at an external variable (external to the *library*, not just to one of the modules within it), you must use the `..got` type to get at it. The `..got` type, instead of giving you the offset from the GOT base to the variable, gives you the offset from the GOT base to a GOT *entry* containing the address of the variable. The linker will set up this GOT entry when it builds the library, and the dynamic linker will place the correct address in it at load time. So to obtain the address of an external variable `extvar` in `EAX`, you would code

```
        mov     eax,[ebx+extvar wrt ..got]
```

This loads the address of `extvar` out of an entry in the GOT. The linker, when it builds the shared library, collects together every relocation of type `..got`, and builds the GOT so as to ensure it has every necessary entry present.

Common variables must also be accessed in this way.

8.2.4 Exporting Symbols to the Library User

If you want to export symbols to the user of the library, you have to declare whether they are functions or data, and if they are data, you have to give the size of the data item. This is because the dynamic linker has to build procedure linkage table entries for any exported functions, and also moves exported data items away from the library's data section in which they were declared.

So to export a function to users of the library, you must use

```
global  func:function           ; declare it as a function

func:   push    ebp

        ; etc.
```

And to export a data item such as an array, you would have to code

```
global  array:data array.end-array      ; give the size too

array:  resd    128
.end:
```

Be careful: If you export a variable to the library user, by declaring it as `GLOBAL` and supplying a size, the variable will end up living in the data section of the main program, rather than in your library's data section, where you declared it. So you will have to access your own global variable with the `..got` mechanism rather than `..gotoff`, as if it were external (which, effectively, it has become).

Equally, if you need to store the address of an exported global in one of your data sections, you can't do it by means of the standard sort of code:

```
dataptr:        dd      global_data_item        ; WRONG
```

NASM will interpret this code as an ordinary relocation, in which `global_data_item` is merely an offset from the beginning of the `.data` section (or whatever); so this reference will end up pointing at your data section instead of at the exported global which resides elsewhere.

Instead of the above code, then, you must write

```
dataptr:        dd      global_data_item wrt ..sym
```

which makes use of the special `WRT` type `..sym` to instruct NASM to search the symbol table for a particular symbol at that address, rather than just relocating by section base.

Either method will work for functions: referring to one of your functions by means of

```
funcptr:        dd      my_function
```

will give the user the address of the code you wrote, whereas

```
funcptr:        dd      my_function wrt .sym
```

will give the address of the procedure linkage table for the function, which is where the calling program will *believe* the function lives. Either address is a valid way to call the function.

8.2.5 Calling Procedures Outside the Library

Calling procedures outside your shared library has to be done by means of a *procedure linkage table*, or PLT. The PLT is placed at a known offset from where the library is loaded, so the library code can make calls to the PLT in a position–independent way. Within the PLT there is code to jump to offsets contained in the GOT, so function calls to other shared libraries or to routines in the main program can be transparently passed off to their real destinations.

To call an external routine, you must use another special PIC relocation type, `WRT ..plt`. This is much easier than the GOT–based ones: you simply replace calls such as `CALL printf` with the PLT–relative version `CALL printf WRT ..plt`.

8.2.6 Generating the Library File

Having written some code modules and assembled them to `.o` files, you then generate your shared library with a command such as

```
ld -shared -o library.so module1.o module2.o        # for ELF
ld -Bshareable -o library.so module1.o module2.o    # for BSD
```

For ELF, if your shared library is going to reside in system directories such as `/usr/lib` or `/lib`, it is usually worth using the `-soname` flag to the linker, to store the final library file name, with a version number, into the library:

```
ld -shared -soname library.so.1 -o library.so.1.2 *.o
```

You would then copy `library.so.1.2` into the library directory, and create `library.so.1` as a symbolic link to it.

Chapter 9: Mixing 16 and 32 Bit Code

This chapter tries to cover some of the issues, largely related to unusual forms of addressing and jump instructions, encountered when writing operating system code such as protected−mode initialisation routines, which require code that operates in mixed segment sizes, such as code in a 16−bit segment trying to modify data in a 32−bit one, or jumps between different−size segments.

9.1 Mixed−Size Jumps

The most common form of mixed−size instruction is the one used when writing a 32−bit OS: having done your setup in 16−bit mode, such as loading the kernel, you then have to boot it by switching into protected mode and jumping to the 32−bit kernel start address. In a fully 32−bit OS, this tends to be the *only* mixed−size instruction you need, since everything before it can be done in pure 16−bit code, and everything after it can be pure 32−bit.

This jump must specify a 48−bit far address, since the target segment is a 32−bit one. However, it must be assembled in a 16−bit segment, so just coding, for example,

```
jmp     0x1234:0x56789ABC       ; wrong!
```

will not work, since the offset part of the address will be truncated to `0x9ABC` and the jump will be an ordinary 16−bit far one.

The Linux kernel setup code gets round the inability of `as86` to generate the required instruction by coding it manually, using `DB` instructions. NASM can go one better than that, by actually generating the right instruction itself. Here's how to do it right:

```
jmp     dword 0x1234:0x56789ABC         ; right
```

The `DWORD` prefix (strictly speaking, it should come *after* the colon, since it is declaring the *offset* field to be a doubleword; but NASM will accept either form, since both are unambiguous) forces the offset part to be treated as far, in the assumption that you are deliberately writing a jump from a 16−bit segment to a 32−bit one.

You can do the reverse operation, jumping from a 32−bit segment to a 16−bit one, by means of the `WORD` prefix:

```
jmp     word 0x8765:0x4321      ; 32 to 16 bit
```

If the `WORD` prefix is specified in 16−bit mode, or the `DWORD` prefix in 32−bit mode, they will be ignored, since each is explicitly forcing NASM into a mode it was in anyway.

9.2 Addressing Between Different−Size Segments

If your OS is mixed 16 and 32−bit, or if you are writing a DOS extender, you are likely to have to deal with some 16−bit segments and some 32−bit ones. At some point, you will probably end up writing code in a 16−bit segment which has to access data in a 32−bit segment, or vice versa.

If the data you are trying to access in a 32−bit segment lies within the first 64K of the segment, you may be able to get away with using an ordinary 16−bit addressing operation for the purpose; but sooner or later, you will want to do 32−bit addressing from 16−bit mode.

The easiest way to do this is to make sure you use a register for the address, since any effective address containing a 32−bit register is forced to be a 32−bit address. So you can do

```
mov     eax,offset_into_32_bit_segment_specified_by_fs
mov     dword [fs:eax],0x11223344
```

This is fine, but slightly cumbersome (since it wastes an instruction and a register) if you already know the precise offset you are aiming at. The x86 architecture does allow 32–bit effective addresses to specify nothing but a 4–byte offset, so why shouldn't NASM be able to generate the best instruction for the purpose?

It can. As in section 9.1, you need only prefix the address with the `DWORD` keyword, and it will be forced to be a 32–bit address:

```
mov     dword [fs:dword my_offset],0x11223344
```

Also as in section 9.1, NASM is not fussy about whether the `DWORD` prefix comes before or after the segment override, so arguably a nicer–looking way to code the above instruction is

```
mov     dword [dword fs:my_offset],0x11223344
```

Don't confuse the `DWORD` prefix *outside* the square brackets, which controls the size of the data stored at the address, with the one `inside` the square brackets which controls the length of the address itself. The two can quite easily be different:

```
mov     word [dword 0x12345678],0x9ABC
```

This moves 16 bits of data to an address specified by a 32–bit offset.

You can also specify `WORD` or `DWORD` prefixes along with the `FAR` prefix to indirect far jumps or calls. For example:

```
call    dword far [fs:word 0x4321]
```

This instruction contains an address specified by a 16–bit offset; it loads a 48–bit far pointer from that (16–bit segment and 32–bit offset), and calls that address.

9.3 Other Mixed–Size Instructions

The other way you might want to access data might be using the string instructions (`LODSx`, `STOSx` and so on) or the `XLATB` instruction. These instructions, since they take no parameters, might seem to have no easy way to make them perform 32–bit addressing when assembled in a 16–bit segment.

This is the purpose of NASM's `a16` and `a32` prefixes. If you are coding `LODSB` in a 16–bit segment but it is supposed to be accessing a string in a 32–bit segment, you should load the desired address into `ESI` and then code

```
a32     lodsb
```

The prefix forces the addressing size to 32 bits, meaning that `LODSB` loads from `[DS:ESI]` instead of `[DS:SI]`. To access a string in a 16–bit segment when coding in a 32–bit one, the corresponding `a16` prefix can be used.

The `a16` and `a32` prefixes can be applied to any instruction in NASM's instruction table, but most of them can generate all the useful forms without them. The prefixes are necessary only for instructions with implicit addressing: `CMPSx` (section B.4.27), `SCASx` (section B.4.286), `LODSx` (section B.4.141), `STOSx` (section B.4.303), `MOVSx` (section B.4.178), `INSx` (section B.4.121), `OUTSx` (section B.4.195), and `XLATB` (section B.4.334). Also, the various push and pop instructions (`PUSHA` and `POPF` as well as the more usual `PUSH` and `POP`) can accept `a16` or `a32` prefixes to force a particular one of `SP` or `ESP` to be used as a stack pointer, in case the stack segment in use is a different size from the code segment.

`PUSH` and `POP`, when applied to segment registers in 32–bit mode, also have the slightly odd behaviour that they push and pop 4 bytes at a time, of which the top two are ignored and the bottom two give the value of the segment register being manipulated. To force the 16–bit behaviour of segment–register push and pop instructions, you can use the operand–size prefix `o16`:

```
o16 push    ss
o16 push    ds
```

This code saves a doubleword of stack space by fitting two segment registers into the space which would normally be consumed by pushing one.

(You can also use the `o32` prefix to force the 32–bit behaviour when in 16–bit mode, but this seems less useful.)

Chapter 10: Troubleshooting

This chapter describes some of the common problems that users have been known to encounter with NASM, and answers them. It also gives instructions for reporting bugs in NASM if you find a difficulty that isn't listed here.

10.1 Common Problems

10.1.1 NASM Generates Inefficient Code

We sometimes get 'bug' reports about NASM generating inefficient, or even 'wrong', code on instructions such as `ADD ESP,8`. This is a deliberate design feature, connected to predictability of output: NASM, on seeing `ADD ESP,8`, will generate the form of the instruction which leaves room for a 32–bit offset. You need to code `ADD ESP,BYTE 8` if you want the space–efficient form of the instruction. This isn't a bug, it's user error: if you prefer to have NASM produce the more efficient code automatically enable optimization with the `-On` option (see section 2.1.16).

10.1.2 My Jumps are Out of Range

Similarly, people complain that when they issue conditional jumps (which are `SHORT` by default) that try to jump too far, NASM reports 'short jump out of range' instead of making the jumps longer.

This, again, is partly a predictability issue, but in fact has a more practical reason as well. NASM has no means of being told what type of processor the code it is generating will be run on; so it cannot decide for itself that it should generate `Jcc NEAR` type instructions, because it doesn't know that it's working for a 386 or above. Alternatively, it could replace the out–of–range short `JNE` instruction with a very short `JE` instruction that jumps over a `JMP NEAR`; this is a sensible solution for processors below a 386, but hardly efficient on processors which have good branch prediction *and* could have used `JNE NEAR` instead. So, once again, it's up to the user, not the assembler, to decide what instructions should be generated. See section 2.1.16.

10.1.3 `ORG` Doesn't Work

People writing boot sector programs in the `bin` format often complain that `ORG` doesn't work the way they'd like: in order to place the `0xAA55` signature word at the end of a 512–byte boot sector, people who are used to MASM tend to code

```
ORG 0

; some boot sector code

ORG 510
DW 0xAA55
```

This is not the intended use of the `ORG` directive in NASM, and will not work. The correct way to solve this problem in NASM is to use the `TIMES` directive, like this:

```
ORG 0

; some boot sector code

TIMES 510-($-$$) DB 0
DW 0xAA55
```

The TIMES directive will insert exactly enough zero bytes into the output to move the assembly point up to 510. This method also has the advantage that if you accidentally fill your boot sector too full, NASM will catch the problem at assembly time and report it, so you won't end up with a boot sector that you have to disassemble to find out what's wrong with it.

10.1.4 TIMES Doesn't Work

The other common problem with the above code is people who write the TIMES line as

```
TIMES 510-$ DB 0
```

by reasoning that $ should be a pure number, just like 510, so the difference between them is also a pure number and can happily be fed to TIMES.

NASM is a *modular* assembler: the various component parts are designed to be easily separable for re-use, so they don't exchange information unnecessarily. In consequence, the bin output format, even though it has been told by the ORG directive that the .text section should start at 0, does not pass that information back to the expression evaluator. So from the evaluator's point of view, $ isn't a pure number: it's an offset from a section base. Therefore the difference between $ and 510 is also not a pure number, but involves a section base. Values involving section bases cannot be passed as arguments to TIMES.

The solution, as in the previous section, is to code the TIMES line in the form

```
TIMES 510-($-$$) DB 0
```

in which $ and $$ are offsets from the same section base, and so their difference is a pure number. This will solve the problem and generate sensible code.

10.2 Bugs

We have never yet released a version of NASM with any *known* bugs. That doesn't usually stop there being plenty we didn't know about, though. Any that you find should be reported firstly via the bugtracker at https://sourceforge.net/projects/nasm/ (click on "Bugs"), or if that fails then through one of the contacts in section 1.2.

Please read section 2.2 first, and don't report the bug if it's listed in there as a deliberate feature. (If you think the feature is badly thought out, feel free to send us reasons why you think it should be changed, but don't just send us mail saying 'This is a bug' if the documentation says we did it on purpose.) Then read section 10.1, and don't bother reporting the bug if it's listed there.

If you do report a bug, *please* give us all of the following information:

- What operating system you're running NASM under. DOS, Linux, NetBSD, Win16, Win32, VMS (I'd be impressed), whatever.
- If you're running NASM under DOS or Win32, tell us whether you've compiled your own executable from the DOS source archive, or whether you were using the standard distribution binaries out of the archive. If you were using a locally built executable, try to reproduce the problem using one of the standard binaries, as this will make it easier for us to reproduce your problem prior to fixing it.
- Which version of NASM you're using, and exactly how you invoked it. Give us the precise command line, and the contents of the NASMENV environment variable if any.
- Which versions of any supplementary programs you're using, and how you invoked them. If the problem only becomes visible at link time, tell us what linker you're using, what version of it you've got, and the exact linker command line. If the problem involves linking against object files generated by a compiler, tell us what compiler, what version, and what command line or

options you used. (If you're compiling in an IDE, please try to reproduce the problem with the command–line version of the compiler.)

- If at all possible, send us a NASM source file which exhibits the problem. If this causes copyright problems (e.g. you can only reproduce the bug in restricted–distribution code) then bear in mind the following two points: firstly, we guarantee that any source code sent to us for the purposes of debugging NASM will be used *only* for the purposes of debugging NASM, and that we will delete all our copies of it as soon as we have found and fixed the bug or bugs in question; and secondly, we would prefer *not* to be mailed large chunks of code anyway. The smaller the file, the better. A three–line sample file that does nothing useful *except* demonstrate the problem is much easier to work with than a fully fledged ten–thousand–line program. (Of course, some errors *do* only crop up in large files, so this may not be possible.)
- A description of what the problem actually *is*. 'It doesn't work' is *not* a helpful description! Please describe exactly what is happening that shouldn't be, or what isn't happening that should. Examples might be: 'NASM generates an error message saying Line 3 for an error that's actually on Line 5'; 'NASM generates an error message that I believe it shouldn't be generating at all'; 'NASM fails to generate an error message that I believe it *should* be generating'; 'the object file produced from this source code crashes my linker'; 'the ninth byte of the output file is 66 and I think it should be 77 instead'.
- If you believe the output file from NASM to be faulty, send it to us. That allows us to determine whether our own copy of NASM generates the same file, or whether the problem is related to portability issues between our development platforms and yours. We can handle binary files mailed to us as MIME attachments, uuencoded, and even BinHex. Alternatively, we may be able to provide an FTP site you can upload the suspect files to; but mailing them is easier for us.
- Any other information or data files that might be helpful. If, for example, the problem involves NASM failing to generate an object file while TASM can generate an equivalent file without trouble, then send us *both* object files, so we can see what TASM is doing differently from us.

Appendix A: Ndisasm

The Netwide Disassembler, NDISASM

A.1 Introduction

The Netwide Disassembler is a small companion program to the Netwide Assembler, NASM. It seemed a shame to have an x86 assembler, complete with a full instruction table, and not make as much use of it as possible, so here's a disassembler which shares the instruction table (and some other bits of code) with NASM.

The Netwide Disassembler does nothing except to produce disassemblies of *binary* source files. NDISASM does not have any understanding of object file formats, like `objdump`, and it will not understand `DOS .EXE` files like `debug` will. It just disassembles.

A.2 Getting Started: Installation

See section 1.3 for installation instructions. NDISASM, like NASM, has a `man page` which you may want to put somewhere useful, if you are on a Unix system.

A.3 Running NDISASM

To disassemble a file, you will typically use a command of the form

```
ndisasm [-b16 | -b32] filename
```

NDISASM can disassemble 16–bit code or 32–bit code equally easily, provided of course that you remember to specify which it is to work with. If no `-b` switch is present, NDISASM works in 16–bit mode by default. The `-u` switch (for USE32) also invokes 32–bit mode.

Two more command line options are `-r` which reports the version number of NDISASM you are running, and `-h` which gives a short summary of command line options.

A.3.1 COM Files: Specifying an Origin

To disassemble a `DOS .COM` file correctly, a disassembler must assume that the first instruction in the file is loaded at address `0x100`, rather than at zero. NDISASM, which assumes by default that any file you give it is loaded at zero, will therefore need to be informed of this.

The `-o` option allows you to declare a different origin for the file you are disassembling. Its argument may be expressed in any of the NASM numeric formats: decimal by default, if it begins with '`$`' or '`0x`' or ends in '`H`' it's `hex`, if it ends in '`Q`' it's `octal`, and if it ends in '`B`' it's `binary`.

Hence, to disassemble a `.COM` file:

```
ndisasm -o100h filename.com
```

will do the trick.

A.3.2 Code Following Data: Synchronisation

Suppose you are disassembling a file which contains some data which isn't machine code, and *then* contains some machine code. NDISASM will faithfully plough through the data section, producing machine instructions wherever it can (although most of them will look bizarre, and some may have

unusual prefixes, e.g. '`FS OR AX,0x240A`'), and generating 'DB' instructions ever so often if it's totally stumped. Then it will reach the code section.

Supposing NDISASM has just finished generating a strange machine instruction from part of the data section, and its file position is now one byte *before* the beginning of the code section. It's entirely possible that another spurious instruction will get generated, starting with the final byte of the data section, and then the correct first instruction in the code section will not be seen because the starting point skipped over it. This isn't really ideal.

To avoid this, you can specify a '`synchronisation`' point, or indeed as many synchronisation points as you like (although NDISASM can only handle 8192 sync points internally). The definition of a sync point is this: NDISASM guarantees to hit sync points exactly during disassembly. If it is thinking about generating an instruction which would cause it to jump over a sync point, it will discard that instruction and output a '`db`' instead. So it *will* start disassembly exactly from the sync point, and so you *will* see all the instructions in your code section.

Sync points are specified using the `-s` option: they are measured in terms of the program origin, not the file position. So if you want to synchronise after 32 bytes of a `.COM` file, you would have to do

```
ndisasm -o100h -s120h file.com
```

rather than

```
ndisasm -o100h -s20h file.com
```

As stated above, you can specify multiple sync markers if you need to, just by repeating the `-s` option.

A.3.3 Mixed Code and Data: Automatic (Intelligent) Synchronisation

Suppose you are disassembling the boot sector of a `DOS` floppy (maybe it has a virus, and you need to understand the virus so that you know what kinds of damage it might have done you). Typically, this will contain a `JMP` instruction, then some data, then the rest of the code. So there is a very good chance of NDISASM being *misaligned* when the data ends and the code begins. Hence a sync point is needed.

On the other hand, why should you have to specify the sync point manually? What you'd do in order to find where the sync point would be, surely, would be to read the `JMP` instruction, and then to use its target address as a sync point. So can NDISASM do that for you?

The answer, of course, is yes: using either of the synonymous switches `-a` (for automatic sync) or `-i` (for intelligent sync) will enable `auto-sync` mode. Auto-sync mode automatically generates a sync point for any forward-referring PC-relative jump or call instruction that NDISASM encounters. (Since NDISASM is one-pass, if it encounters a PC-relative jump whose target has already been processed, there isn't much it can do about it...)

Only PC-relative jumps are processed, since an absolute jump is either through a register (in which case NDISASM doesn't know what the register contains) or involves a segment address (in which case the target code isn't in the same segment that NDISASM is working in, and so the sync point can't be placed anywhere useful).

For some kinds of file, this mechanism will automatically put sync points in all the right places, and save you from having to place any sync points manually. However, it should be stressed that auto-sync mode is *not* guaranteed to catch all the sync points, and you may still have to place some manually.

Auto-sync mode doesn't prevent you from declaring manual sync points: it just adds automatically generated ones to the ones you provide. It's perfectly feasible to specify `-i` *and* some `-s` options.

Another caveat with auto−sync mode is that if, by some unpleasant fluke, something in your data section should disassemble to a PC−relative call or jump instruction, NDISASM may obediently place a sync point in a totally random place, for example in the middle of one of the instructions in your code section. So you may end up with a wrong disassembly even if you use auto−sync. Again, there isn't much I can do about this. If you have problems, you'll have to use manual sync points, or use the `-k` option (documented below) to suppress disassembly of the data area.

A.3.4 Other Options

The `-e` option skips a header on the file, by ignoring the first N bytes. This means that the header is *not* counted towards the disassembly offset: if you give `-e10 -o10`, disassembly will start at byte 10 in the file, and this will be given offset 10, not 20.

The `-k` option is provided with two comma−separated numeric arguments, the first of which is an assembly offset and the second is a number of bytes to skip. This *will* count the skipped bytes towards the assembly offset: its use is to suppress disassembly of a data section which wouldn't contain anything you wanted to see anyway.

A.4 Bugs and Improvements

There are no known bugs. However, any you find, with patches if possible, should be sent to `jules@dsf.org.uk` or `anakin@pobox.com`, or to the developer's site at `https://sourceforge.net/projects/nasm/` and we'll try to fix them. Feel free to send contributions and new features as well.

Future plans include awareness of which processors certain instructions will run on, and marking of instructions that are too advanced for some processor (or are `FPU` instructions, or are undocumented opcodes, or are privileged protected−mode instructions, or whatever).

That's All Folks!

I hope NDISASM is of some use to somebody. Including me. :−)

I don't recommend taking NDISASM apart to see how an efficient disassembler works, because as far as I know, it isn't an efficient one anyway. You have been warned.

Appendix B: x86 Instruction Reference

This appendix provides a complete list of the machine instructions which NASM will assemble, and a short description of the function of each one.

It is not intended to be exhaustive documentation on the fine details of the instructions' function, such as which exceptions they can trigger: for such documentation, you should go to Intel's Web site, `http://developer.intel.com/design/Pentium4/manuals/`.

Instead, this appendix is intended primarily to provide documentation on the way the instructions may be used within NASM. For example, looking up `LOOP` will tell you that NASM allows `CX` or `ECX` to be specified as an optional second argument to the `LOOP` instruction, to enforce which of the two possible counter registers should be used if the default is not the one desired.

The instructions are not quite listed in alphabetical order, since groups of instructions with similar functions are lumped together in the same entry. Most of them don't move very far from their alphabetic position because of this.

B.1 Key to Operand Specifications

The instruction descriptions in this appendix specify their operands using the following notation:

- Registers: `reg8` denotes an 8–bit general purpose register, `reg16` denotes a 16–bit general purpose register, and `reg32` a 32–bit one. `fpureg` denotes one of the eight FPU stack registers, `mmxreg` denotes one of the eight 64–bit MMX registers, and `segreg` denotes a segment register. In addition, some registers (such as `AL`, `DX` or `ECX`) may be specified explicitly.
- Immediate operands: `imm` denotes a generic immediate operand. `imm8`, `imm16` and `imm32` are used when the operand is intended to be a specific size. For some of these instructions, NASM needs an explicit specifier: for example, `ADD ESP,16` could be interpreted as either `ADD r/m32,imm32` or `ADD r/m32,imm8`. NASM chooses the former by default, and so you must specify `ADD ESP,BYTE 16` for the latter.
- Memory references: `mem` denotes a generic memory reference; `mem8`, `mem16`, `mem32`, `mem64` and `mem80` are used when the operand needs to be a specific size. Again, a specifier is needed in some cases: `DEC [address]` is ambiguous and will be rejected by NASM. You must specify `DEC BYTE [address]`, `DEC WORD [address]` or `DEC DWORD [address]` instead.
- Restricted memory references: one form of the `MOV` instruction allows a memory address to be specified *without* allowing the normal range of register combinations and effective address processing. This is denoted by `memoffs8`, `memoffs16` and `memoffs32`.
- Register or memory choices: many instructions can accept either a register *or* a memory reference as an operand. `r/m8` is a shorthand for `reg8/mem8`; similarly `r/m16` and `r/m32`. `r/m64` is MMX–related, and is a shorthand for `mmxreg/mem64`.

B.2 Key to Opcode Descriptions

This appendix also provides the opcodes which NASM will generate for each form of each instruction. The opcodes are listed in the following way:

- A hex number, such as `3F`, indicates a fixed byte containing that number.

- A hex number followed by `+r`, such as `C8+r`, indicates that one of the operands to the instruction is a register, and the 'register value' of that register should be added to the hex number to produce the generated byte. For example, EDX has register value 2, so the code `C8+r`, when the register operand is EDX, generates the hex byte `CA`. Register values for specific registers are given in section B.2.1.
- A hex number followed by `+cc`, such as `40+cc`, indicates that the instruction name has a condition code suffix, and the numeric representation of the condition code should be added to the hex number to produce the generated byte. For example, the code `40+cc`, when the instruction contains the `NE` condition, generates the hex byte `45`. Condition codes and their numeric representations are given in section B.2.2.
- A slash followed by a digit, such as `/2`, indicates that one of the operands to the instruction is a memory address or register (denoted `mem` or `r/m`, with an optional size). This is to be encoded as an effective address, with a ModR/M byte, an optional SIB byte, and an optional displacement, and the spare (register) field of the ModR/M byte should be the digit given (which will be from 0 to 7, so it fits in three bits). The encoding of effective addresses is given in section B.2.5.
- The code `/r` combines the above two: it indicates that one of the operands is a memory address or `r/m`, and another is a register, and that an effective address should be generated with the spare (register) field in the ModR/M byte being equal to the 'register value' of the register operand. The encoding of effective addresses is given in section B.2.5; register values are given in section B.2.1.
- The codes `ib`, `iw` and `id` indicate that one of the operands to the instruction is an immediate value, and that this is to be encoded as a byte, little−endian word or little−endian doubleword respectively.
- The codes `rb`, `rw` and `rd` indicate that one of the operands to the instruction is an immediate value, and that the *difference* between this value and the address of the end of the instruction is to be encoded as a byte, word or doubleword respectively. Where the form `rw/rd` appears, it indicates that either `rw` or `rd` should be used according to whether assembly is being performed in `BITS 16` or `BITS 32` state respectively.
- The codes `ow` and `od` indicate that one of the operands to the instruction is a reference to the contents of a memory address specified as an immediate value: this encoding is used in some forms of the `MOV` instruction in place of the standard effective−address mechanism. The displacement is encoded as a word or doubleword. Again, `ow/od` denotes that `ow` or `od` should be chosen according to the `BITS` setting.
- The codes `o16` and `o32` indicate that the given form of the instruction should be assembled with operand size 16 or 32 bits. In other words, `o16` indicates a `66` prefix in `BITS 32` state, but generates no code in `BITS 16` state; and `o32` indicates a `66` prefix in `BITS 16` state but generates nothing in `BITS 32`.
- The codes `a16` and `a32`, similarly to `o16` and `o32`, indicate the address size of the given form of the instruction. Where this does not match the `BITS` setting, a `67` prefix is required.

B.2.1 Register Values

Where an instruction requires a register value, it is already implicit in the encoding of the rest of the instruction what type of register is intended: an 8−bit general−purpose register, a segment register, a debug register, an MMX register, or whatever. Therefore there is no problem with registers of different types sharing an encoding value.

The encodings for the various classes of register are:

- 8−bit general registers: `AL` is 0, `CL` is 1, `DL` is 2, `BL` is 3, `AH` is 4, `CH` is 5, `DH` is 6, and `BH` is 7.

- 16–bit general registers: AX is 0, CX is 1, DX is 2, BX is 3, SP is 4, BP is 5, SI is 6, and DI is 7.
- 32–bit general registers: EAX is 0, ECX is 1, EDX is 2, EBX is 3, ESP is 4, EBP is 5, ESI is 6, and EDI is 7.
- Segment registers: ES is 0, CS is 1, SS is 2, DS is 3, FS is 4, and GS is 5.
- Floating–point registers: ST0 is 0, ST1 is 1, ST2 is 2, ST3 is 3, ST4 is 4, ST5 is 5, ST6 is 6, and ST7 is 7.
- 64–bit MMX registers: MM0 is 0, MM1 is 1, MM2 is 2, MM3 is 3, MM4 is 4, MM5 is 5, MM6 is 6, and MM7 is 7.
- Control registers: CR0 is 0, CR2 is 2, CR3 is 3, and CR4 is 4.
- Debug registers: DR0 is 0, DR1 is 1, DR2 is 2, DR3 is 3, DR6 is 6, and DR7 is 7.
- Test registers: TR3 is 3, TR4 is 4, TR5 is 5, TR6 is 6, and TR7 is 7.

(Note that wherever a register name contains a number, that number is also the register value for that register.)

B.2.2 Condition Codes

The available condition codes are given here, along with their numeric representations as part of opcodes. Many of these condition codes have synonyms, so several will be listed at a time.

In the following descriptions, the word ‘either’, when applied to two possible trigger conditions, is used to mean ‘either or both’. If ‘either but not both’ is meant, the phrase ‘exactly one of’ is used.

- O is 0 (trigger if the overflow flag is set); NO is 1.
- B, C and NAE are 2 (trigger if the carry flag is set); AE, NB and NC are 3.
- E and Z are 4 (trigger if the zero flag is set); NE and NZ are 5.
- BE and NA are 6 (trigger if either of the carry or zero flags is set); A and NBE are 7.
- S is 8 (trigger if the sign flag is set); NS is 9.
- P and PE are 10 (trigger if the parity flag is set); NP and PO are 11.
- L and NGE are 12 (trigger if exactly one of the sign and overflow flags is set); GE and NL are 13.
- LE and NG are 14 (trigger if either the zero flag is set, or exactly one of the sign and overflow flags is set); G and NLE are 15.

Note that in all cases, the sense of a condition code may be reversed by changing the low bit of the numeric representation.

For details of when an instruction sets each of the status flags, see the individual instruction, plus the Status Flags reference in section B.2.4

B.2.3 SSE Condition Predicates

The condition predicates for SSE comparison instructions are the codes used as part of the opcode, to determine what form of comparison is being carried out. In each case, the imm8 value is the final byte of the opcode encoding, and the predicate is the code used as part of the mnemonic for the instruction (equivalent to the "cc" in an integer instruction that used a condition code). The instructions that use this will give details of what the various mnemonics are, this table is used to help you work out details of what is happening.

```
Predi-  imm8  Description Relation where:   Emula- Result   QNaN
 cate  Encod-             A Is 1st Operand  tion   if NaN   Signal
```

	ing		B Is 2nd Operand		Operand	Invalid
EQ	000B	equal	A = B		False	No
LT	001B	less-than	A < B		False	Yes
LE	010B	less-than-or-equal	A <= B		False	Yes
---	----	greater than	A > B	Swap Operands, Use LT	False	Yes
---	----	greater-than-or-equal	A >= B	Swap Operands, Use LE	False	Yes
UNORD	011B	unordered	A, B = Unordered		True	No
NEQ	100B	not-equal	A != B		True	No
NLT	101B	not-less-than	NOT(A < B)		True	Yes
NLE	110B	not-less-than-or-equal	NOT(A <= B)		True	Yes
---	----	not-greater than	NOT(A > B)	Swap Operands, Use NLT	True	Yes
---	----	not-greater than-or-equal	NOT(A >= B)	Swap Operands, Use NLE	True	Yes
ORD	111B	ordered	A , B = Ordered		False	No

The unordered relationship is true when at least one of the two values being compared is a NaN or in an unsupported format.

Note that the comparisons which are listed as not having a predicate or encoding can only be achieved through software emulation, as described in the "emulation" column. Note in particular that an instruction such as `greater-than` is not the same as `NLE`, as, unlike with the `CMP` instruction, it has to take into account the possibility of one operand containing a NaN or an unsupported numeric format.

B.2.4 Status Flags

The status flags provide some information about the result of the arithmetic instructions. This information can be used by conditional instructions (such a `Jcc` and `CMOVcc`) as well as by some of the other instructions (such as `ADC` and `INTO`).

There are 6 status flags:

`CF - Carry flag.`

Set if an arithmetic operation generates a carry or a borrow out of the most−significant bit of the result; cleared otherwise. This flag indicates an overflow condition for unsigned−integer arithmetic. It is also used in multiple−precision arithmetic.

`PF - Parity flag.`

Set if the least−significant byte of the result contains an even number of 1 bits; cleared otherwise.

`AF - Adjust flag.`

Set if an arithmetic operation generates a carry or a borrow out of bit 3 of the result; cleared otherwise. This flag is used in binary−coded decimal (BCD) arithmetic.

`ZF - Zero flag.`

Set if the result is zero; cleared otherwise.

`SF - Sign flag.`

Set equal to the most−significant bit of the result, which is the sign bit of a signed integer. (0 indicates a positive value and 1 indicates a negative value.)

`OF - Overflow flag.`

Set if the integer result is too large a positive number or too small a negative number (excluding the sign−bit) to fit in the destination operand; cleared otherwise. This flag indicates an overflow condition for signed−integer (two's complement) arithmetic.

B.2.5 Effective Address Encoding: ModR/M and SIB

An effective address is encoded in up to three parts: a ModR/M byte, an optional SIB byte, and an optional byte, word or doubleword displacement field.

The ModR/M byte consists of three fields: the `mod` field, ranging from 0 to 3, in the upper two bits of the byte, the `r/m` field, ranging from 0 to 7, in the lower three bits, and the spare (register) field in the middle (bit 3 to bit 5). The spare field is not relevant to the effective address being encoded, and either contains an extension to the instruction opcode or the register value of another operand.

The ModR/M system can be used to encode a direct register reference rather than a memory access. This is always done by setting the `mod` field to 3 and the `r/m` field to the register value of the register in question (it must be a general−purpose register, and the size of the register must already be implicit in the encoding of the rest of the instruction). In this case, the SIB byte and displacement field are both absent.

In 16−bit addressing mode (either `BITS 16` with no `67` prefix, or `BITS 32` with a `67` prefix), the SIB byte is never used. The general rules for `mod` and `r/m` (there is an exception, given below) are:

- The `mod` field gives the length of the displacement field: 0 means no displacement, 1 means one byte, and 2 means two bytes.
- The `r/m` field encodes the combination of registers to be added to the displacement to give the accessed address: 0 means `BX+SI`, 1 means `BX+DI`, 2 means `BP+SI`, 3 means `BP+DI`, 4 means `SI` only, 5 means `DI` only, 6 means `BP` only, and 7 means `BX` only.

However, there is a special case:

- If `mod` is 0 and `r/m` is 6, the effective address encoded is not `[BP]` as the above rules would suggest, but instead `[disp16]`: the displacement field is present and is two bytes long, and no registers are added to the displacement.

Therefore the effective address `[BP]` cannot be encoded as efficiently as `[BX]`; so if you code `[BP]` in a program, NASM adds a notional 8−bit zero displacement, and sets `mod` to 1, `r/m` to 6, and the one−byte displacement field to 0.

In 32−bit addressing mode (either `BITS 16` with a `67` prefix, or `BITS 32` with no `67` prefix) the general rules (again, there are exceptions) for `mod` and `r/m` are:

- The `mod` field gives the length of the displacement field: 0 means no displacement, 1 means one byte, and 2 means four bytes.
- If only one register is to be added to the displacement, and it is not `ESP`, the `r/m` field gives its register value, and the SIB byte is absent. If the `r/m` field is 4 (which would encode `ESP`), the SIB byte is present and gives the combination and scaling of registers to be added to the displacement.

If the SIB byte is present, it describes the combination of registers (an optional base register, and an optional index register scaled by multiplication by 1, 2, 4 or 8) to be added to the displacement. The SIB byte is divided into the `scale` field, in the top two bits, the `index` field in the next three, and the `base` field in the bottom three. The general rules are:

- The `base` field encodes the register value of the base register.
- The `index` field encodes the register value of the index register, unless it is 4, in which case no index register is used (so `ESP` cannot be used as an index register).
- The `scale` field encodes the multiplier by which the index register is scaled before adding it to the base and displacement: 0 encodes a multiplier of 1, 1 encodes 2, 2 encodes 4 and 3 encodes 8.

The exceptions to the 32−bit encoding rules are:

- If `mod` is 0 and `r/m` is 5, the effective address encoded is not `[EBP]` as the above rules would suggest, but instead `[disp32]`: the displacement field is present and is four bytes long, and no registers are added to the displacement.
- If `mod` is 0, `r/m` is 4 (meaning the SIB byte is present) and `base` is 4, the effective address encoded is not `[EBP+index]` as the above rules would suggest, but instead `[disp32+index]`: the displacement field is present and is four bytes long, and there is no base register (but the index register is still processed in the normal way).

B.3 Key to Instruction Flags

Given along with each instruction in this appendix is a set of flags, denoting the type of the instruction. The types are as follows:

- `8086`, `186`, `286`, `386`, `486`, `PENT` and `P6` denote the lowest processor type that supports the instruction. Most instructions run on all processors above the given type; those that do not are documented. The Pentium II contains no additional instructions beyond the P6 (Pentium Pro); from the point of view of its instruction set, it can be thought of as a P6 with MMX capability.
- `3DNOW` indicates that the instruction is a 3DNow! one, and will run on the AMD K6−2 and later processors. ATHLON extensions to the 3DNow! instruction set are documented as such.
- `CYRIX` indicates that the instruction is specific to Cyrix processors, for example the extra MMX instructions in the Cyrix extended MMX instruction set.
- `FPU` indicates that the instruction is a floating−point one, and will only run on machines with a coprocessor (automatically including 486DX, Pentium and above).
- `KATMAI` indicates that the instruction was introduced as part of the Katmai New Instruction set. These instructions are available on the Pentium III and later processors. Those which are not specifically SSE instructions are also available on the AMD Athlon.

- `MMX` indicates that the instruction is an MMX one, and will run on MMX−capable Pentium processors and the Pentium II.
- `PRIV` indicates that the instruction is a protected−mode management instruction. Many of these may only be used in protected mode, or only at privilege level zero.
- `SSE` and `SSE2` indicate that the instruction is a Streaming SIMD Extension instruction. These instructions operate on multiple values in a single operation. SSE was introduced with the Pentium III and SSE2 was introduced with the Pentium 4.
- `UNDOC` indicates that the instruction is an undocumented one, and not part of the official Intel Architecture; it may or may not be supported on any given machine.
- `WILLAMETTE` indicates that the instruction was introduced as part of the new instruction set in the Pentium 4 and Intel Xeon processors. These instructions are also known as SSE2 instructions.

B.4 x86 Instruction Set

B.4.1 AAA, AAS, AAM, AAD: ASCII Adjustments

```
AAA                             ; 37                    [8086]

AAS                             ; 3F                    [8086]

AAD                             ; D5 0A                 [8086]
AAD imm                         ; D5 ib                 [8086]

AAM                             ; D4 0A                 [8086]
AAM imm                         ; D4 ib                 [8086]
```

These instructions are used in conjunction with the add, subtract, multiply and divide instructions to perform binary−coded decimal arithmetic in *unpacked* (one BCD digit per byte – easy to translate to and from `ASCII`, hence the instruction names) form. There are also packed BCD instructions `DAA` and `DAS`: see section B.4.57.

- `AAA` (ASCII Adjust After Addition) should be used after a one−byte `ADD` instruction whose destination was the `AL` register: by means of examining the value in the low nibble of `AL` and also the auxiliary carry flag `AF`, it determines whether the addition has overflowed, and adjusts it (and sets the carry flag) if so. You can add long BCD strings together by doing `ADD`/`AAA` on the low digits, then doing `ADC`/`AAA` on each subsequent digit.
- `AAS` (ASCII Adjust AL After Subtraction) works similarly to `AAA`, but is for use after `SUB` instructions rather than `ADD`.
- `AAM` (ASCII Adjust AX After Multiply) is for use after you have multiplied two decimal digits together and left the result in `AL`: it divides `AL` by ten and stores the quotient in `AH`, leaving the remainder in `AL`. The divisor 10 can be changed by specifying an operand to the instruction: a particularly handy use of this is `AAM 16`, causing the two nibbles in `AL` to be separated into `AH` and `AL`.
- `AAD` (ASCII Adjust AX Before Division) performs the inverse operation to `AAM`: it multiplies `AH` by ten, adds it to `AL`, and sets `AH` to zero. Again, the multiplier 10 can be changed.

B.4.2 ADC: Add with Carry

```
ADC r/m8,reg8                   ; 10 /r                 [8086]
ADC r/m16,reg16                 ; o16 11 /r             [8086]
ADC r/m32,reg32                 ; o32 11 /r             [386]
```

```
ADC reg8,r/m8                   ; 12 /r                 [8086]
ADC reg16,r/m16                 ; o16 13 /r             [8086]
ADC reg32,r/m32                 ; o32 13 /r             [386]

ADC r/m8,imm8                   ; 80 /2 ib              [8086]
ADC r/m16,imm16                 ; o16 81 /2 iw          [8086]
ADC r/m32,imm32                 ; o32 81 /2 id          [386]

ADC r/m16,imm8                  ; o16 83 /2 ib          [8086]
ADC r/m32,imm8                  ; o32 83 /2 ib          [386]

ADC AL,imm8                     ; 14 ib                 [8086]
ADC AX,imm16                    ; o16 15 iw             [8086]
ADC EAX,imm32                   ; o32 15 id             [386]
```

ADC performs integer addition: it adds its two operands together, plus the value of the carry flag, and leaves the result in its destination (first) operand. The destination operand can be a register or a memory location. The source operand can be a register, a memory location or an immediate value.

The flags are set according to the result of the operation: in particular, the carry flag is affected and can be used by a subsequent ADC instruction.

In the forms with an 8−bit immediate second operand and a longer first operand, the second operand is considered to be signed, and is sign−extended to the length of the first operand. In these cases, the BYTE qualifier is necessary to force NASM to generate this form of the instruction.

To add two numbers without also adding the contents of the carry flag, use ADD (section B.4.3).

B.4.3 ADD: Add Integers

```
ADD r/m8,reg8                   ; 00 /r                 [8086]
ADD r/m16,reg16                 ; o16 01 /r             [8086]
ADD r/m32,reg32                 ; o32 01 /r             [386]

ADD reg8,r/m8                   ; 02 /r                 [8086]
ADD reg16,r/m16                 ; o16 03 /r             [8086]
ADD reg32,r/m32                 ; o32 03 /r             [386]

ADD r/m8,imm8                   ; 80 /7 ib              [8086]
ADD r/m16,imm16                 ; o16 81 /7 iw          [8086]
ADD r/m32,imm32                 ; o32 81 /7 id          [386]

ADD r/m16,imm8                  ; o16 83 /7 ib          [8086]
ADD r/m32,imm8                  ; o32 83 /7 ib          [386]

ADD AL,imm8                     ; 04 ib                 [8086]
ADD AX,imm16                    ; o16 05 iw             [8086]
ADD EAX,imm32                   ; o32 05 id             [386]
```

ADD performs integer addition: it adds its two operands together, and leaves the result in its destination (first) operand. The destination operand can be a register or a memory location. The source operand can be a register, a memory location or an immediate value.

The flags are set according to the result of the operation: in particular, the carry flag is affected and can be used by a subsequent ADC instruction.

In the forms with an 8−bit immediate second operand and a longer first operand, the second operand is considered to be signed, and is sign−extended to the length of the first operand. In these cases, the BYTE qualifier is necessary to force NASM to generate this form of the instruction.

B.4.4 ADDPD: ADD Packed Double−Precision FP Values

```
ADDPD xmm1,xmm2/mem128          ; 66 0F 58 /r      [WILLAMETTE,SSE2]
```

ADDPD performs addition on each of two packed double−precision FP value pairs.

```
dst[0−63]   := dst[0−63]   + src[0−63],
dst[64−127] := dst[64−127] + src[64−127].
```

The destination is an XMM register. The source operand can be either an XMM register or a 128−bit memory location.

B.4.5 ADDPS: ADD Packed Single−Precision FP Values

```
ADDPS xmm1,xmm2/mem128          ; 0F 58 /r         [KATMAI,SSE]
```

ADDPS performs addition on each of four packed single−precision FP value pairs

```
dst[0−31]   := dst[0−31]   + src[0−31],
dst[32−63]  := dst[32−63]  + src[32−63],
dst[64−95]  := dst[64−95]  + src[64−95],
dst[96−127] := dst[96−127] + src[96−127].
```

The destination is an XMM register. The source operand can be either an XMM register or a 128−bit memory location.

B.4.6 ADDSD: ADD Scalar Double−Precision FP Values

```
ADDSD xmm1,xmm2/mem64           ; F2 0F 58 /r      [KATMAI,SSE]
```

ADDSD adds the low double−precision FP values from the source and destination operands and stores the double−precision FP result in the destination operand.

```
dst[0−63]   := dst[0−63] + src[0−63],
dst[64−127) remains unchanged.
```

The destination is an XMM register. The source operand can be either an XMM register or a 64−bit memory location.

B.4.7 ADDSS: ADD Scalar Single−Precision FP Values

```
ADDSS xmm1,xmm2/mem32           ; F3 0F 58 /r      [WILLAMETTE,SSE2]
```

ADDSS adds the low single−precision FP values from the source and destination operands and stores the single−precision FP result in the destination operand.

```
dst[0−31]   := dst[0−31] + src[0−31],
dst[32−127] remains unchanged.
```

The destination is an XMM register. The source operand can be either an XMM register or a 32−bit memory location.

B.4.8 AND: Bitwise AND

```
AND r/m8,reg8                   ; 20 /r                [8086]
AND r/m16,reg16                 ; o16 21 /r            [8086]
AND r/m32,reg32                 ; o32 21 /r            [386]

AND reg8,r/m8                   ; 22 /r                [8086]
AND reg16,r/m16                 ; o16 23 /r            [8086]
AND reg32,r/m32                 ; o32 23 /r            [386]
```

```
AND r/m8,imm8                   ; 80 /4 ib              [8086]
AND r/m16,imm16                 ; o16 81 /4 iw          [8086]
AND r/m32,imm32                 ; o32 81 /4 id          [386]

AND r/m16,imm8                  ; o16 83 /4 ib          [8086]
AND r/m32,imm8                  ; o32 83 /4 ib          [386]

AND AL,imm8                     ; 24 ib                 [8086]
AND AX,imm16                    ; o16 25 iw             [8086]
AND EAX,imm32                   ; o32 25 id             [386]
```

AND performs a bitwise AND operation between its two operands (i.e. each bit of the result is 1 if and only if the corresponding bits of the two inputs were both 1), and stores the result in the destination (first) operand. The destination operand can be a register or a memory location. The source operand can be a register, a memory location or an immediate value.

In the forms with an 8–bit immediate second operand and a longer first operand, the second operand is considered to be signed, and is sign–extended to the length of the first operand. In these cases, the BYTE qualifier is necessary to force NASM to generate this form of the instruction.

The MMX instruction PAND (see section B.4.202) performs the same operation on the 64–bit MMX registers.

B.4.9 ANDNPD: Bitwise Logical AND NOT of Packed Double–Precision FP Values

```
ANDNPD xmm1,xmm2/mem128         ; 66 0F 55 /r       [WILLAMETTE,SSE2]
```

ANDNPD inverts the bits of the two double–precision floating–point values in the destination register, and then performs a logical AND between the two double–precision floating–point values in the source operand and the temporary inverted result, storing the result in the destination register.

```
dst[0-63]   := src[0-63]   AND NOT dst[0-63],
dst[64-127] := src[64-127] AND NOT dst[64-127].
```

The destination is an XMM register. The source operand can be either an XMM register or a 128–bit memory location.

B.4.10 ANDNPS: Bitwise Logical AND NOT of Packed Single–Precision FP Values

```
ANDNPS xmm1,xmm2/mem128         ; 0F 55 /r          [KATMAI,SSE]
```

ANDNPS inverts the bits of the four single–precision floating–point values in the destination register, and then performs a logical AND between the four single–precision floating–point values in the source operand and the temporary inverted result, storing the result in the destination register.

```
dst[0-31]   := src[0-31]   AND NOT dst[0-31],
dst[32-63]  := src[32-63]  AND NOT dst[32-63],
dst[64-95]  := src[64-95]  AND NOT dst[64-95],
dst[96-127] := src[96-127] AND NOT dst[96-127].
```

The destination is an XMM register. The source operand can be either an XMM register or a 128–bit memory location.

B.4.11 ANDPD: Bitwise Logical AND For Single FP

```
ANDPD xmm1,xmm2/mem128          ; 66 0F 54 /r       [WILLAMETTE,SSE2]
```

ANDPD performs a bitwise logical AND of the two double–precision floating point values in the source and destination operand, and stores the result in the destination register.

```
dst[0-63]   := src[0-63]   AND dst[0-63],
dst[64-127] := src[64-127] AND dst[64-127].
```

The destination is an XMM register. The source operand can be either an XMM register or a 128–bit memory location.

B.4.12 ANDPS: Bitwise Logical AND For Single FP

```
ANDPS xmm1,xmm2/mem128          ; 0F 54 /r          [KATMAI,SSE]
```

ANDPS performs a bitwise logical AND of the four single–precision floating point values in the source and destination operand, and stores the result in the destination register.

```
dst[0-31]   := src[0-31]   AND dst[0-31],
dst[32-63]  := src[32-63]  AND dst[32-63],
dst[64-95]  := src[64-95]  AND dst[64-95],
dst[96-127] := src[96-127] AND dst[96-127].
```

The destination is an XMM register. The source operand can be either an XMM register or a 128–bit memory location.

B.4.13 ARPL: Adjust RPL Field of Selector

```
ARPL r/m16,reg16                ; 63 /r                  [286,PRIV]
```

ARPL expects its two word operands to be segment selectors. It adjusts the RPL (requested privilege level – stored in the bottom two bits of the selector) field of the destination (first) operand to ensure that it is no less (i.e. no more privileged than) the RPL field of the source operand. The zero flag is set if and only if a change had to be made.

B.4.14 BOUND: Check Array Index against Bounds

```
BOUND reg16,mem                 ; o16 62 /r              [186]
BOUND reg32,mem                 ; o32 62 /r              [386]
```

BOUND expects its second operand to point to an area of memory containing two signed values of the same size as its first operand (i.e. two words for the 16–bit form; two doublewords for the 32–bit form). It performs two signed comparisons: if the value in the register passed as its first operand is less than the first of the in–memory values, or is greater than or equal to the second, it throws a BR exception. Otherwise, it does nothing.

B.4.15 BSF, BSR: Bit Scan

```
BSF reg16,r/m16                 ; o16 0F BC /r           [386]
BSF reg32,r/m32                 ; o32 0F BC /r           [386]

BSR reg16,r/m16                 ; o16 0F BD /r           [386]
BSR reg32,r/m32                 ; o32 0F BD /r           [386]
```

- BSF searches for the least significant set bit in its source (second) operand, and if it finds one, stores the index in its destination (first) operand. If no set bit is found, the contents of the destination operand are undefined. If the source operand is zero, the zero flag is set.
- BSR performs the same function, but searches from the top instead, so it finds the most significant set bit.

Bit indices are from 0 (least significant) to 15 or 31 (most significant). The destination operand can only be a register. The source operand can be a register or a memory location.

B.4.16 BSWAP: Byte Swap

```
BSWAP reg32                     ; o32 0F C8+r            [486]
```

BSWAP swaps the order of the four bytes of a 32–bit register: bits 0–7 exchange places with bits 24–31, and bits 8–15 swap with bits 16–23. There is no explicit 16–bit equivalent: to byte–swap

AX, BX, CX or DX, XCHG can be used. When BSWAP is used with a 16–bit register, the result is undefined.

B.4.17 BT, BTC, BTR, BTS: Bit Test

```
BT r/m16,reg16                  ; o16 0F A3 /r          [386]
BT r/m32,reg32                  ; o32 0F A3 /r          [386]
BT r/m16,imm8                   ; o16 0F BA /4 ib       [386]
BT r/m32,imm8                   ; o32 0F BA /4 ib       [386]

BTC r/m16,reg16                 ; o16 0F BB /r          [386]
BTC r/m32,reg32                 ; o32 0F BB /r          [386]
BTC r/m16,imm8                  ; o16 0F BA /7 ib       [386]
BTC r/m32,imm8                  ; o32 0F BA /7 ib       [386]

BTR r/m16,reg16                 ; o16 0F B3 /r          [386]
BTR r/m32,reg32                 ; o32 0F B3 /r          [386]
BTR r/m16,imm8                  ; o16 0F BA /6 ib       [386]
BTR r/m32,imm8                  ; o32 0F BA /6 ib       [386]

BTS r/m16,reg16                 ; o16 0F AB /r          [386]
BTS r/m32,reg32                 ; o32 0F AB /r          [386]
BTS r/m16,imm                   ; o16 0F BA /5 ib       [386]
BTS r/m32,imm                   ; o32 0F BA /5 ib       [386]
```

These instructions all test one bit of their first operand, whose index is given by the second operand, and store the value of that bit into the carry flag. Bit indices are from 0 (least significant) to 15 or 31 (most significant).

In addition to storing the original value of the bit into the carry flag, BTR also resets (clears) the bit in the operand itself. BTS sets the bit, and BTC complements the bit. BT does not modify its operands.

The destination can be a register or a memory location. The source can be a register or an immediate value.

If the destination operand is a register, the bit offset should be in the range 0–15 (for 16–bit operands) or 0–31 (for 32–bit operands). An immediate value outside these ranges will be taken modulo 16/32 by the processor.

If the destination operand is a memory location, then an immediate bit offset follows the same rules as for a register. If the bit offset is in a register, then it can be anything within the signed range of the register used (ie, for a 32–bit operand, it can be (–2^31) to (2^31 – 1)

B.4.18 CALL: Call Subroutine

```
CALL imm                        ; E8 rw/rd              [8086]
CALL imm:imm16                  ; o16 9A iw iw          [8086]
CALL imm:imm32                  ; o32 9A id iw          [386]
CALL FAR mem16                  ; o16 FF /3             [8086]
CALL FAR mem32                  ; o32 FF /3             [386]
CALL r/m16                      ; o16 FF /2             [8086]
CALL r/m32                      ; o32 FF /2             [386]
```

CALL calls a subroutine, by means of pushing the current instruction pointer (IP) and optionally CS as well on the stack, and then jumping to a given address.

CS is pushed as well as IP if and only if the call is a far call, i.e. a destination segment address is specified in the instruction. The forms involving two colon–separated arguments are far calls; so are the CALL FAR mem forms.

The immediate near call takes one of two forms (`call imm16/imm32`, determined by the current segment size limit. For 16–bit operands, you would use `CALL 0x1234`, and for 32–bit operands you would use `CALL 0x12345678`. The value passed as an operand is a relative offset.

You can choose between the two immediate far call forms (`CALL imm:imm`) by the use of the `WORD` and `DWORD` keywords: `CALL WORD 0x1234:0x5678`) or `CALL DWORD 0x1234:0x56789abc`.

The `CALL FAR mem` forms execute a far call by loading the destination address out of memory. The address loaded consists of 16 or 32 bits of offset (depending on the operand size), and 16 bits of segment. The operand size may be overridden using `CALL WORD FAR mem` or `CALL DWORD FAR mem`.

The `CALL r/m` forms execute a near call (within the same segment), loading the destination address out of memory or out of a register. The keyword `NEAR` may be specified, for clarity, in these forms, but is not necessary. Again, operand size can be overridden using `CALL WORD mem` or `CALL DWORD mem`.

As a convenience, NASM does not require you to call a far procedure symbol by coding the cumbersome `CALL SEG routine:routine`, but instead allows the easier synonym `CALL FAR routine`.

The `CALL r/m` forms given above are near calls; NASM will accept the `NEAR` keyword (e.g. `CALL NEAR [address]`), even though it is not strictly necessary.

B.4.19 CBW, CWD, CDQ, CWDE: Sign Extensions

```
CBW                             ; o16 98            [8086]
CWDE                            ; o32 98            [386]

CWD                             ; o16 99            [8086]
CDQ                             ; o32 99            [386]
```

All these instructions sign–extend a short value into a longer one, by replicating the top bit of the original value to fill the extended one.

`CBW` extends `AL` into `AX` by repeating the top bit of `AL` in every bit of `AH`. `CWDE` extends `AX` into `EAX`. `CWD` extends `AX` into `DX:AX` by repeating the top bit of `AX` throughout `DX`, and `CDQ` extends `EAX` into `EDX:EAX`.

B.4.20 CLC, CLD, CLI, CLTS: Clear Flags

```
CLC                             ; F8                [8086]
CLD                             ; FC                [8086]
CLI                             ; FA                [8086]
CLTS                            ; 0F 06             [286,PRIV]
```

These instructions clear various flags. `CLC` clears the carry flag; `CLD` clears the direction flag; `CLI` clears the interrupt flag (thus disabling interrupts); and `CLTS` clears the task–switched (`TS`) flag in `CR0`.

To set the carry, direction, or interrupt flags, use the `STC`, `STD` and `STI` instructions (section B.4.301). To invert the carry flag, use `CMC` (section B.4.22).

B.4.21 CLFLUSH: Flush Cache Line

```
CLFLUSH mem                     ; 0F AE /7          [WILLAMETTE,SSE2]
```

`CLFLUSH` invalidates the cache line that contains the linear address specified by the source operand from all levels of the processor cache hierarchy (data and instruction). If, at any level of the cache

hierarchy, the line is inconsistent with memory (dirty) it is written to memory before invalidation. The source operand points to a byte−sized memory location.

Although CLFLUSH is flagged SSE2 and above, it may not be present on all processors which have SSE2 support, and it may be supported on other processors; the CPUID instruction (section B.4.34) will return a bit which indicates support for the CLFLUSH instruction.

B.4.22 CMC: Complement Carry Flag

```
CMC                             ; F5                    [8086]
```

CMC changes the value of the carry flag: if it was 0, it sets it to 1, and vice versa.

B.4.23 CMOVcc: Conditional Move

```
CMOVcc reg16,r/m16              ; o16 0F 40+cc /r       [P6]
CMOVcc reg32,r/m32              ; o32 0F 40+cc /r       [P6]
```

CMOV moves its source (second) operand into its destination (first) operand if the given condition code is satisfied; otherwise it does nothing.

For a list of condition codes, see section B.2.2.

Although the CMOV instructions are flagged P6 and above, they may not be supported by all Pentium Pro processors; the CPUID instruction (section B.4.34) will return a bit which indicates whether conditional moves are supported.

B.4.24 CMP: Compare Integers

```
CMP r/m8,reg8                   ; 38 /r                 [8086]
CMP r/m16,reg16                 ; o16 39 /r             [8086]
CMP r/m32,reg32                 ; o32 39 /r             [386]

CMP reg8,r/m8                   ; 3A /r                 [8086]
CMP reg16,r/m16                 ; o16 3B /r             [8086]
CMP reg32,r/m32                 ; o32 3B /r             [386]

CMP r/m8,imm8                   ; 80 /7 ib              [8086]
CMP r/m16,imm16                 ; o16 81 /7 iw          [8086]
CMP r/m32,imm32                 ; o32 81 /7 id          [386]

CMP r/m16,imm8                  ; o16 83 /7 ib          [8086]
CMP r/m32,imm8                  ; o32 83 /7 ib          [386]

CMP AL,imm8                     ; 3C ib                 [8086]
CMP AX,imm16                    ; o16 3D iw             [8086]
CMP EAX,imm32                   ; o32 3D id             [386]
```

CMP performs a 'mental' subtraction of its second operand from its first operand, and affects the flags as if the subtraction had taken place, but does not store the result of the subtraction anywhere.

In the forms with an 8−bit immediate second operand and a longer first operand, the second operand is considered to be signed, and is sign−extended to the length of the first operand. In these cases, the BYTE qualifier is necessary to force NASM to generate this form of the instruction.

The destination operand can be a register or a memory location. The source can be a register, memory location or an immediate value of the same size as the destination.

B.4.25 CMPccPD: Packed Double−Precision FP Compare

```
CMPPD xmm1,xmm2/mem128,imm8     ; 66 0F C2 /r ib  [WILLAMETTE,SSE2]
```

```
CMPEQPD xmm1,xmm2/mem128       ; 66 0F C2 /r 00  [WILLAMETTE,SSE2]
CMPLTPD xmm1,xmm2/mem128       ; 66 0F C2 /r 01  [WILLAMETTE,SSE2]
CMPLEPD xmm1,xmm2/mem128       ; 66 0F C2 /r 02  [WILLAMETTE,SSE2]
CMPUNORDPD xmm1,xmm2/mem128    ; 66 0F C2 /r 03  [WILLAMETTE,SSE2]
CMPNEQPD xmm1,xmm2/mem128      ; 66 0F C2 /r 04  [WILLAMETTE,SSE2]
CMPNLTPD xmm1,xmm2/mem128      ; 66 0F C2 /r 05  [WILLAMETTE,SSE2]
CMPNLEPD xmm1,xmm2/mem128      ; 66 0F C2 /r 06  [WILLAMETTE,SSE2]
CMPORDPD xmm1,xmm2/mem128      ; 66 0F C2 /r 07  [WILLAMETTE,SSE2]
```

The CMPccPD instructions compare the two packed double–precision FP values in the source and destination operands, and returns the result of the comparison in the destination register. The result of each comparison is a quadword mask of all 1s (comparison true) or all 0s (comparison false).

The destination is an XMM register. The source can be either an XMM register or a 128–bit memory location.

The third operand is an 8–bit immediate value, of which the low 3 bits define the type of comparison. For ease of programming, the 8 two–operand pseudo–instructions are provided, with the third operand already filled in. The Condition Predicates are:

```
EQ      0   Equal
LT      1   Less-than
LE      2   Less-than-or-equal
UNORD   3   Unordered
NE      4   Not-equal
NLT     5   Not-less-than
NLE     6   Not-less-than-or-equal
ORD     7   Ordered
```

For more details of the comparison predicates, and details of how to emulate the "greater–than" equivalents, see section B.2.3

B.4.26 CMPccPS: Packed Single–Precision FP Compare

```
CMPPS xmm1,xmm2/mem128,imm8    ; 0F C2 /r ib     [KATMAI,SSE]

CMPEQPS xmm1,xmm2/mem128       ; 0F C2 /r 00     [KATMAI,SSE]
CMPLTPS xmm1,xmm2/mem128       ; 0F C2 /r 01     [KATMAI,SSE]
CMPLEPS xmm1,xmm2/mem128       ; 0F C2 /r 02     [KATMAI,SSE]
CMPUNORDPS xmm1,xmm2/mem128    ; 0F C2 /r 03     [KATMAI,SSE]
CMPNEQPS xmm1,xmm2/mem128      ; 0F C2 /r 04     [KATMAI,SSE]
CMPNLTPS xmm1,xmm2/mem128      ; 0F C2 /r 05     [KATMAI,SSE]
CMPNLEPS xmm1,xmm2/mem128      ; 0F C2 /r 06     [KATMAI,SSE]
CMPORDPS xmm1,xmm2/mem128      ; 0F C2 /r 07     [KATMAI,SSE]
```

The CMPccPS instructions compare the two packed single–precision FP values in the source and destination operands, and returns the result of the comparison in the destination register. The result of each comparison is a doubleword mask of all 1s (comparison true) or all 0s (comparison false).

The destination is an XMM register. The source can be either an XMM register or a 128–bit memory location.

The third operand is an 8–bit immediate value, of which the low 3 bits define the type of comparison. For ease of programming, the 8 two–operand pseudo–instructions are provided, with the third operand already filled in. The Condition Predicates are:

```
EQ      0   Equal
LT      1   Less-than
LE      2   Less-than-or-equal
```

```
UNORD   3   Unordered
NE      4   Not-equal
NLT     5   Not-less-than
NLE     6   Not-less-than-or-equal
ORD     7   Ordered
```

For more details of the comparison predicates, and details of how to emulate the "greater–than" equivalents, see section B.2.3

B.4.27 CMPSB, CMPSW, CMPSD: Compare Strings

```
CMPSB                           ; A6                    [8086]
CMPSW                           ; o16 A7                [8086]
CMPSD                           ; o32 A7                [386]
```

CMPSB compares the byte at [DS:SI] or [DS:ESI] with the byte at [ES:DI] or [ES:EDI], and sets the flags accordingly. It then increments or decrements (depending on the direction flag: increments if the flag is clear, decrements if it is set) SI and DI (or ESI and EDI).

The registers used are SI and DI if the address size is 16 bits, and ESI and EDI if it is 32 bits. If you need to use an address size not equal to the current BITS setting, you can use an explicit a16 or a32 prefix.

The segment register used to load from [SI] or [ESI] can be overridden by using a segment register name as a prefix (for example, ES CMPSB). The use of ES for the load from [DI] or [EDI] cannot be overridden.

CMPSW and CMPSD work in the same way, but they compare a word or a doubleword instead of a byte, and increment or decrement the addressing registers by 2 or 4 instead of 1.

The REPE and REPNE prefixes (equivalently, REPZ and REPNZ) may be used to repeat the instruction up to CX (or ECX – again, the address size chooses which) times until the first unequal or equal byte is found.

B.4.28 CMPccSD: Scalar Double–Precision FP Compare

```
CMPSD xmm1,xmm2/mem64,imm8      ; F2 0F C2 /r ib  [WILLAMETTE,SSE2]

CMPEQSD xmm1,xmm2/mem64         ; F2 0F C2 /r 00  [WILLAMETTE,SSE2]
CMPLTSD xmm1,xmm2/mem64         ; F2 0F C2 /r 01  [WILLAMETTE,SSE2]
CMPLESD xmm1,xmm2/mem64         ; F2 0F C2 /r 02  [WILLAMETTE,SSE2]
CMPUNORDSD xmm1,xmm2/mem64      ; F2 0F C2 /r 03  [WILLAMETTE,SSE2]
CMPNEQSD xmm1,xmm2/mem64        ; F2 0F C2 /r 04  [WILLAMETTE,SSE2]
CMPNLTSD xmm1,xmm2/mem64        ; F2 0F C2 /r 05  [WILLAMETTE,SSE2]
CMPNLESD xmm1,xmm2/mem64        ; F2 0F C2 /r 06  [WILLAMETTE,SSE2]
CMPORDSD xmm1,xmm2/mem64        ; F2 0F C2 /r 07  [WILLAMETTE,SSE2]
```

The CMPccSD instructions compare the low–order double–precision FP values in the source and destination operands, and returns the result of the comparison in the destination register. The result of each comparison is a quadword mask of all 1s (comparison true) or all 0s (comparison false).

The destination is an XMM register. The source can be either an XMM register or a 128–bit memory location.

The third operand is an 8–bit immediate value, of which the low 3 bits define the type of comparison. For ease of programming, the 8 two–operand pseudo–instructions are provided, with the third operand already filled in. The Condition Predicates are:

```
EQ      0   Equal
LT      1   Less-than
```

```
LE      2   Less-than-or-equal
UNORD   3   Unordered
NE      4   Not-equal
NLT     5   Not-less-than
NLE     6   Not-less-than-or-equal
ORD     7   Ordered
```

For more details of the comparison predicates, and details of how to emulate the "greater−than" equivalents, see section B.2.3

B.4.29 CMPccSS: Scalar Single−Precision FP Compare

```
CMPSS xmm1,xmm2/mem32,imm8     ; F3 0F C2 /r ib  [KATMAI,SSE]

CMPEQSS xmm1,xmm2/mem32        ; F3 0F C2 /r 00  [KATMAI,SSE]
CMPLTSS xmm1,xmm2/mem32        ; F3 0F C2 /r 01  [KATMAI,SSE]
CMPLESS xmm1,xmm2/mem32        ; F3 0F C2 /r 02  [KATMAI,SSE]
CMPUNORDSS xmm1,xmm2/mem32     ; F3 0F C2 /r 03  [KATMAI,SSE]
CMPNEQSS xmm1,xmm2/mem32       ; F3 0F C2 /r 04  [KATMAI,SSE]
CMPNLTSS xmm1,xmm2/mem32       ; F3 0F C2 /r 05  [KATMAI,SSE]
CMPNLESS xmm1,xmm2/mem32       ; F3 0F C2 /r 06  [KATMAI,SSE]
CMPORDSS xmm1,xmm2/mem32       ; F3 0F C2 /r 07  [KATMAI,SSE]
```

The CMPccSS instructions compare the low−order single−precision FP values in the source and destination operands, and returns the result of the comparison in the destination register. The result of each comparison is a doubleword mask of all 1s (comparison true) or all 0s (comparison false).

The destination is an XMM register. The source can be either an XMM register or a 128−bit memory location.

The third operand is an 8−bit immediate value, of which the low 3 bits define the type of comparison. For ease of programming, the 8 two−operand pseudo−instructions are provided, with the third operand already filled in. The Condition Predicates are:

```
EQ      0   Equal
LT      1   Less-than
LE      2   Less-than-or-equal
UNORD   3   Unordered
NE      4   Not-equal
NLT     5   Not-less-than
NLE     6   Not-less-than-or-equal
ORD     7   Ordered
```

For more details of the comparison predicates, and details of how to emulate the "greater−than" equivalents, see section B.2.3

B.4.30 CMPXCHG, CMPXCHG486: Compare and Exchange

```
CMPXCHG r/m8,reg8              ; 0F B0 /r              [PENT]
CMPXCHG r/m16,reg16            ; o16 0F B1 /r          [PENT]
CMPXCHG r/m32,reg32            ; o32 0F B1 /r          [PENT]

CMPXCHG486 r/m8,reg8           ; 0F A6 /r              [486,UNDOC]
CMPXCHG486 r/m16,reg16         ; o16 0F A7 /r          [486,UNDOC]
CMPXCHG486 r/m32,reg32         ; o32 0F A7 /r          [486,UNDOC]
```

These two instructions perform exactly the same operation; however, apparently some (not all) 486 processors support it under a non−standard opcode, so NASM provides the undocumented CMPXCHG486 form to generate the non−standard opcode.

CMPXCHG compares its destination (first) operand to the value in AL, AX or EAX (depending on the operand size of the instruction). If they are equal, it copies its source (second) operand into the destination and sets the zero flag. Otherwise, it clears the zero flag and copies the destination register to AL, AX or EAX.

The destination can be either a register or a memory location. The source is a register.

CMPXCHG is intended to be used for atomic operations in multitasking or multiprocessor environments. To safely update a value in shared memory, for example, you might load the value into EAX, load the updated value into EBX, and then execute the instruction LOCK CMPXCHG [value],EBX. If value has not changed since being loaded, it is updated with your desired new value, and the zero flag is set to let you know it has worked. (The LOCK prefix prevents another processor doing anything in the middle of this operation: it guarantees atomicity.) However, if another processor has modified the value in between your load and your attempted store, the store does not happen, and you are notified of the failure by a cleared zero flag, so you can go round and try again.

B.4.31 CMPXCHG8B: Compare and Exchange Eight Bytes

```
CMPXCHG8B mem                   ; 0F C7 /1              [PENT]
```

This is a larger and more unwieldy version of CMPXCHG: it compares the 64–bit (eight–byte) value stored at [mem] with the value in EDX:EAX. If they are equal, it sets the zero flag and stores ECX:EBX into the memory area. If they are unequal, it clears the zero flag and stores the memory contents into EDX:EAX.

CMPXCHG8B can be used with the LOCK prefix, to allow atomic execution. This is useful in multi–processor and multi–tasking environments.

B.4.32 COMISD: Scalar Ordered Double–Precision FP Compare and Set EFLAGS

```
COMISD xmm1,xmm2/mem64          ; 66 0F 2F /r     [WILLAMETTE,SSE2]
```

COMISD compares the low–order double–precision FP value in the two source operands. ZF, PF and CF are set according to the result. OF, AF and AF are cleared. The unordered result is returned if either source is a NaN (QNaN or SNaN).

The destination operand is an XMM register. The source can be either an XMM register or a memory location.

The flags are set according to the following rules:

```
Result          Flags        Values

UNORDERED:      ZF,PF,CF <-- 111;
GREATER_THAN:   ZF,PF,CF <-- 000;
LESS_THAN:      ZF,PF,CF <-- 001;
EQUAL:          ZF,PF,CF <-- 100;
```

B.4.33 COMISS: Scalar Ordered Single–Precision FP Compare and Set EFLAGS

```
COMISS xmm1,xmm2/mem32          ; 66 0F 2F /r     [KATMAI,SSE]
```

COMISS compares the low–order single–precision FP value in the two source operands. ZF, PF and CF are set according to the result. OF, AF and AF are cleared. The unordered result is returned if either source is a NaN (QNaN or SNaN).

The destination operand is an XMM register. The source can be either an XMM register or a memory location.

The flags are set according to the following rules:

```
Result          Flags        Values

UNORDERED:      ZF,PF,CF <-- 111;
GREATER_THAN:   ZF,PF,CF <-- 000;
LESS_THAN:      ZF,PF,CF <-- 001;
EQUAL:          ZF,PF,CF <-- 100;
```

B.4.34 CPUID: Get CPU Identification Code

```
CPUID                           ; 0F A2                 [PENT]
```

CPUID returns various information about the processor it is being executed on. It fills the four registers EAX, EBX, ECX and EDX with information, which varies depending on the input contents of EAX.

CPUID also acts as a barrier to serialise instruction execution: executing the CPUID instruction guarantees that all the effects (memory modification, flag modification, register modification) of previous instructions have been completed before the next instruction gets fetched.

The information returned is as follows:

- If EAX is zero on input, EAX on output holds the maximum acceptable input value of EAX, and EBX:EDX:ECX contain the string "GenuineIntel" (or not, if you have a clone processor). That is to say, EBX contains "Genu" (in NASM's own sense of character constants, described in section 3.4.2), EDX contains "ineI" and ECX contains "ntel".
- If EAX is one on input, EAX on output contains version information about the processor, and EDX contains a set of feature flags, showing the presence and absence of various features. For example, bit 8 is set if the CMPXCHG8B instruction (section B.4.31) is supported, bit 15 is set if the conditional move instructions (section B.4.23 and section B.4.72) are supported, and bit 23 is set if MMX instructions are supported.
- If EAX is two on input, EAX, EBX, ECX and EDX all contain information about caches and TLBs (Translation Lookahead Buffers).

For more information on the data returned from CPUID, see the documentation from Intel and other processor manufacturers.

B.4.35 CVTDQ2PD: Packed Signed INT32 to Packed Double–Precision FP Conversion

```
CVTDQ2PD xmm1,xmm2/mem64        ; F3 0F E6 /r       [WILLAMETTE,SSE2]
```

CVTDQ2PD converts two packed signed doublewords from the source operand to two packed double–precision FP values in the destination operand.

The destination operand is an XMM register. The source can be either an XMM register or a 64–bit memory location. If the source is a register, the packed integers are in the low quadword.

B.4.36 CVTDQ2PS: Packed Signed INT32 to Packed Single–Precision FP Conversion

```
CVTDQ2PS xmm1,xmm2/mem128       ; 0F 5B /r          [WILLAMETTE,SSE2]
```

CVTDQ2PS converts four packed signed doublewords from the source operand to four packed single–precision FP values in the destination operand.

The destination operand is an XMM register. The source can be either an XMM register or a 128–bit memory location.

For more details of this instruction, see the Intel Processor manuals.

B.4.37 CVTPD2DQ: Packed Double–Precision FP to Packed Signed INT32 Conversion

```
CVTPD2DQ xmm1,xmm2/mem128       ; F2 0F E6 /r       [WILLAMETTE,SSE2]
```

CVTPD2DQ converts two packed double−precision FP values from the source operand to two packed signed doublewords in the low quadword of the destination operand. The high quadword of the destination is set to all 0s.

The destination operand is an XMM register. The source can be either an XMM register or a 128−bit memory location.

For more details of this instruction, see the Intel Processor manuals.

B.4.38 CVTPD2PI: Packed Double−Precision FP to Packed Signed INT32 Conversion

```
CVTPD2PI mm,xmm/mem128          ; 66 0F 2D /r      [WILLAMETTE,SSE2]
```

CVTPD2PI converts two packed double−precision FP values from the source operand to two packed signed doublewords in the destination operand.

The destination operand is an MMX register. The source can be either an XMM register or a 128−bit memory location.

For more details of this instruction, see the Intel Processor manuals.

B.4.39 CVTPD2PS: Packed Double−Precision FP to Packed Single−Precision FP Conversion

```
CVTPD2PS xmm1,xmm2/mem128       ; 66 0F 5A /r      [WILLAMETTE,SSE2]
```

CVTPD2PS converts two packed double−precision FP values from the source operand to two packed single−precision FP values in the low quadword of the destination operand. The high quadword of the destination is set to all 0s.

The destination operand is an XMM register. The source can be either an XMM register or a 128−bit memory location.

For more details of this instruction, see the Intel Processor manuals.

B.4.40 CVTPI2PD: Packed Signed INT32 to Packed Double−Precision FP Conversion

```
CVTPI2PD xmm,mm/mem64           ; 66 0F 2A /r      [WILLAMETTE,SSE2]
```

CVTPI2PD converts two packed signed doublewords from the source operand to two packed double−precision FP values in the destination operand.

The destination operand is an XMM register. The source can be either an MMX register or a 64−bit memory location.

For more details of this instruction, see the Intel Processor manuals.

B.4.41 CVTPI2PS: Packed Signed INT32 to Packed Single−FP Conversion

```
CVTPI2PS xmm,mm/mem64           ; 0F 2A /r         [KATMAI,SSE]
```

CVTPI2PS converts two packed signed doublewords from the source operand to two packed single−precision FP values in the low quadword of the destination operand. The high quadword of the destination remains unchanged.

The destination operand is an XMM register. The source can be either an MMX register or a 64−bit memory location.

For more details of this instruction, see the Intel Processor manuals.

B.4.42 CVTPS2DQ: Packed Single−Precision FP to Packed Signed INT32 Conversion

```
CVTPS2DQ xmm1,xmm2/mem128       ; 66 0F 5B /r      [WILLAMETTE,SSE2]
```

CVTPS2DQ converts four packed single–precision FP values from the source operand to four packed signed doublewords in the destination operand.

The destination operand is an XMM register. The source can be either an XMM register or a 128–bit memory location.

For more details of this instruction, see the Intel Processor manuals.

B.4.43 CVTPS2PD: Packed Single–Precision FP to Packed Double–Precision FP Conversion

```
CVTPS2PD xmm1,xmm2/mem64        ; 0F 5A /r          [WILLAMETTE,SSE2]
```

CVTPS2PD converts two packed single–precision FP values from the source operand to two packed double–precision FP values in the destination operand.

The destination operand is an XMM register. The source can be either an XMM register or a 64–bit memory location. If the source is a register, the input values are in the low quadword.

For more details of this instruction, see the Intel Processor manuals.

B.4.44 CVTPS2PI: Packed Single–Precision FP to Packed Signed INT32 Conversion

```
CVTPS2PI mm,xmm/mem64           ; 0F 2D /r          [KATMAI,SSE]
```

CVTPS2PI converts two packed single–precision FP values from the source operand to two packed signed doublewords in the destination operand.

The destination operand is an MMX register. The source can be either an XMM register or a 64–bit memory location. If the source is a register, the input values are in the low quadword.

For more details of this instruction, see the Intel Processor manuals.

B.4.45 CVTSD2SI: Scalar Double–Precision FP to Signed INT32 Conversion

```
CVTSD2SI reg32,xmm/mem64        ; F2 0F 2D /r       [WILLAMETTE,SSE2]
```

CVTSD2SI converts a double–precision FP value from the source operand to a signed doubleword in the destination operand.

The destination operand is a general purpose register. The source can be either an XMM register or a 64–bit memory location. If the source is a register, the input value is in the low quadword.

For more details of this instruction, see the Intel Processor manuals.

B.4.46 CVTSD2SS: Scalar Double–Precision FP to Scalar Single–Precision FP Conversion

```
CVTSD2SS xmm1,xmm2/mem64        ; F2 0F 5A /r       [KATMAI,SSE]
```

CVTSD2SS converts a double–precision FP value from the source operand to a single–precision FP value in the low doubleword of the destination operand. The upper 3 doublewords are left unchanged.

The destination operand is an XMM register. The source can be either an XMM register or a 64–bit memory location. If the source is a register, the input value is in the low quadword.

For more details of this instruction, see the Intel Processor manuals.

B.4.47 CVTSI2SD: Signed INT32 to Scalar Double–Precision FP Conversion

```
CVTSI2SD xmm,r/m32              ; F2 0F 2A /r       [WILLAMETTE,SSE2]
```

CVTSI2SD converts a signed doubleword from the source operand to a double–precision FP value in the low quadword of the destination operand. The high quadword is left unchanged.

The destination operand is an XMM register. The source can be either a general purpose register or a 32–bit memory location.

For more details of this instruction, see the Intel Processor manuals.

B.4.48 CVTSI2SS: Signed INT32 to Scalar Single–Precision FP Conversion

```
CVTSI2SS xmm,r/m32              ; F3 0F 2A /r      [KATMAI,SSE]
```

CVTSI2SS converts a signed doubleword from the source operand to a single–precision FP value in the low doubleword of the destination operand. The upper 3 doublewords are left unchanged.

The destination operand is an XMM register. The source can be either a general purpose register or a 32–bit memory location.

For more details of this instruction, see the Intel Processor manuals.

B.4.49 CVTSS2SD: Scalar Single–Precision FP to Scalar Double–Precision FP Conversion

```
CVTSS2SD xmm1,xmm2/mem32        ; F3 0F 5A /r      [WILLAMETTE,SSE2]
```

CVTSS2SD converts a single–precision FP value from the source operand to a double–precision FP value in the low quadword of the destination operand. The upper quadword is left unchanged.

The destination operand is an XMM register. The source can be either an XMM register or a 32–bit memory location. If the source is a register, the input value is contained in the low doubleword.

For more details of this instruction, see the Intel Processor manuals.

B.4.50 CVTSS2SI: Scalar Single–Precision FP to Signed INT32 Conversion

```
CVTSS2SI reg32,xmm/mem32        ; F3 0F 2D /r      [KATMAI,SSE]
```

CVTSS2SI converts a single–precision FP value from the source operand to a signed doubleword in the destination operand.

The destination operand is a general purpose register. The source can be either an XMM register or a 32–bit memory location. If the source is a register, the input value is in the low doubleword.

For more details of this instruction, see the Intel Processor manuals.

B.4.51 CVTTPD2DQ: Packed Double–Precision FP to Packed Signed INT32 Conversion with Truncation

```
CVTTPD2DQ xmm1,xmm2/mem128      ; 66 0F E6 /r      [WILLAMETTE,SSE2]
```

CVTTPD2DQ converts two packed double–precision FP values in the source operand to two packed single–precision FP values in the destination operand. If the result is inexact, it is truncated (rounded toward zero). The high quadword is set to all 0s.

The destination operand is an XMM register. The source can be either an XMM register or a 128–bit memory location.

For more details of this instruction, see the Intel Processor manuals.

B.4.52 CVTTPD2PI: Packed Double–Precision FP to Packed Signed INT32 Conversion with Truncation

```
CVTTPD2PI mm,xmm/mem128         ; 66 0F 2C /r      [WILLAMETTE,SSE2]
```

CVTTPD2PI converts two packed double–precision FP values in the source operand to two packed single–precision FP values in the destination operand. If the result is inexact, it is truncated (rounded toward zero).

The destination operand is an MMX register. The source can be either an XMM register or a 128–bit memory location.

For more details of this instruction, see the Intel Processor manuals.

B.4.53 CVTTPS2DQ: Packed Single–Precision FP to Packed Signed INT32 Conversion with Truncation

```
CVTTPS2DQ xmm1,xmm2/mem128     ; F3 0F 5B /r      [WILLAMETTE,SSE2]
```

CVTTPS2DQ converts four packed single–precision FP values in the source operand to four packed signed doublewords in the destination operand. If the result is inexact, it is truncated (rounded toward zero).

The destination operand is an XMM register. The source can be either an XMM register or a 128–bit memory location.

For more details of this instruction, see the Intel Processor manuals.

B.4.54 CVTTPS2PI: Packed Single–Precision FP to Packed Signed INT32 Conversion with Truncation

```
CVTTPS2PI mm,xmm/mem64          ; 0F 2C /r        [KATMAI,SSE]
```

CVTTPS2PI converts two packed single–precision FP values in the source operand to two packed signed doublewords in the destination operand. If the result is inexact, it is truncated (rounded toward zero). If the source is a register, the input values are in the low quadword.

The destination operand is an MMX register. The source can be either an XMM register or a 64–bit memory location. If the source is a register, the input value is in the low quadword.

For more details of this instruction, see the Intel Processor manuals.

B.4.55 CVTTSD2SI: Scalar Double–Precision FP to Signed INT32 Conversion with Truncation

```
CVTTSD2SI reg32,xmm/mem64       ; F2 0F 2C /r     [WILLAMETTE,SSE2]
```

CVTTSD2SI converts a double–precision FP value in the source operand to a signed doubleword in the destination operand. If the result is inexact, it is truncated (rounded toward zero).

The destination operand is a general purpose register. The source can be either an XMM register or a 64–bit memory location. If the source is a register, the input value is in the low quadword.

For more details of this instruction, see the Intel Processor manuals.

B.4.56 CVTTSS2SI: Scalar Single–Precision FP to Signed INT32 Conversion with Truncation

```
CVTTSD2SI reg32,xmm/mem32       ; F3 0F 2C /r     [KATMAI,SSE]
```

CVTTSS2SI converts a single–precision FP value in the source operand to a signed doubleword in the destination operand. If the result is inexact, it is truncated (rounded toward zero).

The destination operand is a general purpose register. The source can be either an XMM register or a 32–bit memory location. If the source is a register, the input value is in the low doubleword.

For more details of this instruction, see the Intel Processor manuals.

B.4.57 DAA, DAS: Decimal Adjustments

```
DAA                             ; 27                [8086]
DAS                             ; 2F                [8086]
```

These instructions are used in conjunction with the add and subtract instructions to perform binary−coded decimal arithmetic in *packed* (one BCD digit per nibble) form. For the unpacked equivalents, see section B.4.1.

DAA should be used after a one−byte ADD instruction whose destination was the AL register: by means of examining the value in the AL and also the auxiliary carry flag AF, it determines whether either digit of the addition has overflowed, and adjusts it (and sets the carry and auxiliary−carry flags) if so. You can add long BCD strings together by doing ADD/DAA on the low two digits, then doing ADC/DAA on each subsequent pair of digits.

DAS works similarly to DAA, but is for use after SUB instructions rather than ADD.

B.4.58 DEC: Decrement Integer

```
DEC reg16                       ; o16 48+r              [8086]
DEC reg32                       ; o32 48+r              [386]
DEC r/m8                        ; FE /1                 [8086]
DEC r/m16                       ; o16 FF /1             [8086]
DEC r/m32                       ; o32 FF /1             [386]
```

DEC subtracts 1 from its operand. It does *not* affect the carry flag: to affect the carry flag, use SUB something,1 (see section B.4.305). DEC affects all the other flags according to the result.

This instruction can be used with a LOCK prefix to allow atomic execution.

See also INC (section B.4.120).

B.4.59 DIV: Unsigned Integer Divide

```
DIV r/m8                        ; F6 /6                 [8086]
DIV r/m16                       ; o16 F7 /6             [8086]
DIV r/m32                       ; o32 F7 /6             [386]
```

DIV performs unsigned integer division. The explicit operand provided is the divisor; the dividend and destination operands are implicit, in the following way:

- For DIV r/m8, AX is divided by the given operand; the quotient is stored in AL and the remainder in AH.
- For DIV r/m16, DX:AX is divided by the given operand; the quotient is stored in AX and the remainder in DX.
- For DIV r/m32, EDX:EAX is divided by the given operand; the quotient is stored in EAX and the remainder in EDX.

Signed integer division is performed by the IDIV instruction: see section B.4.117.

B.4.60 DIVPD: Packed Double−Precision FP Divide

```
DIVPD xmm1,xmm2/mem128          ; 66 0F 5E /r     [WILLAMETTE,SSE2]
```

DIVPD divides the two packed double−precision FP values in the destination operand by the two packed double−precision FP values in the source operand, and stores the packed double−precision results in the destination register.

The destination is an XMM register. The source operand can be either an XMM register or a 128−bit memory location.

```
dst[0−63]   := dst[0−63]   / src[0−63],
dst[64−127] := dst[64−127] / src[64−127].
```

B.4.61 DIVPS: Packed Single–Precision FP Divide

```
DIVPS xmm1,xmm2/mem128          ; 0F 5E /r          [KATMAI,SSE]
```

DIVPS divides the four packed single–precision FP values in the destination operand by the four packed single–precision FP values in the source operand, and stores the packed single–precision results in the destination register.

The destination is an XMM register. The source operand can be either an XMM register or a 128–bit memory location.

```
dst[0–31]   := dst[0–31]   / src[0–31],
dst[32–63]  := dst[32–63]  / src[32–63],
dst[64–95]  := dst[64–95]  / src[64–95],
dst[96–127] := dst[96–127] / src[96–127].
```

B.4.62 DIVSD: Scalar Double–Precision FP Divide

```
DIVSD xmm1,xmm2/mem64           ; F2 0F 5E /r       [WILLAMETTE,SSE2]
```

DIVSD divides the low–order double–precision FP value in the destination operand by the low–order double–precision FP value in the source operand, and stores the double–precision result in the destination register.

The destination is an XMM register. The source operand can be either an XMM register or a 64–bit memory location.

```
dst[0–63]   := dst[0–63] / src[0–63],
dst[64–127] remains unchanged.
```

B.4.63 DIVSS: Scalar Single–Precision FP Divide

```
DIVSS xmm1,xmm2/mem32           ; F3 0F 5E /r       [KATMAI,SSE]
```

DIVSS divides the low–order single–precision FP value in the destination operand by the low–order single–precision FP value in the source operand, and stores the single–precision result in the destination register.

The destination is an XMM register. The source operand can be either an XMM register or a 32–bit memory location.

```
dst[0–31]   := dst[0–31] / src[0–31],
dst[32–127] remains unchanged.
```

B.4.64 EMMS: Empty MMX State

```
EMMS                            ; 0F 77                 [PENT,MMX]
```

EMMS sets the FPU tag word (marking which floating–point registers are available) to all ones, meaning all registers are available for the FPU to use. It should be used after executing MMX instructions and before executing any subsequent floating–point operations.

B.4.65 ENTER: Create Stack Frame

```
ENTER imm,imm                   ; C8 iw ib              [186]
```

ENTER constructs a stack frame for a high–level language procedure call. The first operand (the iw in the opcode definition above refers to the first operand) gives the amount of stack space to allocate for local variables; the second (the ib above) gives the nesting level of the procedure (for languages like Pascal, with nested procedures).

The function of ENTER, with a nesting level of zero, is equivalent to

```
PUSH EBP              ; or PUSH BP          in 16 bits
MOV EBP,ESP           ; or MOV BP,SP        in 16 bits
SUB ESP,operand1      ; or SUB SP,operand1  in 16 bits
```

This creates a stack frame with the procedure parameters accessible upwards from EBP, and local variables accessible downwards from EBP.

With a nesting level of one, the stack frame created is 4 (or 2) bytes bigger, and the value of the final frame pointer EBP is accessible in memory at [EBP-4].

This allows ENTER, when called with a nesting level of two, to look at the stack frame described by the *previous* value of EBP, find the frame pointer at offset –4 from that, and push it along with its new frame pointer, so that when a level–two procedure is called from within a level–one procedure, [EBP-4] holds the frame pointer of the most recent level–one procedure call and [EBP-8] holds that of the most recent level–two call. And so on, for nesting levels up to 31.

Stack frames created by ENTER can be destroyed by the LEAVE instruction: see section B.4.136.

B.4.66 F2XM1: Calculate 2**X–1

```
F2XM1                 ; D9 F0               [8086,FPU]
```

F2XM1 raises 2 to the power of ST0, subtracts one, and stores the result back into ST0. The initial contents of ST0 must be a number in the range –1.0 to +1.0.

B.4.67 FABS: Floating–Point Absolute Value

```
FABS                  ; D9 E1               [8086,FPU]
```

FABS computes the absolute value of ST0,by clearing the sign bit, and stores the result back in ST0.

B.4.68 FADD, FADDP: Floating–Point Addition

```
FADD mem32            ; D8 /0               [8086,FPU]
FADD mem64            ; DC /0               [8086,FPU]

FADD fpureg           ; D8 C0+r             [8086,FPU]
FADD ST0,fpureg       ; D8 C0+r             [8086,FPU]

FADD TO fpureg        ; DC C0+r             [8086,FPU]
FADD fpureg,ST0       ; DC C0+r             [8086,FPU]

FADDP fpureg          ; DE C0+r             [8086,FPU]
FADDP fpureg,ST0      ; DE C0+r             [8086,FPU]
```

- FADD, given one operand, adds the operand to ST0 and stores the result back in ST0. If the operand has the TO modifier, the result is stored in the register given rather than in ST0.
- FADDP performs the same function as FADD TO, but pops the register stack after storing the result.

The given two–operand forms are synonyms for the one–operand forms.

To add an integer value to ST0, use the c{FIADD} instruction (section B.4.80)

B.4.69 FBLD, FBSTP: BCD Floating–Point Load and Store

```
FBLD mem80            ; DF /4               [8086,FPU]
FBSTP mem80           ; DF /6               [8086,FPU]
```

FBLD loads an 80–bit (ten–byte) packed binary–coded decimal number from the given memory address, converts it to a real, and pushes it on the register stack. FBSTP stores the value of ST0, in packed BCD, at the given address and then pops the register stack.

B.4.70 FCHS: Floating–Point Change Sign

```
FCHS                            ; D9 E0                 [8086,FPU]
```

FCHS negates the number in ST0, by inverting the sign bit: negative numbers become positive, and vice versa.

B.4.71 FCLEX, FNCLEX: Clear Floating–Point Exceptions

```
FCLEX                           ; 9B DB E2              [8086,FPU]
FNCLEX                          ; DB E2                 [8086,FPU]
```

FCLEX clears any floating–point exceptions which may be pending. FNCLEX does the same thing but doesn't wait for previous floating–point operations (including the *handling* of pending exceptions) to finish first.

B.4.72 FCMOVcc: Floating–Point Conditional Move

```
FCMOVB fpureg                   ; DA C0+r               [P6,FPU]
FCMOVB ST0,fpureg               ; DA C0+r               [P6,FPU]

FCMOVE fpureg                   ; DA C8+r               [P6,FPU]
FCMOVE ST0,fpureg               ; DA C8+r               [P6,FPU]

FCMOVBE fpureg                  ; DA D0+r               [P6,FPU]
FCMOVBE ST0,fpureg              ; DA D0+r               [P6,FPU]

FCMOVU fpureg                   ; DA D8+r               [P6,FPU]
FCMOVU ST0,fpureg               ; DA D8+r               [P6,FPU]

FCMOVNB fpureg                  ; DB C0+r               [P6,FPU]
FCMOVNB ST0,fpureg              ; DB C0+r               [P6,FPU]

FCMOVNE fpureg                  ; DB C8+r               [P6,FPU]
FCMOVNE ST0,fpureg              ; DB C8+r               [P6,FPU]

FCMOVNBE fpureg                 ; DB D0+r               [P6,FPU]
FCMOVNBE ST0,fpureg             ; DB D0+r               [P6,FPU]

FCMOVNU fpureg                  ; DB D8+r               [P6,FPU]
FCMOVNU ST0,fpureg              ; DB D8+r               [P6,FPU]
```

The FCMOV instructions perform conditional move operations: each of them moves the contents of the given register into ST0 if its condition is satisfied, and does nothing if not.

The conditions are not the same as the standard condition codes used with conditional jump instructions. The conditions B, BE, NB, NBE, E and NE are exactly as normal, but none of the other standard ones are supported. Instead, the condition U and its counterpart NU are provided; the U condition is satisfied if the last two floating–point numbers compared were *unordered*, i.e. they were not equal but neither one could be said to be greater than the other, for example if they were NaNs. (The flag state which signals this is the setting of the parity flag: so the U condition is notionally equivalent to PE, and NU is equivalent to PO.)

The FCMOV conditions test the main processor's status flags, not the FPU status flags, so using FCMOV directly after FCOM will not work. Instead, you should either use FCOMI which writes directly to the main CPU flags word, or use FSTSW to extract the FPU flags.

Although the FCMOV instructions are flagged P6 above, they may not be supported by all Pentium Pro processors; the CPUID instruction (section B.4.34) will return a bit which indicates whether conditional moves are supported.

B.4.73 FCOM, FCOMP, FCOMPP, FCOMI, FCOMIP: Floating–Point Compare

```
FCOM mem32                      ; D8 /2                 [8086,FPU]
FCOM mem64                      ; DC /2                 [8086,FPU]
FCOM fpureg                     ; D8 D0+r               [8086,FPU]
FCOM ST0,fpureg                 ; D8 D0+r               [8086,FPU]

FCOMP mem32                     ; D8 /3                 [8086,FPU]
FCOMP mem64                     ; DC /3                 [8086,FPU]
FCOMP fpureg                    ; D8 D8+r               [8086,FPU]
FCOMP ST0,fpureg                ; D8 D8+r               [8086,FPU]

FCOMPP                          ; DE D9                 [8086,FPU]

FCOMI fpureg                    ; DB F0+r               [P6,FPU]
FCOMI ST0,fpureg                ; DB F0+r               [P6,FPU]

FCOMIP fpureg                   ; DF F0+r               [P6,FPU]
FCOMIP ST0,fpureg               ; DF F0+r               [P6,FPU]
```

FCOM compares ST0 with the given operand, and sets the FPU flags accordingly. ST0 is treated as the left–hand side of the comparison, so that the carry flag is set (for a 'less–than' result) if ST0 is less than the given operand.

FCOMP does the same as FCOM, but pops the register stack afterwards. FCOMPP compares ST0 with ST1 and then pops the register stack twice.

FCOMI and FCOMIP work like the corresponding forms of FCOM and FCOMP, but write their results directly to the CPU flags register rather than the FPU status word, so they can be immediately followed by conditional jump or conditional move instructions.

The FCOM instructions differ from the FUCOM instructions (section B.4.108) only in the way they handle quiet NaNs: FUCOM will handle them silently and set the condition code flags to an 'unordered' result, whereas FCOM will generate an exception.

B.4.74 FCOS: Cosine

```
FCOS                            ; D9 FF                 [386,FPU]
```

FCOS computes the cosine of ST0 (in radians), and stores the result in ST0. The absolute value of ST0 must be less than 2**63.

See also FSINCOS (section B.4.100).

B.4.75 FDECSTP: Decrement Floating–Point Stack Pointer

```
FDECSTP                         ; D9 F6                 [8086,FPU]
```

FDECSTP decrements the 'top' field in the floating–point status word. This has the effect of rotating the FPU register stack by one, as if the contents of ST7 had been pushed on the stack. See also FINCSTP (section B.4.85).

B.4.76 FxDISI, FxENI: Disable and Enable Floating–Point Interrupts

```
FDISI                           ; 9B DB E1              [8086,FPU]
FNDISI                          ; DB E1                 [8086,FPU]

FENI                            ; 9B DB E0              [8086,FPU]
FNENI                           ; DB E0                 [8086,FPU]
```

FDISI and FENI disable and enable floating–point interrupts. These instructions are only meaningful on original 8087 processors: the 287 and above treat them as no–operation instructions.

FNDISI and FNENI do the same thing as FDISI and FENI respectively, but without waiting for the floating−point processor to finish what it was doing first.

B.4.77 FDIV, FDIVP, FDIVR, FDIVRP: Floating−Point Division

```
FDIV mem32                      ; D8 /6                 [8086,FPU]
FDIV mem64                      ; DC /6                 [8086,FPU]

FDIV fpureg                     ; D8 F0+r               [8086,FPU]
FDIV ST0,fpureg                 ; D8 F0+r               [8086,FPU]

FDIV TO fpureg                  ; DC F8+r               [8086,FPU]
FDIV fpureg,ST0                 ; DC F8+r               [8086,FPU]

FDIVR mem32                     ; D8 /7                 [8086,FPU]
FDIVR mem64                     ; DC /7                 [8086,FPU]

FDIVR fpureg                    ; D8 F8+r               [8086,FPU]
FDIVR ST0,fpureg                ; D8 F8+r               [8086,FPU]

FDIVR TO fpureg                 ; DC F0+r               [8086,FPU]
FDIVR fpureg,ST0                ; DC F0+r               [8086,FPU]

FDIVP fpureg                    ; DE F8+r               [8086,FPU]
FDIVP fpureg,ST0                ; DE F8+r               [8086,FPU]

FDIVRP fpureg                   ; DE F0+r               [8086,FPU]
FDIVRP fpureg,ST0               ; DE F0+r               [8086,FPU]
```

- FDIV divides ST0 by the given operand and stores the result back in ST0, unless the TO qualifier is given, in which case it divides the given operand by ST0 and stores the result in the operand.
- FDIVR does the same thing, but does the division the other way up: so if TO is not given, it divides the given operand by ST0 and stores the result in ST0, whereas if TO is given it divides ST0 by its operand and stores the result in the operand.
- FDIVP operates like FDIV TO, but pops the register stack once it has finished.
- FDIVRP operates like FDIVR TO, but pops the register stack once it has finished.

For FP/Integer divisions, see FIDIV (section B.4.82).

B.4.78 FEMMS: Faster Enter/Exit of the MMX or floating−point state

```
FEMMS                           ; 0F 0E             [PENT,3DNOW]
```

FEMMS can be used in place of the EMMS instruction on processors which support the 3DNow! instruction set. Following execution of FEMMS, the state of the MMX/FP registers is undefined, and this allows a faster context switch between FP and MMX instructions. The FEMMS instruction can also be used *before* executing MMX instructions

B.4.79 FFREE: Flag Floating−Point Register as Unused

```
FFREE fpureg                    ; DD C0+r               [8086,FPU]
FFREEP fpureg                   ; DF C0+r               [286,FPU,UNDOC]
```

FFREE marks the given register as being empty.

FFREEP marks the given register as being empty, and then pops the register stack.

B.4.80 FIADD: Floating–Point/Integer Addition

```
FIADD mem16                     ; DE /0                 [8086,FPU]
FIADD mem32                     ; DA /0                 [8086,FPU]
```

FIADD adds the 16–bit or 32–bit integer stored in the given memory location to ST0, storing the result in ST0.

B.4.81 FICOM, FICOMP: Floating–Point/Integer Compare

```
FICOM mem16                     ; DE /2                 [8086,FPU]
FICOM mem32                     ; DA /2                 [8086,FPU]

FICOMP mem16                    ; DE /3                 [8086,FPU]
FICOMP mem32                    ; DA /3                 [8086,FPU]
```

FICOM compares ST0 with the 16–bit or 32–bit integer stored in the given memory location, and sets the FPU flags accordingly. FICOMP does the same, but pops the register stack afterwards.

B.4.82 FIDIV, FIDIVR: Floating–Point/Integer Division

```
FIDIV mem16                     ; DE /6                 [8086,FPU]
FIDIV mem32                     ; DA /6                 [8086,FPU]

FIDIVR mem16                    ; DE /7                 [8086,FPU]
FIDIVR mem32                    ; DA /7                 [8086,FPU]
```

FIDIV divides ST0 by the 16–bit or 32–bit integer stored in the given memory location, and stores the result in ST0. FIDIVR does the division the other way up: it divides the integer by ST0, but still stores the result in ST0.

B.4.83 FILD, FIST, FISTP: Floating–Point/Integer Conversion

```
FILD mem16                      ; DF /0                 [8086,FPU]
FILD mem32                      ; DB /0                 [8086,FPU]
FILD mem64                      ; DF /5                 [8086,FPU]

FIST mem16                      ; DF /2                 [8086,FPU]
FIST mem32                      ; DB /2                 [8086,FPU]

FISTP mem16                     ; DF /3                 [8086,FPU]
FISTP mem32                     ; DB /3                 [8086,FPU]
FISTP mem64                     ; DF /7                 [8086,FPU]
```

FILD loads an integer out of a memory location, converts it to a real, and pushes it on the FPU register stack. FIST converts ST0 to an integer and stores that in memory; FISTP does the same as FIST, but pops the register stack afterwards.

B.4.84 FIMUL: Floating–Point/Integer Multiplication

```
FIMUL mem16                     ; DE /1                 [8086,FPU]
FIMUL mem32                     ; DA /1                 [8086,FPU]
```

FIMUL multiplies ST0 by the 16–bit or 32–bit integer stored in the given memory location, and stores the result in ST0.

B.4.85 FINCSTP: Increment Floating–Point Stack Pointer

```
FINCSTP                         ; D9 F7                 [8086,FPU]
```

FINCSTP increments the ‘top’ field in the floating–point status word. This has the effect of rotating the FPU register stack by one, as if the register stack had been popped; however, unlike the popping

of the stack performed by many FPU instructions, it does not flag the new ST7 (previously ST0) as empty. See also FDECSTP (section B.4.75).

B.4.86 FINIT, FNINIT: Initialise Floating–Point Unit

```
FINIT                           ; 9B DB E3              [8086,FPU]
FNINIT                          ; DB E3                 [8086,FPU]
```

FINIT initialises the FPU to its default state. It flags all registers as empty, without actually change their values, clears the top of stack pointer. FNINIT does the same, without first waiting for pending exceptions to clear.

B.4.87 FISUB: Floating–Point/Integer Subtraction

```
FISUB mem16                     ; DE /4                 [8086,FPU]
FISUB mem32                     ; DA /4                 [8086,FPU]

FISUBR mem16                    ; DE /5                 [8086,FPU]
FISUBR mem32                    ; DA /5                 [8086,FPU]
```

FISUB subtracts the 16–bit or 32–bit integer stored in the given memory location from ST0, and stores the result in ST0. FISUBR does the subtraction the other way round, i.e. it subtracts ST0 from the given integer, but still stores the result in ST0.

B.4.88 FLD: Floating–Point Load

```
FLD mem32                       ; D9 /0                 [8086,FPU]
FLD mem64                       ; DD /0                 [8086,FPU]
FLD mem80                       ; DB /5                 [8086,FPU]
FLD fpureg                      ; D9 C0+r               [8086,FPU]
```

FLD loads a floating–point value out of the given register or memory location, and pushes it on the FPU register stack.

B.4.89 FLDxx: Floating–Point Load Constants

```
FLD1                            ; D9 E8                 [8086,FPU]
FLDL2E                          ; D9 EA                 [8086,FPU]
FLDL2T                          ; D9 E9                 [8086,FPU]
FLDLG2                          ; D9 EC                 [8086,FPU]
FLDLN2                          ; D9 ED                 [8086,FPU]
FLDPI                           ; D9 EB                 [8086,FPU]
FLDZ                            ; D9 EE                 [8086,FPU]
```

These instructions push specific standard constants on the FPU register stack.

Instruction	Constant pushed
FLD1	1
FLDL2E	base-2 logarithm of e
FLDL2T	base-2 log of 10
FLDLG2	base-10 log of 2
FLDLN2	base-e log of 2
FLDPI	pi
FLDZ	zero

B.4.90 FLDCW: Load Floating–Point Control Word

```
FLDCW mem16                     ; D9 /5                 [8086,FPU]
```

FLDCW loads a 16–bit value out of memory and stores it into the FPU control word (governing things like the rounding mode, the precision, and the exception masks). See also FSTCW (section B.4.103). If exceptions are enabled and you don't want to generate one, use FCLEX or FNCLEX (section B.4.71) before loading the new control word.

B.4.91 FLDENV: Load Floating–Point Environment

```
FLDENV mem                      ; D9 /4                 [8086,FPU]
```

FLDENV loads the FPU operating environment (control word, status word, tag word, instruction pointer, data pointer and last opcode) from memory. The memory area is 14 or 28 bytes long, depending on the CPU mode at the time. See also FSTENV (section B.4.104).

B.4.92 FMUL, FMULP: Floating–Point Multiply

```
FMUL mem32                      ; D8 /1                 [8086,FPU]
FMUL mem64                      ; DC /1                 [8086,FPU]

FMUL fpureg                     ; D8 C8+r               [8086,FPU]
FMUL ST0,fpureg                 ; D8 C8+r               [8086,FPU]

FMUL TO fpureg                  ; DC C8+r               [8086,FPU]
FMUL fpureg,ST0                 ; DC C8+r               [8086,FPU]

FMULP fpureg                    ; DE C8+r               [8086,FPU]
FMULP fpureg,ST0                ; DE C8+r               [8086,FPU]
```

FMUL multiplies ST0 by the given operand, and stores the result in ST0, unless the TO qualifier is used in which case it stores the result in the operand. FMULP performs the same operation as FMUL TO, and then pops the register stack.

B.4.93 FNOP: Floating–Point No Operation

```
FNOP                            ; D9 D0                 [8086,FPU]
```

FNOP does nothing.

B.4.94 FPATAN, FPTAN: Arctangent and Tangent

```
FPATAN                          ; D9 F3                 [8086,FPU]
FPTAN                           ; D9 F2                 [8086,FPU]
```

FPATAN computes the arctangent, in radians, of the result of dividing ST1 by ST0, stores the result in ST1, and pops the register stack. It works like the C atan2 function, in that changing the sign of both ST0 and ST1 changes the output value by pi (so it performs true rectangular–to–polar coordinate conversion, with ST1 being the Y coordinate and ST0 being the X coordinate, not merely an arctangent).

FPTAN computes the tangent of the value in ST0 (in radians), and stores the result back into ST0.

The absolute value of ST0 must be less than 2**63.

B.4.95 FPREM, FPREM1: Floating–Point Partial Remainder

```
FPREM                           ; D9 F8                 [8086,FPU]
FPREM1                          ; D9 F5                 [386,FPU]
```

These instructions both produce the remainder obtained by dividing ST0 by ST1. This is calculated, notionally, by dividing ST0 by ST1, rounding the result to an integer, multiplying by ST1 again, and computing the value which would need to be added back on to the result to get back to the original value in ST0.

The two instructions differ in the way the notional round–to–integer operation is performed. FPREM does it by rounding towards zero, so that the remainder it returns always has the same sign as the original value in ST0; FPREM1 does it by rounding to the nearest integer, so that the remainder always has at most half the magnitude of ST1.

Both instructions calculate *partial* remainders, meaning that they may not manage to provide the final result, but might leave intermediate results in ST0 instead. If this happens, they will set the C2 flag in the FPU status word; therefore, to calculate a remainder, you should repeatedly execute FPREM or FPREM1 until C2 becomes clear.

B.4.96 FRNDINT: Floating–Point Round to Integer

```
FRNDINT                         ; D9 FC                 [8086,FPU]
```

FRNDINT rounds the contents of ST0 to an integer, according to the current rounding mode set in the FPU control word, and stores the result back in ST0.

B.4.97 FSAVE, FRSTOR: Save/Restore Floating–Point State

```
FSAVE mem                       ; 9B DD /6              [8086,FPU]
FNSAVE mem                      ; DD /6                 [8086,FPU]

FRSTOR mem                      ; DD /4                 [8086,FPU]
```

FSAVE saves the entire floating–point unit state, including all the information saved by FSTENV (section B.4.104) plus the contents of all the registers, to a 94 or 108 byte area of memory (depending on the CPU mode). FRSTOR restores the floating–point state from the same area of memory.

FNSAVE does the same as FSAVE, without first waiting for pending floating–point exceptions to clear.

B.4.98 FSCALE: Scale Floating–Point Value by Power of Two

```
FSCALE                          ; D9 FD                 [8086,FPU]
```

FSCALE scales a number by a power of two: it rounds ST1 towards zero to obtain an integer, then multiplies ST0 by two to the power of that integer, and stores the result in ST0.

B.4.99 FSETPM: Set Protected Mode

```
FSETPM                          ; DB E4                 [286,FPU]
```

This instruction initialises protected mode on the 287 floating–point coprocessor. It is only meaningful on that processor: the 387 and above treat the instruction as a no–operation.

B.4.100 FSIN, FSINCOS: Sine and Cosine

```
FSIN                            ; D9 FE                 [386,FPU]
FSINCOS                         ; D9 FB                 [386,FPU]
```

FSIN calculates the sine of ST0 (in radians) and stores the result in ST0. FSINCOS does the same, but then pushes the cosine of the same value on the register stack, so that the sine ends up in ST1 and the cosine in ST0. FSINCOS is faster than executing FSIN and FCOS (see section B.4.74) in succession.

The absolute value of ST0 must be less than 2**63.

B.4.101 FSQRT: Floating–Point Square Root

```
FSQRT                           ; D9 FA                 [8086,FPU]
```

FSQRT calculates the square root of ST0 and stores the result in ST0.

B.4.102 FST, FSTP: Floating–Point Store

```
FST mem32                       ; D9 /2                 [8086,FPU]
FST mem64                       ; DD /2                 [8086,FPU]
FST fpureg                      ; DD D0+r               [8086,FPU]

FSTP mem32                      ; D9 /3                 [8086,FPU]
FSTP mem64                      ; DD /3                 [8086,FPU]
FSTP mem80                      ; DB /7                 [8086,FPU]
FSTP fpureg                     ; DD D8+r               [8086,FPU]
```

FST stores the value in ST0 into the given memory location or other FPU register. FSTP does the same, but then pops the register stack.

B.4.103 FSTCW: Store Floating–Point Control Word

```
FSTCW mem16                     ; 9B D9 /7              [8086,FPU]
FNSTCW mem16                    ; D9 /7                 [8086,FPU]
```

FSTCW stores the FPU control word (governing things like the rounding mode, the precision, and the exception masks) into a 2–byte memory area. See also FLDCW (section B.4.90).

FNSTCW does the same thing as FSTCW, without first waiting for pending floating–point exceptions to clear.

B.4.104 FSTENV: Store Floating–Point Environment

```
FSTENV mem                      ; 9B D9 /6              [8086,FPU]
FNSTENV mem                     ; D9 /6                 [8086,FPU]
```

FSTENV stores the FPU operating environment (control word, status word, tag word, instruction pointer, data pointer and last opcode) into memory. The memory area is 14 or 28 bytes long, depending on the CPU mode at the time. See also FLDENV (section B.4.91).

FNSTENV does the same thing as FSTENV, without first waiting for pending floating–point exceptions to clear.

B.4.105 FSTSW: Store Floating–Point Status Word

```
FSTSW mem16                     ; 9B DD /7              [8086,FPU]
FSTSW AX                        ; 9B DF E0              [286,FPU]

FNSTSW mem16                    ; DD /7                 [8086,FPU]
FNSTSW AX                       ; DF E0                 [286,FPU]
```

FSTSW stores the FPU status word into AX or into a 2–byte memory area.

FNSTSW does the same thing as FSTSW, without first waiting for pending floating–point exceptions to clear.

B.4.106 FSUB, FSUBP, FSUBR, FSUBRP: Floating–Point Subtract

```
FSUB mem32                      ; D8 /4                 [8086,FPU]
FSUB mem64                      ; DC /4                 [8086,FPU]

FSUB fpureg                     ; D8 E0+r               [8086,FPU]
FSUB ST0,fpureg                 ; D8 E0+r               [8086,FPU]

FSUB TO fpureg                  ; DC E8+r               [8086,FPU]
FSUB fpureg,ST0                 ; DC E8+r               [8086,FPU]
```

```
FSUBR mem32                     ; D8 /5                 [8086,FPU]
FSUBR mem64                     ; DC /5                 [8086,FPU]

FSUBR fpureg                    ; D8 E8+r               [8086,FPU]
FSUBR ST0,fpureg                ; D8 E8+r               [8086,FPU]

FSUBR TO fpureg                 ; DC E0+r               [8086,FPU]
FSUBR fpureg,ST0                ; DC E0+r               [8086,FPU]

FSUBP fpureg                    ; DE E8+r               [8086,FPU]
FSUBP fpureg,ST0                ; DE E8+r               [8086,FPU]

FSUBRP fpureg                   ; DE E0+r               [8086,FPU]
FSUBRP fpureg,ST0               ; DE E0+r               [8086,FPU]
```

- `FSUB` subtracts the given operand from `ST0` and stores the result back in `ST0`, unless the `TO` qualifier is given, in which case it subtracts `ST0` from the given operand and stores the result in the operand.
- `FSUBR` does the same thing, but does the subtraction the other way up: so if `TO` is not given, it subtracts `ST0` from the given operand and stores the result in `ST0`, whereas if `TO` is given it subtracts its operand from `ST0` and stores the result in the operand.
- `FSUBP` operates like `FSUB TO`, but pops the register stack once it has finished.
- `FSUBRP` operates like `FSUBR TO`, but pops the register stack once it has finished.

B.4.107 `FTST`: Test `ST0` Against Zero

```
FTST                            ; D9 E4                 [8086,FPU]
```

`FTST` compares `ST0` with zero and sets the FPU flags accordingly. `ST0` is treated as the left–hand side of the comparison, so that a 'less–than' result is generated if `ST0` is negative.

B.4.108 `FUCOMxx`: Floating–Point Unordered Compare

```
FUCOM fpureg                    ; DD E0+r               [386,FPU]
FUCOM ST0,fpureg                ; DD E0+r               [386,FPU]

FUCOMP fpureg                   ; DD E8+r               [386,FPU]
FUCOMP ST0,fpureg               ; DD E8+r               [386,FPU]

FUCOMPP                         ; DA E9                 [386,FPU]

FUCOMI fpureg                   ; DB E8+r               [P6,FPU]
FUCOMI ST0,fpureg               ; DB E8+r               [P6,FPU]

FUCOMIP fpureg                  ; DF E8+r               [P6,FPU]
FUCOMIP ST0,fpureg              ; DF E8+r               [P6,FPU]
```

- `FUCOM` compares `ST0` with the given operand, and sets the FPU flags accordingly. `ST0` is treated as the left–hand side of the comparison, so that the carry flag is set (for a 'less–than' result) if `ST0` is less than the given operand.
- `FUCOMP` does the same as `FUCOM`, but pops the register stack afterwards. `FUCOMPP` compares `ST0` with `ST1` and then pops the register stack twice.
- `FUCOMI` and `FUCOMIP` work like the corresponding forms of `FUCOM` and `FUCOMP`, but write their results directly to the CPU flags register rather than the FPU status word, so they can be immediately followed by conditional jump or conditional move instructions.

The FUCOM instructions differ from the FCOM instructions (section B.4.73) only in the way they handle quiet NaNs: FUCOM will handle them silently and set the condition code flags to an 'unordered' result, whereas FCOM will generate an exception.

B.4.109 FXAM: Examine Class of Value in ST0

```
FXAM                            ; D9 E5                 [8086,FPU]
```

FXAM sets the FPU flags C3, C2 and C0 depending on the type of value stored in ST0:

Register contents	Flags
Unsupported format	000
NaN	001
Finite number	010
Infinity	011
Zero	100
Empty register	101
Denormal	110

Additionally, the C1 flag is set to the sign of the number.

B.4.110 FXCH: Floating–Point Exchange

```
FXCH                            ; D9 C9                 [8086,FPU]
FXCH fpureg                     ; D9 C8+r               [8086,FPU]
FXCH fpureg,ST0                 ; D9 C8+r               [8086,FPU]
FXCH ST0,fpureg                 ; D9 C8+r               [8086,FPU]
```

FXCH exchanges ST0 with a given FPU register. The no–operand form exchanges ST0 with ST1.

B.4.111 FXRSTOR: Restore FP, MMX and SSE State

```
FXRSTOR memory                  ; 0F AE /1                [P6,SSE,FPU]
```

The FXRSTOR instruction reloads the FPU, MMX and SSE state (environment and registers), from the 512 byte memory area defined by the source operand. This data should have been written by a previous FXSAVE.

B.4.112 FXSAVE: Store FP, MMX and SSE State

```
FXSAVE memory                   ; 0F AE /0          [P6,SSE,FPU]
```

FXSAVEThe FXSAVE instruction writes the current FPU, MMX and SSE technology states (environment and registers), to the 512 byte memory area defined by the destination operand. It does this without checking for pending unmasked floating–point exceptions (similar to the operation of FNSAVE).

Unlike the FSAVE/FNSAVE instructions, the processor retains the contents of the FPU, MMX and SSE state in the processor after the state has been saved. This instruction has been optimised to maximize floating–point save performance.

B.4.113 FXTRACT: Extract Exponent and Significand

```
FXTRACT                         ; D9 F4                 [8086,FPU]
```

FXTRACT separates the number in ST0 into its exponent and significand (mantissa), stores the exponent back into ST0, and then pushes the significand on the register stack (so that the significand ends up in ST0, and the exponent in ST1).

B.4.114 FYL2X, FYL2XP1: Compute Y times Log2(X) or Log2(X+1)

```
FYL2X                           ; D9 F1                 [8086,FPU]
FYL2XP1                         ; D9 F9                 [8086,FPU]
```

FYL2X multiplies ST1 by the base−2 logarithm of ST0, stores the result in ST1, and pops the register stack (so that the result ends up in ST0). ST0 must be non−zero and positive.

FYL2XP1 works the same way, but replacing the base−2 log of ST0 with that of ST0 plus one. This time, ST0 must have magnitude no greater than 1 minus half the square root of two.

B.4.115 HLT: Halt Processor

```
HLT                             ; F4                    [8086,PRIV]
```

HLT puts the processor into a halted state, where it will perform no more operations until restarted by an interrupt or a reset.

On the 286 and later processors, this is a privileged instruction.

B.4.116 IBTS: Insert Bit String

```
IBTS r/m16,reg16                ; o16 0F A7 /r          [386,UNDOC]
IBTS r/m32,reg32                ; o32 0F A7 /r          [386,UNDOC]
```

The implied operation of this instruction is:

```
IBTS r/m16,AX,CL,reg16
IBTS r/m32,EAX,CL,reg32
```

Writes a bit string from the source operand to the destination. CL indicates the number of bits to be copied, from the low bits of the source. (E)AX indicates the low order bit offset in the destination that is written to. For example, if CL is set to 4 and AX (for 16−bit code) is set to 5, bits 0−3 of src will be copied to bits 5−8 of dst. This instruction is very poorly documented, and I have been unable to find any official source of documentation on it.

IBTS is supported only on the early Intel 386s, and conflicts with the opcodes for CMPXCHG486 (on early Intel 486s). NASM supports it only for completeness. Its counterpart is XBTS (see section B.4.332).

B.4.117 IDIV: Signed Integer Divide

```
IDIV r/m8                       ; F6 /7                 [8086]
IDIV r/m16                      ; o16 F7 /7             [8086]
IDIV r/m32                      ; o32 F7 /7             [386]
```

IDIV performs signed integer division. The explicit operand provided is the divisor; the dividend and destination operands are implicit, in the following way:

- For IDIV r/m8, AX is divided by the given operand; the quotient is stored in AL and the remainder in AH.
- For IDIV r/m16, DX:AX is divided by the given operand; the quotient is stored in AX and the remainder in DX.
- For IDIV r/m32, EDX:EAX is divided by the given operand; the quotient is stored in EAX and the remainder in EDX.

Unsigned integer division is performed by the DIV instruction: see section B.4.59.

B.4.118 IMUL: Signed Integer Multiply

```
IMUL r/m8                       ; F6 /5                 [8086]
IMUL r/m16                      ; o16 F7 /5             [8086]
IMUL r/m32                      ; o32 F7 /5             [386]

IMUL reg16,r/m16                ; o16 0F AF /r          [386]
IMUL reg32,r/m32                ; o32 0F AF /r          [386]

IMUL reg16,imm8                 ; o16 6B /r ib          [186]
IMUL reg16,imm16                ; o16 69 /r iw          [186]
IMUL reg32,imm8                 ; o32 6B /r ib          [386]
IMUL reg32,imm32                ; o32 69 /r id          [386]

IMUL reg16,r/m16,imm8           ; o16 6B /r ib          [186]
IMUL reg16,r/m16,imm16          ; o16 69 /r iw          [186]
IMUL reg32,r/m32,imm8           ; o32 6B /r ib          [386]
IMUL reg32,r/m32,imm32          ; o32 69 /r id          [386]
```

IMUL performs signed integer multiplication. For the single−operand form, the other operand and destination are implicit, in the following way:

- For IMUL r/m8, AL is multiplied by the given operand; the product is stored in AX.
- For IMUL r/m16, AX is multiplied by the given operand; the product is stored in DX:AX.
- For IMUL r/m32, EAX is multiplied by the given operand; the product is stored in EDX:EAX.

The two−operand form multiplies its two operands and stores the result in the destination (first) operand. The three−operand form multiplies its last two operands and stores the result in the first operand.

The two−operand form with an immediate second operand is in fact a shorthand for the three−operand form, as can be seen by examining the opcode descriptions: in the two−operand form, the code /r takes both its register and r/m parts from the same operand (the first one).

In the forms with an 8−bit immediate operand and another longer source operand, the immediate operand is considered to be signed, and is sign−extended to the length of the other source operand. In these cases, the BYTE qualifier is necessary to force NASM to generate this form of the instruction.

Unsigned integer multiplication is performed by the MUL instruction: see section B.4.184.

B.4.119 IN: Input from I/O Port

```
IN AL,imm8                      ; E4 ib                 [8086]
IN AX,imm8                      ; o16 E5 ib             [8086]
IN EAX,imm8                     ; o32 E5 ib             [386]
IN AL,DX                        ; EC                    [8086]
IN AX,DX                        ; o16 ED                [8086]
IN EAX,DX                       ; o32 ED                [386]
```

IN reads a byte, word or doubleword from the specified I/O port, and stores it in the given destination register. The port number may be specified as an immediate value if it is between 0 and 255, and otherwise must be stored in DX. See also OUT (section B.4.194).

B.4.120 INC: Increment Integer

```
INC reg16                       ; o16 40+r              [8086]
INC reg32                       ; o32 40+r              [386]
INC r/m8                        ; FE /0                 [8086]
```

```
INC r/m16                       ; o16 FF /0             [8086]
INC r/m32                       ; o32 FF /0             [386]
```

INC adds 1 to its operand. It does *not* affect the carry flag: to affect the carry flag, use ADD something,1 (see section B.4.3). INC affects all the other flags according to the result.

This instruction can be used with a LOCK prefix to allow atomic execution.

See also DEC (section B.4.58).

B.4.121 INSB, INSW, INSD: Input String from I/O Port

```
INSB                            ; 6C                    [186]
INSW                            ; o16 6D                [186]
INSD                            ; o32 6D                [386]
```

INSB inputs a byte from the I/O port specified in DX and stores it at [ES:DI] or [ES:EDI]. It then increments or decrements (depending on the direction flag: increments if the flag is clear, decrements if it is set) DI or EDI.

The register used is DI if the address size is 16 bits, and EDI if it is 32 bits. If you need to use an address size not equal to the current BITS setting, you can use an explicit a16 or a32 prefix.

Segment override prefixes have no effect for this instruction: the use of ES for the load from [DI] or [EDI] cannot be overridden.

INSW and INSD work in the same way, but they input a word or a doubleword instead of a byte, and increment or decrement the addressing register by 2 or 4 instead of 1.

The REP prefix may be used to repeat the instruction CX (or ECX – again, the address size chooses which) times.

See also OUTSB, OUTSW and OUTSD (section B.4.195).

B.4.122 INT: Software Interrupt

```
INT imm8                        ; CD ib                 [8086]
```

INT causes a software interrupt through a specified vector number from 0 to 255.

The code generated by the INT instruction is always two bytes long: although there are short forms for some INT instructions, NASM does not generate them when it sees the INT mnemonic. In order to generate single–byte breakpoint instructions, use the INT3 or INT1 instructions (see section B.4.123) instead.

B.4.123 INT3, INT1, ICEBP, INT01: Breakpoints

```
INT1                            ; F1                    [P6]
ICEBP                           ; F1                    [P6]
INT01                           ; F1                    [P6]

INT3                            ; CC                    [8086]
INT03                           ; CC                    [8086]
```

INT1 and INT3 are short one–byte forms of the instructions INT 1 and INT 3 (see section B.4.122). They perform a similar function to their longer counterparts, but take up less code space. They are used as breakpoints by debuggers.

- INT1, and its alternative synonyms INT01 and ICEBP, is an instruction used by in–circuit emulators (ICEs). It is present, though not documented, on some processors down to the 286, but is only documented for the Pentium Pro. INT3 is the instruction normally used as a breakpoint by debuggers.

- INT3, and its synonym INT03, is not precisely equivalent to INT 3: the short form, since it is designed to be used as a breakpoint, bypasses the normal IOPL checks in virtual–8086 mode, and also does not go through interrupt redirection.

B.4.124 INTO: Interrupt if Overflow

```
INTO                            ; CE                    [8086]
```

INTO performs an INT 4 software interrupt (see section B.4.122) if and only if the overflow flag is set.

B.4.125 INVD: Invalidate Internal Caches

```
INVD                            ; 0F 08                 [486]
```

INVD invalidates and empties the processor's internal caches, and causes the processor to instruct external caches to do the same. It does not write the contents of the caches back to memory first: any modified data held in the caches will be lost. To write the data back first, use WBINVD (section B.4.328).

B.4.126 INVLPG: Invalidate TLB Entry

```
INVLPG mem                      ; 0F 01 /7              [486]
```

INVLPG invalidates the translation lookahead buffer (TLB) entry associated with the supplied memory address.

B.4.127 IRET, IRETW, IRETD: Return from Interrupt

```
IRET                            ; CF                    [8086]
IRETW                           ; o16 CF                [8086]
IRETD                           ; o32 CF                [386]
```

IRET returns from an interrupt (hardware or software) by means of popping IP (or EIP), CS and the flags off the stack and then continuing execution from the new CS:IP.

IRETW pops IP, CS and the flags as 2 bytes each, taking 6 bytes off the stack in total. IRETD pops EIP as 4 bytes, pops a further 4 bytes of which the top two are discarded and the bottom two go into CS, and pops the flags as 4 bytes as well, taking 12 bytes off the stack.

IRET is a shorthand for either IRETW or IRETD, depending on the default BITS setting at the time.

B.4.128 Jcc: Conditional Branch

```
Jcc imm                         ; 70+cc rb              [8086]
Jcc NEAR imm                    ; 0F 80+cc rw/rd        [386]
```

The conditional jump instructions execute a near (same segment) jump if and only if their conditions are satisfied. For example, JNZ jumps only if the zero flag is not set.

The ordinary form of the instructions has only a 128–byte range; the NEAR form is a 386 extension to the instruction set, and can span the full size of a segment. NASM will not override your choice of jump instruction: if you want Jcc NEAR, you have to use the NEAR keyword.

The SHORT keyword is allowed on the first form of the instruction, for clarity, but is not necessary.

For details of the condition codes, see section B.2.2.

B.4.129 JCXZ, JECXZ: Jump if CX/ECX Zero

```
JCXZ imm                        ; a16 E3 rb             [8086]
JECXZ imm                       ; a32 E3 rb             [386]
```

JCXZ performs a short jump (with maximum range 128 bytes) if and only if the contents of the CX register is 0. JECXZ does the same thing, but with ECX.

B.4.130 JMP: Jump

```
JMP imm                         ; E9 rw/rd              [8086]
JMP SHORT imm                   ; EB rb                 [8086]
JMP imm:imm16                   ; o16 EA iw iw          [8086]
JMP imm:imm32                   ; o32 EA id iw          [386]
JMP FAR mem                     ; o16 FF /5             [8086]
JMP FAR mem32                   ; o32 FF /5             [386]
JMP r/m16                       ; o16 FF /4             [8086]
JMP r/m32                       ; o32 FF /4             [386]
```

JMP jumps to a given address. The address may be specified as an absolute segment and offset, or as a relative jump within the current segment.

JMP SHORT imm has a maximum range of 128 bytes, since the displacement is specified as only 8 bits, but takes up less code space. NASM does not choose when to generate JMP SHORT for you: you must explicitly code SHORT every time you want a short jump.

You can choose between the two immediate far jump forms (JMP imm:imm) by the use of the WORD and DWORD keywords: JMP WORD 0x1234:0x5678) or JMP DWORD 0x1234:0x56789abc.

The JMP FAR mem forms execute a far jump by loading the destination address out of memory. The address loaded consists of 16 or 32 bits of offset (depending on the operand size), and 16 bits of segment. The operand size may be overridden using JMP WORD FAR mem or JMP DWORD FAR mem.

The JMP r/m forms execute a near jump (within the same segment), loading the destination address out of memory or out of a register. The keyword NEAR may be specified, for clarity, in these forms, but is not necessary. Again, operand size can be overridden using JMP WORD mem or JMP DWORD mem.

As a convenience, NASM does not require you to jump to a far symbol by coding the cumbersome JMP SEG routine:routine, but instead allows the easier synonym JMP FAR routine.

The CALL r/m forms given above are near calls; NASM will accept the NEAR keyword (e.g. CALL NEAR [address]), even though it is not strictly necessary.

B.4.131 LAHF: Load AH from Flags

```
LAHF                            ; 9F                    [8086]
```

LAHF sets the AH register according to the contents of the low byte of the flags word.

The operation of LAHF is:

```
AH <-- SF:ZF:0:AF:0:PF:1:CF
```

See also SAHF (section B.4.282).

B.4.132 LAR: Load Access Rights

```
LAR reg16,r/m16                 ; o16 0F 02 /r          [286,PRIV]
LAR reg32,r/m32                 ; o32 0F 02 /r          [286,PRIV]
```

LAR takes the segment selector specified by its source (second) operand, finds the corresponding segment descriptor in the GDT or LDT, and loads the access–rights byte of the descriptor into its destination (first) operand.

B.4.133 LDMXCSR: Load Streaming SIMD Extension Control/Status

```
LDMXCSR mem32                   ; 0F AE /2          [KATMAI,SSE]
```

LDMXCSR loads 32–bits of data from the specified memory location into the MXCSR control/status register. MXCSR is used to enable masked/unmasked exception handling, to set rounding modes, to set flush–to–zero mode, and to view exception status flags.

For details of the MXCSR register, see the Intel processor docs.

See also STMXCSR (section B.4.302

B.4.134 LDS, LES, LFS, LGS, LSS: Load Far Pointer

```
LDS reg16,mem                   ; o16 C5 /r             [8086]
LDS reg32,mem                   ; o32 C5 /r             [386]

LES reg16,mem                   ; o16 C4 /r             [8086]
LES reg32,mem                   ; o32 C4 /r             [386]

LFS reg16,mem                   ; o16 0F B4 /r          [386]
LFS reg32,mem                   ; o32 0F B4 /r          [386]

LGS reg16,mem                   ; o16 0F B5 /r          [386]
LGS reg32,mem                   ; o32 0F B5 /r          [386]

LSS reg16,mem                   ; o16 0F B2 /r          [386]
LSS reg32,mem                   ; o32 0F B2 /r          [386]
```

These instructions load an entire far pointer (16 or 32 bits of offset, plus 16 bits of segment) out of memory in one go. LDS, for example, loads 16 or 32 bits from the given memory address into the given register (depending on the size of the register), then loads the *next* 16 bits from memory into DS. LES, LFS, LGS and LSS work in the same way but use the other segment registers.

B.4.135 LEA: Load Effective Address

```
LEA reg16,mem                   ; o16 8D /r             [8086]
LEA reg32,mem                   ; o32 8D /r             [386]
```

LEA, despite its syntax, does not access memory. It calculates the effective address specified by its second operand as if it were going to load or store data from it, but instead it stores the calculated address into the register specified by its first operand. This can be used to perform quite complex calculations (e.g. LEA EAX,[EBX+ECX*4+100]) in one instruction.

LEA, despite being a purely arithmetic instruction which accesses no memory, still requires square brackets around its second operand, as if it were a memory reference.

The size of the calculation is the current *address* size, and the size that the result is stored as is the current *operand* size. If the address and operand size are not the same, then if the addressing mode was 32–bits, the low 16–bits are stored, and if the address was 16–bits, it is zero–extended to 32–bits before storing.

B.4.136 LEAVE: Destroy Stack Frame

```
LEAVE                           ; C9                    [186]
```

LEAVE destroys a stack frame of the form created by the ENTER instruction (see section B.4.65). It is functionally equivalent to MOV ESP,EBP followed by POP EBP (or MOV SP,BP followed by POP BP in 16–bit mode).

B.4.137 LFENCE: Load Fence

```
LFENCE                          ; 0F AE /5          [WILLAMETTE,SSE2]
```

LFENCE performs a serialising operation on all loads from memory that were issued before the LFENCE instruction. This guarantees that all memory reads before the LFENCE instruction are visible before any reads after the LFENCE instruction.

LFENCE is ordered respective to other LFENCE instruction, MFENCE, any memory read and any other serialising instruction (such as CPUID).

Weakly ordered memory types can be used to achieve higher processor performance through such techniques as out−of−order issue and speculative reads. The degree to which a consumer of data recognizes or knows that the data is weakly ordered varies among applications and may be unknown to the producer of this data. The LFENCE instruction provides a performance−efficient way of ensuring load ordering between routines that produce weakly−ordered results and routines that consume that data.

LFENCE uses the following ModRM encoding:

```
Mod (7:6)         = 11B
Reg/Opcode (5:3)  = 101B
R/M (2:0)         = 000B
```

All other ModRM encodings are defined to be reserved, and use of these encodings risks incompatibility with future processors.

See also SFENCE (section B.4.288) and MFENCE (section B.4.151).

B.4.138 LGDT, LIDT, LLDT: Load Descriptor Tables

```
LGDT mem                        ; 0F 01 /2              [286,PRIV]
LIDT mem                        ; 0F 01 /3              [286,PRIV]
LLDT r/m16                      ; 0F 00 /2              [286,PRIV]
```

LGDT and LIDT both take a 6−byte memory area as an operand: they load a 32−bit linear address and a 16−bit size limit from that area (in the opposite order) into the GDTR (global descriptor table register) or IDTR (interrupt descriptor table register). These are the only instructions which directly use *linear* addresses, rather than segment/offset pairs.

LLDT takes a segment selector as an operand. The processor looks up that selector in the GDT and stores the limit and base address given there into the LDTR (local descriptor table register).

See also SGDT, SIDT and SLDT (section B.4.289).

B.4.139 LMSW: Load/Store Machine Status Word

```
LMSW r/m16                      ; 0F 01 /6              [286,PRIV]
```

LMSW loads the bottom four bits of the source operand into the bottom four bits of the CR0 control register (or the Machine Status Word, on 286 processors). See also SMSW (section B.4.296).

B.4.140 LOADALL, LOADALL286: Load Processor State

```
LOADALL                         ; 0F 07                 [386,UNDOC]
LOADALL286                      ; 0F 05                 [286,UNDOC]
```

This instruction, in its two different−opcode forms, is apparently supported on most 286 processors, some 386 and possibly some 486. The opcode differs between the 286 and the 386.

The function of the instruction is to load all information relating to the state of the processor out of a block of memory: on the 286, this block is located implicitly at absolute address `0x800`, and on the 386 and 486 it is at `[ES:EDI]`.

B.4.141 LODSB, LODSW, LODSD: Load from String

```
LODSB                           ; AC                    [8086]
LODSW                           ; o16 AD                [8086]
LODSD                           ; o32 AD                [386]
```

`LODSB` loads a byte from `[DS:SI]` or `[DS:ESI]` into `AL`. It then increments or decrements (depending on the direction flag: increments if the flag is clear, decrements if it is set) `SI` or `ESI`.

The register used is `SI` if the address size is 16 bits, and `ESI` if it is 32 bits. If you need to use an address size not equal to the current `BITS` setting, you can use an explicit `a16` or `a32` prefix.

The segment register used to load from `[SI]` or `[ESI]` can be overridden by using a segment register name as a prefix (for example, `ES LODSB`).

`LODSW` and `LODSD` work in the same way, but they load a word or a doubleword instead of a byte, and increment or decrement the addressing registers by 2 or 4 instead of 1.

B.4.142 LOOP, LOOPE, LOOPZ, LOOPNE, LOOPNZ: Loop with Counter

```
LOOP imm                        ; E2 rb                 [8086]
LOOP imm,CX                     ; a16 E2 rb             [8086]
LOOP imm,ECX                    ; a32 E2 rb             [386]

LOOPE imm                       ; E1 rb                 [8086]
LOOPE imm,CX                    ; a16 E1 rb             [8086]
LOOPE imm,ECX                   ; a32 E1 rb             [386]
LOOPZ imm                       ; E1 rb                 [8086]
LOOPZ imm,CX                    ; a16 E1 rb             [8086]
LOOPZ imm,ECX                   ; a32 E1 rb             [386]

LOOPNE imm                      ; E0 rb                 [8086]
LOOPNE imm,CX                   ; a16 E0 rb             [8086]
LOOPNE imm,ECX                  ; a32 E0 rb             [386]
LOOPNZ imm                      ; E0 rb                 [8086]
LOOPNZ imm,CX                   ; a16 E0 rb             [8086]
LOOPNZ imm,ECX                  ; a32 E0 rb             [386]
```

`LOOP` decrements its counter register (either `CX` or `ECX` – if one is not specified explicitly, the `BITS` setting dictates which is used) by one, and if the counter does not become zero as a result of this operation, it jumps to the given label. The jump has a range of 128 bytes.

`LOOPE` (or its synonym `LOOPZ`) adds the additional condition that it only jumps if the counter is nonzero *and* the zero flag is set. Similarly, `LOOPNE` (and `LOOPNZ`) jumps only if the counter is nonzero and the zero flag is clear.

B.4.143 LSL: Load Segment Limit

```
LSL reg16,r/m16                 ; o16 0F 03 /r          [286,PRIV]
LSL reg32,r/m32                 ; o32 0F 03 /r          [286,PRIV]
```

`LSL` is given a segment selector in its source (second) operand; it computes the segment limit value by loading the segment limit field from the associated segment descriptor in the `GDT` or `LDT`. (This involves shifting left by 12 bits if the segment limit is page−granular, and not if it is byte−granular; so you end up with a byte limit in either case.) The segment limit obtained is then loaded into the destination (first) operand.

B.4.144 LTR: Load Task Register

```
LTR r/m16                       ; 0F 00 /3              [286,PRIV]
```

LTR looks up the segment base and limit in the GDT or LDT descriptor specified by the segment selector given as its operand, and loads them into the Task Register.

B.4.145 MASKMOVDQU: Byte Mask Write

```
MASKMOVDQU xmm1,xmm2            ; 66 0F F7 /r       [WILLAMETTE,SSE2]
```

MASKMOVDQU stores data from xmm1 to the location specified by ES:(E)DI. The size of the store depends on the address−size attribute. The most significant bit in each byte of the mask register xmm2 is used to selectively write the data (0 = no write, 1 = write) on a per−byte basis.

B.4.146 MASKMOVQ: Byte Mask Write

```
MASKMOVQ mm1,mm2                ; 0F F7 /r          [KATMAI,MMX]
```

MASKMOVQ stores data from mm1 to the location specified by ES:(E)DI. The size of the store depends on the address−size attribute. The most significant bit in each byte of the mask register mm2 is used to selectively write the data (0 = no write, 1 = write) on a per−byte basis.

B.4.147 MAXPD: Return Packed Double−Precision FP Maximum

```
MAXPD xmm1,xmm2/m128            ; 66 0F 5F /r       [WILLAMETTE,SSE2]
```

MAXPD performs a SIMD compare of the packed double−precision FP numbers from xmm1 and xmm2/mem, and stores the maximum values of each pair of values in xmm1. If the values being compared are both zeroes, source2 (xmm2/m128) would be returned. If source2 (xmm2/m128) is an SNaN, this SNaN is forwarded unchanged to the destination (i.e., a QNaN version of the SNaN is not returned).

B.4.148 MAXPS: Return Packed Single−Precision FP Maximum

```
MAXPS xmm1,xmm2/m128            ; 0F 5F /r          [KATMAI,SSE]
```

MAXPS performs a SIMD compare of the packed single−precision FP numbers from xmm1 and xmm2/mem, and stores the maximum values of each pair of values in xmm1. If the values being compared are both zeroes, source2 (xmm2/m128) would be returned. If source2 (xmm2/m128) is an SNaN, this SNaN is forwarded unchanged to the destination (i.e., a QNaN version of the SNaN is not returned).

B.4.149 MAXSD: Return Scalar Double−Precision FP Maximum

```
MAXSD xmm1,xmm2/m64             ; F2 0F 5F /r       [WILLAMETTE,SSE2]
```

MAXSD compares the low−order double−precision FP numbers from xmm1 and xmm2/mem, and stores the maximum value in xmm1. If the values being compared are both zeroes, source2 (xmm2/m64) would be returned. If source2 (xmm2/m64) is an SNaN, this SNaN is forwarded unchanged to the destination (i.e., a QNaN version of the SNaN is not returned). The high quadword of the destination is left unchanged.

B.4.150 MAXSS: Return Scalar Single−Precision FP Maximum

```
MAXSS xmm1,xmm2/m32             ; F3 0F 5F /r       [KATMAI,SSE]
```

MAXSS compares the low−order single−precision FP numbers from xmm1 and xmm2/mem, and stores the maximum value in xmm1. If the values being compared are both zeroes, source2 (xmm2/m32) would be returned. If source2 (xmm2/m32) is an SNaN, this SNaN is forwarded unchanged to the destination (i.e., a QNaN version of the SNaN is not returned). The high three doublewords of the destination are left unchanged.

B.4.151 MFENCE: Memory Fence

```
MFENCE                          ; 0F AE /6          [WILLAMETTE,SSE2]
```

MFENCE performs a serialising operation on all loads from memory and writes to memory that were issued before the MFENCE instruction. This guarantees that all memory reads and writes before the MFENCE instruction are completed before any reads and writes after the MFENCE instruction.

MFENCE is ordered respective to other MFENCE instructions, LFENCE, SFENCE, any memory read and any other serialising instruction (such as CPUID).

Weakly ordered memory types can be used to achieve higher processor performance through such techniques as out−of−order issue, speculative reads, write−combining, and write−collapsing. The degree to which a consumer of data recognizes or knows that the data is weakly ordered varies among applications and may be unknown to the producer of this data. The MFENCE instruction provides a performance−efficient way of ensuring load and store ordering between routines that produce weakly−ordered results and routines that consume that data.

MFENCE uses the following ModRM encoding:

```
Mod (7:6)         = 11B
Reg/Opcode (5:3)  = 110B
R/M (2:0)         = 000B
```

All other ModRM encodings are defined to be reserved, and use of these encodings risks incompatibility with future processors.

See also LFENCE (section B.4.137) and SFENCE (section B.4.288).

B.4.152 MINPD: Return Packed Double−Precision FP Minimum

```
MINPD xmm1,xmm2/m128            ; 66 0F 5D /r       [WILLAMETTE,SSE2]
```

MINPD performs a SIMD compare of the packed double−precision FP numbers from xmm1 and xmm2/mem, and stores the minimum values of each pair of values in xmm1. If the values being compared are both zeroes, source2 (xmm2/m128) would be returned. If source2 (xmm2/m128) is an SNaN, this SNaN is forwarded unchanged to the destination (i.e., a QNaN version of the SNaN is not returned).

B.4.153 MINPS: Return Packed Single−Precision FP Minimum

```
MINPS xmm1,xmm2/m128            ; 0F 5D /r          [KATMAI,SSE]
```

MINPS performs a SIMD compare of the packed single−precision FP numbers from xmm1 and xmm2/mem, and stores the minimum values of each pair of values in xmm1. If the values being compared are both zeroes, source2 (xmm2/m128) would be returned. If source2 (xmm2/m128) is an SNaN, this SNaN is forwarded unchanged to the destination (i.e., a QNaN version of the SNaN is not returned).

B.4.154 MINSD: Return Scalar Double−Precision FP Minimum

```
MINSD xmm1,xmm2/m64             ; F2 0F 5D /r       [WILLAMETTE,SSE2]
```

MINSD compares the low−order double−precision FP numbers from xmm1 and xmm2/mem, and stores the minimum value in xmm1. If the values being compared are both zeroes, source2 (xmm2/m64) would be returned. If source2 (xmm2/m64) is an SNaN, this SNaN is forwarded unchanged to the destination (i.e., a QNaN version of the SNaN is not returned). The high quadword of the destination is left unchanged.

B.4.155 MINSS: Return Scalar Single−Precision FP Minimum

```
MINSS xmm1,xmm2/m32             ; F3 0F 5D /r       [KATMAI,SSE]
```

MINSS compares the low−order single−precision FP numbers from xmm1 and xmm2/mem, and stores the minimum value in xmm1. If the values being compared are both zeroes, source2 (xmm2/m32) would be returned. If source2 (xmm2/m32) is an SNaN, this SNaN is forwarded unchanged to the destination (i.e., a QNaN version of the SNaN is not returned). The high three doublewords of the destination are left unchanged.

B.4.156 MOV: Move Data

```
MOV r/m8,reg8                   ; 88 /r                 [8086]
MOV r/m16,reg16                 ; o16 89 /r             [8086]
MOV r/m32,reg32                 ; o32 89 /r             [386]
MOV reg8,r/m8                   ; 8A /r                 [8086]
MOV reg16,r/m16                 ; o16 8B /r             [8086]
MOV reg32,r/m32                 ; o32 8B /r             [386]

MOV reg8,imm8                   ; B0+r ib               [8086]
MOV reg16,imm16                 ; o16 B8+r iw           [8086]
MOV reg32,imm32                 ; o32 B8+r id           [386]
MOV r/m8,imm8                   ; C6 /0 ib              [8086]
MOV r/m16,imm16                 ; o16 C7 /0 iw          [8086]
MOV r/m32,imm32                 ; o32 C7 /0 id          [386]

MOV AL,memoffs8                 ; A0 ow/od              [8086]
MOV AX,memoffs16                ; o16 A1 ow/od          [8086]
MOV EAX,memoffs32               ; o32 A1 ow/od          [386]
MOV memoffs8,AL                 ; A2 ow/od              [8086]
MOV memoffs16,AX                ; o16 A3 ow/od          [8086]
MOV memoffs32,EAX               ; o32 A3 ow/od          [386]

MOV r/m16,segreg                ; o16 8C /r             [8086]
MOV r/m32,segreg                ; o32 8C /r             [386]
MOV segreg,r/m16                ; o16 8E /r             [8086]
MOV segreg,r/m32                ; o32 8E /r             [386]

MOV reg32,CR0/2/3/4             ; 0F 20 /r              [386]
MOV reg32,DR0/1/2/3/6/7         ; 0F 21 /r              [386]
MOV reg32,TR3/4/5/6/7           ; 0F 24 /r              [386]
MOV CR0/2/3/4,reg32             ; 0F 22 /r              [386]
MOV DR0/1/2/3/6/7,reg32         ; 0F 23 /r              [386]
MOV TR3/4/5/6/7,reg32           ; 0F 26 /r              [386]
```

MOV copies the contents of its source (second) operand into its destination (first) operand.

In all forms of the MOV instruction, the two operands are the same size, except for moving between a segment register and an r/m32 operand. These instructions are treated exactly like the corresponding 16−bit equivalent (so that, for example, MOV DS,EAX functions identically to MOV DS,AX but saves a prefix when in 32−bit mode), except that when a segment register is moved into a 32−bit destination, the top two bytes of the result are undefined.

MOV may not use CS as a destination.

CR4 is only a supported register on the Pentium and above.

Test registers are supported on 386/486 processors and on some non−Intel Pentium class processors.

B.4.157 MOVAPD: Move Aligned Packed Double−Precision FP Values

```
MOVAPD xmm1,xmm2/mem128         ; 66 0F 28 /r     [WILLAMETTE,SSE2]
MOVAPD xmm1/mem128,xmm2         ; 66 0F 29 /r     [WILLAMETTE,SSE2]
```

MOVAPD moves a double quadword containing 2 packed double−precision FP values from the source operand to the destination. When the source or destination operand is a memory location, it must be aligned on a 16−byte boundary.

To move data in and out of memory locations that are not known to be on 16−byte boundaries, use the MOVUPD instruction (section B.4.182).

B.4.158 MOVAPS: Move Aligned Packed Single−Precision FP Values

```
MOVAPS xmm1,xmm2/mem128        ; 0F 28 /r         [KATMAI,SSE]
MOVAPS xmm1/mem128,xmm2        ; 0F 29 /r         [KATMAI,SSE]
```

MOVAPS moves a double quadword containing 4 packed single−precision FP values from the source operand to the destination. When the source or destination operand is a memory location, it must be aligned on a 16−byte boundary.

To move data in and out of memory locations that are not known to be on 16−byte boundaries, use the MOVUPS instruction (section B.4.183).

B.4.159 MOVD: Move Doubleword to/from MMX Register

```
MOVD mm,r/m32                  ; 0F 6E /r             [PENT,MMX]
MOVD r/m32,mm                  ; 0F 7E /r             [PENT,MMX]
MOVD xmm,r/m32                 ; 66 0F 6E /r      [WILLAMETTE,SSE2]
MOVD r/m32,xmm                 ; 66 0F 7E /r      [WILLAMETTE,SSE2]
```

MOVD copies 32 bits from its source (second) operand into its destination (first) operand. When the destination is a 64−bit MMX register or a 128−bit XMM register, the input value is zero−extended to fill the destination register.

B.4.160 MOVDQ2Q: Move Quadword from XMM to MMX register.

```
MOVDQ2Q mm,xmm                 ; F2 OF D6 /r      [WILLAMETTE,SSE2]
```

MOVDQ2Q moves the low quadword from the source operand to the destination operand.

B.4.161 MOVDQA: Move Aligned Double Quadword

```
MOVDQA xmm1,xmm2/m128          ; 66 OF 6F /r      [WILLAMETTE,SSE2]
MOVDQA xmm1/m128,xmm2          ; 66 OF 7F /r      [WILLAMETTE,SSE2]
```

MOVDQA moves a double quadword from the source operand to the destination operand. When the source or destination operand is a memory location, it must be aligned to a 16−byte boundary.

To move a double quadword to or from unaligned memory locations, use the MOVDQU instruction (section B.4.162).

B.4.162 MOVDQU: Move Unaligned Double Quadword

```
MOVDQU xmm1,xmm2/m128          ; F3 OF 6F /r      [WILLAMETTE,SSE2]
MOVDQU xmm1/m128,xmm2          ; F3 OF 7F /r      [WILLAMETTE,SSE2]
```

MOVDQU moves a double quadword from the source operand to the destination operand. When the source or destination operand is a memory location, the memory may be unaligned.

To move a double quadword to or from known aligned memory locations, use the MOVDQA instruction (section B.4.161).

B.4.163 MOVHLPS: Move Packed Single−Precision FP High to Low

```
MOVHLPS xmm1,xmm2              ; OF 12 /r         [KATMAI,SSE]
```

MOVHLPS moves the two packed single–precision FP values from the high quadword of the source register xmm2 to the low quadword of the destination register, xmm2. The upper quadword of xmm1 is left unchanged.

The operation of this instruction is:

```
dst[0–63]   := src[64–127],
dst[64–127] remains unchanged.
```

B.4.164 MOVHPD: Move High Packed Double–Precision FP

```
MOVHPD xmm,m64              ; 66 OF 16 /r      [WILLAMETTE,SSE2]
MOVHPD m64,xmm              ; 66 OF 17 /r      [WILLAMETTE,SSE2]
```

MOVHPD moves a double–precision FP value between the source and destination operands. One of the operands is a 64–bit memory location, the other is the high quadword of an XMM register.

The operation of this instruction is:

```
mem[0–63]   := xmm[64–127];
```

or

```
xmm[0–63]   remains unchanged;
xmm[64–127] := mem[0–63].
```

B.4.165 MOVHPS: Move High Packed Single–Precision FP

```
MOVHPS xmm,m64              ; 0F 16 /r         [KATMAI,SSE]
MOVHPS m64,xmm              ; 0F 17 /r         [KATMAI,SSE]
```

MOVHPS moves two packed single–precision FP values between the source and destination operands. One of the operands is a 64–bit memory location, the other is the high quadword of an XMM register.

The operation of this instruction is:

```
mem[0–63]   := xmm[64–127];
```

or

```
xmm[0–63]   remains unchanged;
xmm[64–127] := mem[0–63].
```

B.4.166 MOVLHPS: Move Packed Single–Precision FP Low to High

```
MOVLHPS xmm1,xmm2           ; OF 16 /r         [KATMAI,SSE]
```

MOVLHPS moves the two packed single–precision FP values from the low quadword of the source register xmm2 to the high quadword of the destination register, xmm2. The low quadword of xmm1 is left unchanged.

The operation of this instruction is:

```
dst[0–63]   remains unchanged;
dst[64–127] := src[0–63].
```

B.4.167 MOVLPD: Move Low Packed Double–Precision FP

```
MOVLPD xmm,m64              ; 66 OF 12 /r      [WILLAMETTE,SSE2]
MOVLPD m64,xmm              ; 66 OF 13 /r      [WILLAMETTE,SSE2]
```

MOVLPD moves a double–precision FP value between the source and destination operands. One of the operands is a 64–bit memory location, the other is the low quadword of an XMM register.

The operation of this instruction is:

```
mem(0-63)    := xmm(0-63);
```

or

```
xmm(0-63)    := mem(0-63);
xmm(64-127) remains unchanged.
```

B.4.168 MOVLPS: Move Low Packed Single–Precision FP

```
MOVLPS xmm,m64                  ; OF 12 /r          [KATMAI,SSE]
MOVLPS m64,xmm                  ; OF 13 /r          [KATMAI,SSE]
```

MOVLPS moves two packed single–precision FP values between the source and destination operands. One of the operands is a 64–bit memory location, the other is the low quadword of an XMM register.

The operation of this instruction is:

```
mem(0-63)    := xmm(0-63);
```

or

```
xmm(0-63)    := mem(0-63);
xmm(64-127) remains unchanged.
```

B.4.169 MOVMSKPD: Extract Packed Double–Precision FP Sign Mask

```
MOVMSKPD reg32,xmm                ; 66 0F 50 /r    [WILLAMETTE,SSE2]
```

MOVMSKPD inserts a 2–bit mask in r32, formed of the most significant bits of each double–precision FP number of the source operand.

B.4.170 MOVMSKPS: Extract Packed Single–Precision FP Sign Mask

```
MOVMSKPS reg32,xmm                ; 0F 50 /r       [KATMAI,SSE]
```

MOVMSKPS inserts a 4–bit mask in r32, formed of the most significant bits of each single–precision FP number of the source operand.

B.4.171 MOVNTDQ: Move Double Quadword Non Temporal

```
MOVNTDQ m128,xmm                ; 66 0F E7 /r       [WILLAMETTE,SSE2]
```

MOVNTDQ moves the double quadword from the XMM source register to the destination memory location, using a non–temporal hint. This store instruction minimizes cache pollution.

B.4.172 MOVNTI: Move Doubleword Non Temporal

```
MOVNTI m32,reg32                ; 0F C3 /r          [WILLAMETTE,SSE2]
```

MOVNTI moves the doubleword in the source register to the destination memory location, using a non–temporal hint. This store instruction minimizes cache pollution.

B.4.173 MOVNTPD: Move Aligned Four Packed Single–Precision FP Values Non Temporal

```
MOVNTPD m128,xmm                ; 66 0F 2B /r       [WILLAMETTE,SSE2]
```

MOVNTPD moves the double quadword from the XMM source register to the destination memory location, using a non–temporal hint. This store instruction minimizes cache pollution. The memory location must be aligned to a 16–byte boundary.

B.4.174 MOVNTPS: Move Aligned Four Packed Single–Precision FP Values Non Temporal

```
MOVNTPS m128,xmm                ; 0F 2B /r          [KATMAI,SSE]
```

MOVNTPS moves the double quadword from the XMM source register to the destination memory location, using a non–temporal hint. This store instruction minimizes cache pollution. The memory location must be aligned to a 16–byte boundary.

B.4.175 MOVNTQ: Move Quadword Non Temporal

```
MOVNTQ m64,mm                   ; 0F E7 /r          [KATMAI,MMX]
```

MOVNTQ moves the quadword in the MMX source register to the destination memory location, using a non–temporal hint. This store instruction minimizes cache pollution.

B.4.176 MOVQ: Move Quadword to/from MMX Register

```
MOVQ mm1,mm2/m64                ; 0F 6F /r              [PENT,MMX]
MOVQ mm1/m64,mm2                ; 0F 7F /r              [PENT,MMX]

MOVQ xmm1,xmm2/m64              ; F3 0F 7E /r     [WILLAMETTE,SSE2]
MOVQ xmm1/m64,xmm2              ; 66 0F D6 /r     [WILLAMETTE,SSE2]
```

MOVQ copies 64 bits from its source (second) operand into its destination (first) operand. When the source is an XMM register, the low quadword is moved. When the destination is an XMM register, the destination is the low quadword, and the high quadword is cleared.

B.4.177 MOVQ2DQ: Move Quadword from MMX to XMM register.

```
MOVQ2DQ xmm,mm                  ; F3 OF D6 /r     [WILLAMETTE,SSE2]
```

MOVQ2DQ moves the quadword from the source operand to the low quadword of the destination operand, and clears the high quadword.

B.4.178 MOVSB, MOVSW, MOVSD: Move String

```
MOVSB                           ; A4                    [8086]
MOVSW                           ; o16 A5                [8086]
MOVSD                           ; o32 A5                [386]
```

MOVSB copies the byte at [DS:SI] or [DS:ESI] to [ES:DI] or [ES:EDI]. It then increments or decrements (depending on the direction flag: increments if the flag is clear, decrements if it is set) SI and DI (or ESI and EDI).

The registers used are SI and DI if the address size is 16 bits, and ESI and EDI if it is 32 bits. If you need to use an address size not equal to the current BITS setting, you can use an explicit a16 or a32 prefix.

The segment register used to load from [SI] or [ESI] can be overridden by using a segment register name as a prefix (for example, es movsb). The use of ES for the store to [DI] or [EDI] cannot be overridden.

MOVSW and MOVSD work in the same way, but they copy a word or a doubleword instead of a byte, and increment or decrement the addressing registers by 2 or 4 instead of 1.

The REP prefix may be used to repeat the instruction CX (or ECX – again, the address size chooses which) times.

B.4.179 MOVSD: Move Scalar Double–Precision FP Value

```
MOVSD xmm1,xmm2/m64             ; F2 0F 10 /r     [WILLAMETTE,SSE2]
MOVSD xmm1/m64,xmm2             ; F2 0F 11 /r     [WILLAMETTE,SSE2]
```

MOVSD moves a double–precision FP value from the source operand to the destination operand. When the source or destination is a register, the low–order FP value is read or written.

B.4.180 MOVSS: Move Scalar Single–Precision FP Value

```
MOVSS xmm1,xmm2/m32             ; F3 0F 10 /r      [KATMAI,SSE]
MOVSS xmm1/m32,xmm2             ; F3 0F 11 /r      [KATMAI,SSE]
```

MOVSS moves a single–precision FP value from the source operand to the destination operand. When the source or destination is a register, the low–order FP value is read or written.

B.4.181 MOVSX, MOVZX: Move Data with Sign or Zero Extend

```
MOVSX reg16,r/m8                ; o16 0F BE /r          [386]
MOVSX reg32,r/m8                ; o32 0F BE /r          [386]
MOVSX reg32,r/m16               ; o32 0F BF /r          [386]

MOVZX reg16,r/m8                ; o16 0F B6 /r          [386]
MOVZX reg32,r/m8                ; o32 0F B6 /r          [386]
MOVZX reg32,r/m16               ; o32 0F B7 /r          [386]
```

MOVSX sign–extends its source (second) operand to the length of its destination (first) operand, and copies the result into the destination operand. MOVZX does the same, but zero–extends rather than sign–extending.

B.4.182 MOVUPD: Move Unaligned Packed Double–Precision FP Values

```
MOVUPD xmm1,xmm2/mem128         ; 66 0F 10 /r      [WILLAMETTE,SSE2]
MOVUPD xmm1/mem128,xmm2         ; 66 0F 11 /r      [WILLAMETTE,SSE2]
```

MOVUPD moves a double quadword containing 2 packed double–precision FP values from the source operand to the destination. This instruction makes no assumptions about alignment of memory operands.

To move data in and out of memory locations that are known to be on 16–byte boundaries, use the MOVAPD instruction (section B.4.157).

B.4.183 MOVUPS: Move Unaligned Packed Single–Precision FP Values

```
MOVUPS xmm1,xmm2/mem128         ; 0F 10 /r         [KATMAI,SSE]
MOVUPS xmm1/mem128,xmm2         ; 0F 11 /r         [KATMAI,SSE]
```

MOVUPS moves a double quadword containing 4 packed single–precision FP values from the source operand to the destination. This instruction makes no assumptions about alignment of memory operands.

To move data in and out of memory locations that are known to be on 16–byte boundaries, use the MOVAPS instruction (section B.4.158).

B.4.184 MUL: Unsigned Integer Multiply

```
MUL r/m8                        ; F6 /4                 [8086]
MUL r/m16                       ; o16 F7 /4             [8086]
MUL r/m32                       ; o32 F7 /4             [386]
```

MUL performs unsigned integer multiplication. The other operand to the multiplication, and the destination operand, are implicit, in the following way:

- For MUL r/m8, AL is multiplied by the given operand; the product is stored in AX.
- For MUL r/m16, AX is multiplied by the given operand; the product is stored in DX:AX.

- For `MUL r/m32`, `EAX` is multiplied by the given operand; the product is stored in `EDX:EAX`.

Signed integer multiplication is performed by the `IMUL` instruction: see section B.4.118.

B.4.185 `MULPD`: Packed Single–FP Multiply

```
MULPD xmm1,xmm2/mem128          ; 66 0F 59 /r      [WILLAMETTE,SSE2]
```

`MULPD` performs a SIMD multiply of the packed double–precision FP values in both operands, and stores the results in the destination register.

B.4.186 `MULPS`: Packed Single–FP Multiply

```
MULPS xmm1,xmm2/mem128          ; 0F 59 /r         [KATMAI,SSE]
```

`MULPS` performs a SIMD multiply of the packed single–precision FP values in both operands, and stores the results in the destination register.

B.4.187 `MULSD`: Scalar Single–FP Multiply

```
MULSD xmm1,xmm2/mem32           ; F2 0F 59 /r      [WILLAMETTE,SSE2]
```

`MULSD` multiplies the lowest double–precision FP values of both operands, and stores the result in the low quadword of xmm1.

B.4.188 `MULSS`: Scalar Single–FP Multiply

```
MULSS xmm1,xmm2/mem32           ; F3 0F 59 /r      [KATMAI,SSE]
```

`MULSS` multiplies the lowest single–precision FP values of both operands, and stores the result in the low doubleword of xmm1.

B.4.189 `NEG`, `NOT`: Two's and One's Complement

```
NEG r/m8                        ; F6 /3                [8086]
NEG r/m16                       ; o16 F7 /3            [8086]
NEG r/m32                       ; o32 F7 /3            [386]

NOT r/m8                        ; F6 /2                [8086]
NOT r/m16                       ; o16 F7 /2            [8086]
NOT r/m32                       ; o32 F7 /2            [386]
```

`NEG` replaces the contents of its operand by the two's complement negation (invert all the bits and then add one) of the original value. `NOT`, similarly, performs one's complement (inverts all the bits).

B.4.190 `NOP`: No Operation

```
NOP                             ; 90                   [8086]
```

`NOP` performs no operation. Its opcode is the same as that generated by `XCHG AX,AX` or `XCHG EAX,EAX` (depending on the processor mode; see section B.4.333).

B.4.191 `OR`: Bitwise OR

```
OR r/m8,reg8                    ; 08 /r                [8086]
OR r/m16,reg16                  ; o16 09 /r            [8086]
OR r/m32,reg32                  ; o32 09 /r            [386]

OR reg8,r/m8                    ; 0A /r                [8086]
OR reg16,r/m16                  ; o16 0B /r            [8086]
OR reg32,r/m32                  ; o32 0B /r            [386]
```

```
OR r/m8,imm8                    ; 80 /1 ib              [8086]
OR r/m16,imm16                  ; o16 81 /1 iw          [8086]
OR r/m32,imm32                  ; o32 81 /1 id          [386]

OR r/m16,imm8                   ; o16 83 /1 ib          [8086]
OR r/m32,imm8                   ; o32 83 /1 ib          [386]

OR AL,imm8                      ; 0C ib                 [8086]
OR AX,imm16                     ; o16 0D iw             [8086]
OR EAX,imm32                    ; o32 0D id             [386]
```

OR performs a bitwise OR operation between its two operands (i.e. each bit of the result is 1 if and only if at least one of the corresponding bits of the two inputs was 1), and stores the result in the destination (first) operand.

In the forms with an 8–bit immediate second operand and a longer first operand, the second operand is considered to be signed, and is sign–extended to the length of the first operand. In these cases, the BYTE qualifier is necessary to force NASM to generate this form of the instruction.

The MMX instruction POR (see section B.4.247) performs the same operation on the 64–bit MMX registers.

B.4.192 ORPD: Bit–wise Logical OR of Double–Precision FP Data

```
ORPD xmm1,xmm2/m128             ; 66 0F 56 /r     [WILLAMETTE,SSE2]
```

ORPD return a bit–wise logical OR between xmm1 and xmm2/mem, and stores the result in xmm1. If the source operand is a memory location, it must be aligned to a 16–byte boundary.

B.4.193 ORPS: Bit–wise Logical OR of Single–Precision FP Data

```
ORPS xmm1,xmm2/m128             ; 0F 56 /r        [KATMAI,SSE]
```

ORPS return a bit–wise logical OR between xmm1 and xmm2/mem, and stores the result in xmm1. If the source operand is a memory location, it must be aligned to a 16–byte boundary.

B.4.194 OUT: Output Data to I/O Port

```
OUT imm8,AL                     ; E6 ib                 [8086]
OUT imm8,AX                     ; o16 E7 ib             [8086]
OUT imm8,EAX                    ; o32 E7 ib             [386]
OUT DX,AL                       ; EE                    [8086]
OUT DX,AX                       ; o16 EF                [8086]
OUT DX,EAX                      ; o32 EF                [386]
```

OUT writes the contents of the given source register to the specified I/O port. The port number may be specified as an immediate value if it is between 0 and 255, and otherwise must be stored in DX. See also IN (section B.4.119).

B.4.195 OUTSB, OUTSW, OUTSD: Output String to I/O Port

```
OUTSB                           ; 6E                    [186]
OUTSW                           ; o16 6F                [186]
OUTSD                           ; o32 6F                [386]
```

OUTSB loads a byte from [DS:SI] or [DS:ESI] and writes it to the I/O port specified in DX. It then increments or decrements (depending on the direction flag: increments if the flag is clear, decrements if it is set) SI or ESI.

The register used is SI if the address size is 16 bits, and ESI if it is 32 bits. If you need to use an address size not equal to the current BITS setting, you can use an explicit a16 or a32 prefix.

The segment register used to load from [SI] or [ESI] can be overridden by using a segment register name as a prefix (for example, es outsb).

OUTSW and OUTSD work in the same way, but they output a word or a doubleword instead of a byte, and increment or decrement the addressing registers by 2 or 4 instead of 1.

The REP prefix may be used to repeat the instruction CX (or ECX – again, the address size chooses which) times.

B.4.196 PACKSSDW, PACKSSWB, PACKUSWB: Pack Data

```
PACKSSDW mm1,mm2/m64          ; 0F 6B /r              [PENT,MMX]
PACKSSWB mm1,mm2/m64          ; 0F 63 /r              [PENT,MMX]
PACKUSWB mm1,mm2/m64          ; 0F 67 /r              [PENT,MMX]

PACKSSDW xmm1,xmm2/m128       ; 66 0F 6B /r     [WILLAMETTE,SSE2]
PACKSSWB xmm1,xmm2/m128       ; 66 0F 63 /r     [WILLAMETTE,SSE2]
PACKUSWB xmm1,xmm2/m128       ; 66 0F 67 /r     [WILLAMETTE,SSE2]
```

All these instructions start by combining the source and destination operands, and then splitting the result in smaller sections which it then packs into the destination register. The MMX versions pack two 64–bit operands into one 64–bit register, while the SSE versions pack two 128–bit operands into one 128–bit register.

- PACKSSWB splits the combined value into words, and then reduces the words to bytes, using signed saturation. It then packs the bytes into the destination register in the same order the words were in.
- PACKSSDW performs the same operation as PACKSSWB, except that it reduces doublewords to words, then packs them into the destination register.
- PACKUSWB performs the same operation as PACKSSWB, except that it uses unsigned saturation when reducing the size of the elements.

To perform signed saturation on a number, it is replaced by the largest signed number (7FFFh or 7Fh) that *will* fit, and if it is too small it is replaced by the smallest signed number (8000h or 80h) that will fit. To perform unsigned saturation, the input is treated as unsigned, and the input is replaced by the largest unsigned number that will fit.

B.4.197 PADDB, PADDW, PADDD: Add Packed Integers

```
PADDB mm1,mm2/m64             ; 0F FC /r              [PENT,MMX]
PADDW mm1,mm2/m64             ; 0F FD /r              [PENT,MMX]
PADDD mm1,mm2/m64             ; 0F FE /r              [PENT,MMX]

PADDB xmm1,xmm2/m128          ; 66 0F FC /r     [WILLAMETTE,SSE2]
PADDW xmm1,xmm2/m128          ; 66 0F FD /r     [WILLAMETTE,SSE2]
PADDD xmm1,xmm2/m128          ; 66 0F FE /r     [WILLAMETTE,SSE2]
```

PADDx performs packed addition of the two operands, storing the result in the destination (first) operand.

- PADDB treats the operands as packed bytes, and adds each byte individually;
- PADDW treats the operands as packed words;
- PADDD treats its operands as packed doublewords.

When an individual result is too large to fit in its destination, it is wrapped around and the low bits are stored, with the carry bit discarded.

B.4.198 PADDQ: Add Packed Quadword Integers

```
PADDQ mm1,mm2/m64              ; 0F D4 /r              [PENT,MMX]

PADDQ xmm1,xmm2/m128           ; 66 0F D4 /r      [WILLAMETTE,SSE2]
```

PADDQ adds the quadwords in the source and destination operands, and stores the result in the destination register.

When an individual result is too large to fit in its destination, it is wrapped around and the low bits are stored, with the carry bit discarded.

B.4.199 PADDSB, PADDSW: Add Packed Signed Integers With Saturation

```
PADDSB mm1,mm2/m64             ; 0F EC /r              [PENT,MMX]
PADDSW mm1,mm2/m64             ; 0F ED /r              [PENT,MMX]

PADDSB xmm1,xmm2/m128          ; 66 0F EC /r      [WILLAMETTE,SSE2]
PADDSW xmm1,xmm2/m128          ; 66 0F ED /r      [WILLAMETTE,SSE2]
```

PADDSx performs packed addition of the two operands, storing the result in the destination (first) operand. PADDSB treats the operands as packed bytes, and adds each byte individually; and PADDSW treats the operands as packed words.

When an individual result is too large to fit in its destination, a saturated value is stored. The resulting value is the value with the largest magnitude of the same sign as the result which will fit in the available space.

B.4.200 PADDSIW: MMX Packed Addition to Implicit Destination

```
PADDSIW mmxreg,r/m64           ; 0F 51 /r              [CYRIX,MMX]
```

PADDSIW, specific to the Cyrix extensions to the MMX instruction set, performs the same function as PADDSW, except that the result is placed in an implied register.

To work out the implied register, invert the lowest bit in the register number. So PADDSIW MM0,MM2 would put the result in MM1, but PADDSIW MM1,MM2 would put the result in MM0.

B.4.201 PADDUSB, PADDUSW: Add Packed Unsigned Integers With Saturation

```
PADDUSB mm1,mm2/m64            ; 0F DC /r              [PENT,MMX]
PADDUSW mm1,mm2/m64            ; 0F DD /r              [PENT,MMX]

PADDUSB xmm1,xmm2/m128          ; 66 0F DC /r     [WILLAMETTE,SSE2]
PADDUSW xmm1,xmm2/m128          ; 66 0F DD /r     [WILLAMETTE,SSE2]
```

PADDUSx performs packed addition of the two operands, storing the result in the destination (first) operand. PADDUSB treats the operands as packed bytes, and adds each byte individually; and PADDUSW treats the operands as packed words.

When an individual result is too large to fit in its destination, a saturated value is stored. The resulting value is the maximum value that will fit in the available space.

B.4.202 PAND, PANDN: MMX Bitwise AND and AND–NOT

```
PAND mm1,mm2/m64               ; 0F DB /r              [PENT,MMX]
PANDN mm1,mm2/m64              ; 0F DF /r              [PENT,MMX]

PAND xmm1,xmm2/m128            ; 66 0F DB /r      [WILLAMETTE,SSE2]
PANDN xmm1,xmm2/m128           ; 66 0F DF /r      [WILLAMETTE,SSE2]
```

PAND performs a bitwise AND operation between its two operands (i.e. each bit of the result is 1 if and only if the corresponding bits of the two inputs were both 1), and stores the result in the destination (first) operand.

PANDN performs the same operation, but performs a one's complement operation on the destination (first) operand first.

B.4.203 PAUSE: Spin Loop Hint

```
PAUSE                           ; F3 90             [WILLAMETTE,SSE2]
```

PAUSE provides a hint to the processor that the following code is a spin loop. This improves processor performance by bypassing possible memory order violations. On older processors, this instruction operates as a NOP.

B.4.204 PAVEB: MMX Packed Average

```
PAVEB mmxreg,r/m64              ; 0F 50 /r              [CYRIX,MMX]
```

PAVEB, specific to the Cyrix MMX extensions, treats its two operands as vectors of eight unsigned bytes, and calculates the average of the corresponding bytes in the operands. The resulting vector of eight averages is stored in the first operand.

This opcode maps to MOVMSKPS r32, xmm on processors that support the SSE instruction set.

B.4.205 PAVGB PAVGW: Average Packed Integers

```
PAVGB mm1,mm2/m64               ; 0F E0 /r          [KATMAI,MMX]
PAVGW mm1,mm2/m64               ; 0F E3 /r          [KATMAI,MMX,SM]

PAVGB xmm1,xmm2/m128            ; 66 0F E0 /r       [WILLAMETTE,SSE2]
PAVGW xmm1,xmm2/m128            ; 66 0F E3 /r       [WILLAMETTE,SSE2]
```

PAVGB and PAVGW add the unsigned data elements of the source operand to the unsigned data elements of the destination register, then adds 1 to the temporary results. The results of the add are then each independently right−shifted by one bit position. The high order bits of each element are filled with the carry bits of the corresponding sum.

- PAVGB operates on packed unsigned bytes, and
- PAVGW operates on packed unsigned words.

B.4.206 PAVGUSB: Average of unsigned packed 8−bit values

```
PAVGUSB mm1,mm2/m64             ; 0F 0F /r BF           [PENT,3DNOW]
```

PAVGUSB adds the unsigned data elements of the source operand to the unsigned data elements of the destination register, then adds 1 to the temporary results. The results of the add are then each independently right−shifted by one bit position. The high order bits of each element are filled with the carry bits of the corresponding sum.

This instruction performs exactly the same operations as the PAVGB MMX instruction (section B.4.205).

B.4.207 PCMPxx: Compare Packed Integers.

```
PCMPEQB mm1,mm2/m64             ; 0F 74 /r              [PENT,MMX]
PCMPEQW mm1,mm2/m64             ; 0F 75 /r              [PENT,MMX]
PCMPEQD mm1,mm2/m64             ; 0F 76 /r              [PENT,MMX]
```

```
PCMPGTB mm1,mm2/m64             ; 0F 64 /r              [PENT,MMX]
PCMPGTW mm1,mm2/m64             ; 0F 65 /r              [PENT,MMX]
PCMPGTD mm1,mm2/m64             ; 0F 66 /r              [PENT,MMX]

PCMPEQB xmm1,xmm2/m128          ; 66 0F 74 /r     [WILLAMETTE,SSE2]
PCMPEQW xmm1,xmm2/m128          ; 66 0F 75 /r     [WILLAMETTE,SSE2]
PCMPEQD xmm1,xmm2/m128          ; 66 0F 76 /r     [WILLAMETTE,SSE2]

PCMPGTB xmm1,xmm2/m128          ; 66 0F 64 /r     [WILLAMETTE,SSE2]
PCMPGTW xmm1,xmm2/m128          ; 66 0F 65 /r     [WILLAMETTE,SSE2]
PCMPGTD xmm1,xmm2/m128          ; 66 0F 66 /r     [WILLAMETTE,SSE2]
```

The `PCMPxx` instructions all treat their operands as vectors of bytes, words, or doublewords; corresponding elements of the source and destination are compared, and the corresponding element of the destination (first) operand is set to all zeros or all ones depending on the result of the comparison.

- `PCMPxxB` treats the operands as vectors of bytes;
- `PCMPxxW` treats the operands as vectors of words;
- `PCMPxxD` treats the operands as vectors of doublewords;
- `PCMPEQx` sets the corresponding element of the destination operand to all ones if the two elements compared are equal;
- `PCMPGTx` sets the destination element to all ones if the element of the first (destination) operand is greater (treated as a signed integer) than that of the second (source) operand.

B.4.208 `PDISTIB`: MMX Packed Distance and Accumulate with Implied Register

```
PDISTIB mm,m64                  ; 0F 54 /r              [CYRIX,MMX]
```

`PDISTIB`, specific to the Cyrix MMX extensions, treats its two input operands as vectors of eight unsigned bytes. For each byte position, it finds the absolute difference between the bytes in that position in the two input operands, and adds that value to the byte in the same position in the implied output register. The addition is saturated to an unsigned byte in the same way as `PADDUSB`.

To work out the implied register, invert the lowest bit in the register number. So `PDISTIB MM0,M64` would put the result in `MM1`, but `PDISTIB MM1,M64` would put the result in `MM0`.

Note that `PDISTIB` cannot take a register as its second source operand.

Operation:

```
dstI[0-7]     := dstI[0-7]   + ABS(src0[0-7] - src1[0-7]),
dstI[8-15]    := dstI[8-15]  + ABS(src0[8-15] - src1[8-15]),
.......
.......
dstI[56-63]   := dstI[56-63] + ABS(src0[56-63] - src1[56-63]).
```

B.4.209 `PEXTRW`: Extract Word

```
PEXTRW reg32,mm,imm8            ; 0F C5 /r ib     [KATMAI,MMX]
PEXTRW reg32,xmm,imm8           ; 66 0F C5 /r ib  [WILLAMETTE,SSE2]
```

`PEXTRW` moves the word in the source register (second operand) that is pointed to by the count operand (third operand), into the lower half of a 32–bit general purpose register. The upper half of the register is cleared to all 0s.

When the source operand is an MMX register, the two least significant bits of the count specify the source word. When it is an SSE register, the three least significant bits specify the word location.

B.4.210 PF2ID: Packed Single–Precision FP to Integer Convert

```
PF2ID mm1,mm2/m64              ; 0F 0F /r 1D          [PENT,3DNOW]
```

PF2ID converts two single–precision FP values in the source operand to signed 32–bit integers, using truncation, and stores them in the destination operand. Source values that are outside the range supported by the destination are saturated to the largest absolute value of the same sign.

B.4.211 PF2IW: Packed Single–Precision FP to Integer Word Convert

```
PF2IW mm1,mm2/m64              ; 0F 0F /r 1C          [PENT,3DNOW]
```

PF2IW converts two single–precision FP values in the source operand to signed 16–bit integers, using truncation, and stores them in the destination operand. Source values that are outside the range supported by the destination are saturated to the largest absolute value of the same sign.

- In the K6–2 and K6–III, the 16–bit value is zero–extended to 32–bits before storing.
- In the K6–2+, K6–III+ and Athlon processors, the value is sign–extended to 32–bits before storing.

B.4.212 PFACC: Packed Single–Precision FP Accumulate

```
PFACC mm1,mm2/m64              ; 0F 0F /r AE          [PENT,3DNOW]
```

PFACC adds the two single–precision FP values from the destination operand together, then adds the two single–precision FP values from the source operand, and places the results in the low and high doublewords of the destination operand.

The operation is:

```
dst[0–31]   := dst[0–31] + dst[32–63],
dst[32–63]  := src[0–31] + src[32–63].
```

B.4.213 PFADD: Packed Single–Precision FP Addition

```
PFADD mm1,mm2/m64              ; 0F 0F /r 9E          [PENT,3DNOW]
```

PFADD performs addition on each of two packed single–precision FP value pairs.

```
dst[0–31]   := dst[0–31]  + src[0–31],
dst[32–63]  := dst[32–63] + src[32–63].
```

B.4.214 PFCMPxx: Packed Single–Precision FP Compare

```
PFCMPEQ mm1,mm2/m64            ; 0F 0F /r B0          [PENT,3DNOW]
PFCMPGE mm1,mm2/m64            ; 0F 0F /r 90          [PENT,3DNOW]
PFCMPGT mm1,mm2/m64            ; 0F 0F /r A0          [PENT,3DNOW]
```

The PFCMPxx instructions compare the packed single–point FP values in the source and destination operands, and set the destination according to the result. If the condition is true, the destination is set to all 1s, otherwise it's set to all 0s.

- PFCMPEQ tests whether dst == src;
- PFCMPGE tests whether dst >= src;
- PFCMPGT tests whether dst > src.

B.4.215 PFMAX: Packed Single–Precision FP Maximum

```
PFMAX mm1,mm2/m64              ; 0F 0F /r A4          [PENT,3DNOW]
```

PFMAX returns the higher of each pair of single–precision FP values. If the higher value is zero, it is returned as positive zero.

B.4.216 PFMIN: Packed Single–Precision FP Minimum

```
PFMIN mm1,mm2/m64              ; 0F 0F /r 94          [PENT,3DNOW]
```

PFMIN returns the lower of each pair of single–precision FP values. If the lower value is zero, it is returned as positive zero.

B.4.217 PFMUL: Packed Single–Precision FP Multiply

```
PFMUL mm1,mm2/m64              ; 0F 0F /r B4          [PENT,3DNOW]
```

PFMUL returns the product of each pair of single–precision FP values.

```
dst[0–31]  := dst[0–31]  * src[0–31],
dst[32–63] := dst[32–63] * src[32–63].
```

B.4.218 PFNACC: Packed Single–Precision FP Negative Accumulate

```
PFNACC mm1,mm2/m64             ; 0F 0F /r 8A          [PENT,3DNOW]
```

PFNACC performs a negative accumulate of the two single–precision FP values in the source and destination registers. The result of the accumulate from the destination register is stored in the low doubleword of the destination, and the result of the source accumulate is stored in the high doubleword of the destination register.

The operation is:

```
dst[0–31]  := dst[0–31] – dst[32–63],
dst[32–63] := src[0–31] – src[32–63].
```

B.4.219 PFPNACC: Packed Single–Precision FP Mixed Accumulate

```
PFPNACC mm1,mm2/m64            ; 0F 0F /r 8E          [PENT,3DNOW]
```

PFPNACC performs a positive accumulate of the two single–precision FP values in the source register and a negative accumulate of the destination register. The result of the accumulate from the destination register is stored in the low doubleword of the destination, and the result of the source accumulate is stored in the high doubleword of the destination register.

The operation is:

```
dst[0–31]  := dst[0–31] – dst[32–63],
dst[32–63] := src[0–31] + src[32–63].
```

B.4.220 PFRCP: Packed Single–Precision FP Reciprocal Approximation

```
PFRCP mm1,mm2/m64              ; 0F 0F /r 96          [PENT,3DNOW]
```

PFRCP performs a low precision estimate of the reciprocal of the low–order single–precision FP value in the source operand, storing the result in both halves of the destination register. The result is accurate to 14 bits.

For higher precision reciprocals, this instruction should be followed by two more instructions: PFRCPIT1 (section B.4.221) and PFRCPIT2 (section B.4.221). This will result in a 24–bit accuracy. For more details, see the AMD 3DNow! technology manual.

B.4.221 PFRCPIT1: Packed Single–Precision FP Reciprocal, First Iteration Step

```
PFRCPIT1 mm1,mm2/m64          ; 0F 0F /r A6          [PENT,3DNOW]
```

PFRCPIT1 performs the first intermediate step in the calculation of the reciprocal of a single–precision FP value. The first source value (mm1 is the original value, and the second source value (mm2/m64 is the result of a PFRCP instruction.

For the final step in a reciprocal, returning the full 24–bit accuracy of a single–precision FP value, see PFRCPIT2 (section B.4.222). For more details, see the AMD 3DNow! technology manual.

B.4.222 PFRCPIT2: Packed Single–Precision FP Reciprocal/ Reciprocal Square Root, Second Iteration Step

```
PFRCPIT2 mm1,mm2/m64          ; 0F 0F /r B6          [PENT,3DNOW]
```

PFRCPIT2 performs the second and final intermediate step in the calculation of a reciprocal or reciprocal square root, refining the values returned by the PFRCP and PFRSQRT instructions, respectively.

The first source value (mm1) is the output of either a PFRCPIT1 or a PFRSQIT1 instruction, and the second source is the output of either the PFRCP or the PFRSQRT instruction. For more details, see the AMD 3DNow! technology manual.

B.4.223 PFRSQIT1: Packed Single–Precision FP Reciprocal Square Root, First Iteration Step

```
PFRSQIT1 mm1,mm2/m64          ; 0F 0F /r A7          [PENT,3DNOW]
```

PFRSQIT1 performs the first intermediate step in the calculation of the reciprocal square root of a single–precision FP value. The first source value (mm1 is the square of the result of a PFRSQRT instruction, and the second source value (mm2/m64 is the original value.

For the final step in a calculation, returning the full 24–bit accuracy of a single–precision FP value, see PFRCPIT2 (section B.4.222). For more details, see the AMD 3DNow! technology manual.

B.4.224 PFRSQRT: Packed Single–Precision FP Reciprocal Square Root Approximation

```
PFRSQRT mm1,mm2/m64           ; 0F 0F /r 97          [PENT,3DNOW]
```

PFRSQRT performs a low precision estimate of the reciprocal square root of the low–order single–precision FP value in the source operand, storing the result in both halves of the destination register. The result is accurate to 15 bits.

For higher precision reciprocals, this instruction should be followed by two more instructions: PFRSQIT1 (section B.4.223) and PFRCPIT2 (section B.4.221). This will result in a 24–bit accuracy. For more details, see the AMD 3DNow! technology manual.

B.4.225 PFSUB: Packed Single–Precision FP Subtract

```
PFSUB mm1,mm2/m64             ; 0F 0F /r 9A          [PENT,3DNOW]
```

PFSUB subtracts the single–precision FP values in the source from those in the destination, and stores the result in the destination operand.

```
dst[0-31]  := dst[0-31]  - src[0-31],
dst[32-63] := dst[32-63] - src[32-63].
```

B.4.226 PFSUBR: Packed Single–Precision FP Reverse Subtract

```
PFSUBR mm1,mm2/m64            ; 0F 0F /r AA          [PENT,3DNOW]
```

PFSUBR subtracts the single–precision FP values in the destination from those in the source, and stores the result in the destination operand.

```
dst[0-31]  := src[0-31]  - dst[0-31],
dst[32-63] := src[32-63] - dst[32-63].
```

B.4.227 PI2FD: Packed Doubleword Integer to Single–Precision FP Convert

```
PI2FD mm1,mm2/m64              ; 0F 0F /r 0D          [PENT,3DNOW]
```

PF2ID converts two signed 32–bit integers in the source operand to single–precision FP values, using truncation of significant digits, and stores them in the destination operand.

B.4.228 PF2IW: Packed Word Integer to Single–Precision FP Convert

```
PI2FW mm1,mm2/m64              ; 0F 0F /r 0C          [PENT,3DNOW]
```

PF2IW converts two signed 16–bit integers in the source operand to single–precision FP values, and stores them in the destination operand. The input values are in the low word of each doubleword.

B.4.229 PINSRW: Insert Word

```
PINSRW mm,r16/r32/m16,imm8     ;0F C4 /r ib       [KATMAI,MMX]
PINSRW xmm,r16/r32/m16,imm8    ;66 0F C4 /r ib    [WILLAMETTE,SSE2]
```

PINSRW loads a word from a 16–bit register (or the low half of a 32–bit register), or from memory, and loads it to the word position in the destination register, pointed at by the count operand (third operand). If the destination is an MMX register, the low two bits of the count byte are used, if it is an XMM register the low 3 bits are used. The insertion is done in such a way that the other words from the destination register are left untouched.

B.4.230 PMACHRIW: Packed Multiply and Accumulate with Rounding

```
PMACHRIW mm,m64                ; 0F 5E /r             [CYRIX,MMX]
```

PMACHRIW takes two packed 16–bit integer inputs, multiplies the values in the inputs, rounds on bit 15 of each result, then adds bits 15–30 of each result to the corresponding position of the *implied* destination register.

The operation of this instruction is:

```
dstI[0-15]  := dstI[0-15]  + (mm[0-15] *m64[0-15]
                                       + 0x00004000)[15-30],
dstI[16-31] := dstI[16-31] + (mm[16-31]*m64[16-31]
                                       + 0x00004000)[15-30],
dstI[32-47] := dstI[32-47] + (mm[32-47]*m64[32-47]
                                       + 0x00004000)[15-30],
dstI[48-63] := dstI[48-63] + (mm[48-63]*m64[48-63]
                                       + 0x00004000)[15-30].
```

Note that PMACHRIW cannot take a register as its second source operand.

B.4.231 PMADDWD: MMX Packed Multiply and Add

```
PMADDWD mm1,mm2/m64            ; 0F F5 /r             [PENT,MMX]
PMADDWD xmm1,xmm2/m128         ; 66 0F F5 /r      [WILLAMETTE,SSE2]
```

PMADDWD treats its two inputs as vectors of signed words. It multiplies corresponding elements of the two operands, giving doubleword results. These are then added together in pairs and stored in the destination operand.

The operation of this instruction is:

```
dst[0-31]   := (dst[0-15] * src[0-15])
                          + (dst[16-31] * src[16-31]);
dst[32-63]  := (dst[32-47] * src[32-47])
                          + (dst[48-63] * src[48-63]);
```

The following apply to the SSE version of the instruction:

```
dst[64-95]  := (dst[64-79] * src[64-79])
                          + (dst[80-95] * src[80-95]);
dst[96-127] := (dst[96-111] * src[96-111])
                          + (dst[112-127] * src[112-127]).
```

B.4.232 PMAGW: MMX Packed Magnitude

```
PMAGW mm1,mm2/m64              ; 0F 52 /r              [CYRIX,MMX]
```

PMAGW, specific to the Cyrix MMX extensions, treats both its operands as vectors of four signed words. It compares the absolute values of the words in corresponding positions, and sets each word of the destination (first) operand to whichever of the two words in that position had the larger absolute value.

B.4.233 PMAXSW: Packed Signed Integer Word Maximum

```
PMAXSW mm1,mm2/m64             ; 0F EE /r          [KATMAI,MMX]
PMAXSW xmm1,xmm2/m128          ; 66 0F EE /r       [WILLAMETTE,SSE2]
```

PMAXSW compares each pair of words in the two source operands, and for each pair it stores the maximum value in the destination register.

B.4.234 PMAXUB: Packed Unsigned Integer Byte Maximum

```
PMAXUB mm1,mm2/m64             ; 0F DE /r          [KATMAI,MMX]
PMAXUB xmm1,xmm2/m128          ; 66 0F DE /r       [WILLAMETTE,SSE2]
```

PMAXUB compares each pair of bytes in the two source operands, and for each pair it stores the maximum value in the destination register.

B.4.235 PMINSW: Packed Signed Integer Word Minimum

```
PMINSW mm1,mm2/m64             ; 0F EA /r          [KATMAI,MMX]
PMINSW xmm1,xmm2/m128          ; 66 0F EA /r       [WILLAMETTE,SSE2]
```

PMINSW compares each pair of words in the two source operands, and for each pair it stores the minimum value in the destination register.

B.4.236 PMINUB: Packed Unsigned Integer Byte Minimum

```
PMINUB mm1,mm2/m64             ; 0F DA /r          [KATMAI,MMX]
PMINUB xmm1,xmm2/m128          ; 66 0F DA /r       [WILLAMETTE,SSE2]
```

PMINUB compares each pair of bytes in the two source operands, and for each pair it stores the minimum value in the destination register.

B.4.237 PMOVMSKB: Move Byte Mask To Integer

```
PMOVMSKB reg32,mm              ; 0F D7 /r          [KATMAI,MMX]
PMOVMSKB reg32,xmm             ; 66 0F D7 /r       [WILLAMETTE,SSE2]
```

PMOVMSKB returns an 8–bit or 16–bit mask formed of the most significant bits of each byte of source operand (8–bits for an MMX register, 16–bits for an XMM register).

B.4.238 PMULHRWC, PMULHRIW: Multiply Packed 16–bit Integers With Rounding, and Store High Word

```
PMULHRWC mm1,mm2/m64          ; 0F 59 /r              [CYRIX,MMX]
PMULHRIW mm1,mm2/m64          ; 0F 5D /r              [CYRIX,MMX]
```

These instructions take two packed 16–bit integer inputs, multiply the values in the inputs, round on bit 15 of each result, then store bits 15–30 of each result to the corresponding position of the destination register.

- For PMULHRWC, the destination is the first source operand.
- For PMULHRIW, the destination is an implied register (worked out as described for PADDSIW (section B.4.200)).

The operation of this instruction is:

```
dst[0–15]  := (src1[0–15] *src2[0–15]  + 0x00004000)[15–30]
dst[16–31] := (src1[16–31]*src2[16–31] + 0x00004000)[15–30]
dst[32–47] := (src1[32–47]*src2[32–47] + 0x00004000)[15–30]
dst[48–63] := (src1[48–63]*src2[48–63] + 0x00004000)[15–30]
```

See also PMULHRWA (section B.4.239) for a 3DNow! version of this instruction.

B.4.239 PMULHRWA: Multiply Packed 16–bit Integers With Rounding, and Store High Word

```
PMULHRWA mm1,mm2/m64          ; 0F 0F /r B7     [PENT,3DNOW]
```

PMULHRWA takes two packed 16–bit integer inputs, multiplies the values in the inputs, rounds on bit 16 of each result, then stores bits 16–31 of each result to the corresponding position of the destination register.

The operation of this instruction is:

```
dst[0–15]  := (src1[0–15] *src2[0–15]  + 0x00008000)[16–31];
dst[16–31] := (src1[16–31]*src2[16–31] + 0x00008000)[16–31];
dst[32–47] := (src1[32–47]*src2[32–47] + 0x00008000)[16–31];
dst[48–63] := (src1[48–63]*src2[48–63] + 0x00008000)[16–31].
```

See also PMULHRWC (section B.4.238) for a Cyrix version of this instruction.

B.4.240 PMULHUW: Multiply Packed 16–bit Integers, and Store High Word

```
PMULHUW mm1,mm2/m64           ; 0F E4 /r        [KATMAI,MMX]
PMULHUW xmm1,xmm2/m128        ; 66 0F E4 /r     [WILLAMETTE,SSE2]
```

PMULHUW takes two packed unsigned 16–bit integer inputs, multiplies the values in the inputs, then stores bits 16–31 of each result to the corresponding position of the destination register.

B.4.241 PMULHW, PMULLW: Multiply Packed 16–bit Integers, and Store

```
PMULHW mm1,mm2/m64            ; 0F E5 /r              [PENT,MMX]
PMULLW mm1,mm2/m64            ; 0F D5 /r              [PENT,MMX]

PMULHW xmm1,xmm2/m128         ; 66 0F E5 /r     [WILLAMETTE,SSE2]
PMULLW xmm1,xmm2/m128         ; 66 0F D5 /r     [WILLAMETTE,SSE2]
```

PMULxW takes two packed unsigned 16–bit integer inputs, and multiplies the values in the inputs, forming doubleword results.

- PMULHW then stores the top 16 bits of each doubleword in the destination (first) operand;
- PMULLW stores the bottom 16 bits of each doubleword in the destination operand.

B.4.242 PMULUDQ: Multiply Packed Unsigned 32–bit Integers, and Store.

```
PMULUDQ mm1,mm2/m64             ; 0F F4 /r          [WILLAMETTE,SSE2]
PMULUDQ xmm1,xmm2/m128          ; 66 0F F4 /r       [WILLAMETTE,SSE2]
```

PMULUDQ takes two packed unsigned 32–bit integer inputs, and multiplies the values in the inputs, forming quadword results. The source is either an unsigned doubleword in the low doubleword of a 64–bit operand, or it's two unsigned doublewords in the first and third doublewords of a 128–bit operand. This produces either one or two 64–bit results, which are stored in the respective quadword locations of the destination register.

The operation is:

```
dst[0-63]   := dst[0-31]  * src[0-31];
dst[64-127] := dst[64-95] * src[64-95].
```

B.4.243 PMVccZB: MMX Packed Conditional Move

```
PMVZB mmxreg,mem64              ; 0F 58 /r              [CYRIX,MMX]
PMVNZB mmxreg,mem64             ; 0F 5A /r              [CYRIX,MMX]
PMVLZB mmxreg,mem64             ; 0F 5B /r              [CYRIX,MMX]
PMVGEZB mmxreg,mem64            ; 0F 5C /r              [CYRIX,MMX]
```

These instructions, specific to the Cyrix MMX extensions, perform parallel conditional moves. The two input operands are treated as vectors of eight bytes. Each byte of the destination (first) operand is either written from the corresponding byte of the source (second) operand, or left alone, depending on the value of the byte in the *implied* operand (specified in the same way as PADDSIW, in section B.4.200).

- PMVZB performs each move if the corresponding byte in the implied operand is zero;
- PMVNZB moves if the byte is non–zero;
- PMVLZB moves if the byte is less than zero;
- PMVGEZB moves if the byte is greater than or equal to zero.

Note that these instructions cannot take a register as their second source operand.

B.4.244 POP: Pop Data from Stack

```
POP reg16                       ; o16 58+r              [8086]
POP reg32                       ; o32 58+r              [386]

POP r/m16                       ; o16 8F /0             [8086]
POP r/m32                       ; o32 8F /0             [386]

POP CS                          ; 0F                    [8086,UNDOC]
POP DS                          ; 1F                    [8086]
POP ES                          ; 07                    [8086]
POP SS                          ; 17                    [8086]
POP FS                          ; 0F A1                 [386]
POP GS                          ; 0F A9                 [386]
```

POP loads a value from the stack (from [SS:SP] or [SS:ESP]) and then increments the stack pointer.

The address–size attribute of the instruction determines whether SP or ESP is used as the stack pointer: to deliberately override the default given by the BITS setting, you can use an a16 or a32 prefix.

The operand−size attribute of the instruction determines whether the stack pointer is incremented by 2 or 4: this means that segment register pops in `BITS 32` mode will pop 4 bytes off the stack and discard the upper two of them. If you need to override that, you can use an `o16` or `o32` prefix.

The above opcode listings give two forms for general−purpose register pop instructions: for example, `POP BX` has the two forms `5B` and `8F C3`. NASM will always generate the shorter form when given `POP BX`. NDISASM will disassemble both.

`POP CS` is not a documented instruction, and is not supported on any processor above the 8086 (since they use `0Fh` as an opcode prefix for instruction set extensions). However, at least some 8086 processors do support it, and so NASM generates it for completeness.

B.4.245 POPAx: Pop All General−Purpose Registers

```
POPA                            ; 61                    [186]
POPAW                           ; o16 61                [186]
POPAD                           ; o32 61                [386]
```

- `POPAW` pops a word from the stack into each of, successively, `DI`, `SI`, `BP`, nothing (it discards a word from the stack which was a placeholder for `SP`), `BX`, `DX`, `CX` and `AX`. It is intended to reverse the operation of `PUSHAW` (see section B.4.264), but it ignores the value for `SP` that was pushed on the stack by `PUSHAW`.
- `POPAD` pops twice as much data, and places the results in `EDI`, `ESI`, `EBP`, nothing (placeholder for `ESP`), `EBX`, `EDX`, `ECX` and `EAX`. It reverses the operation of `PUSHAD`.

`POPA` is an alias mnemonic for either `POPAW` or `POPAD`, depending on the current `BITS` setting.

Note that the registers are popped in reverse order of their numeric values in opcodes (see section B.2.1).

B.4.246 POPFx: Pop Flags Register

```
POPF                            ; 9D                    [8086]
POPFW                           ; o16 9D                [8086]
POPFD                           ; o32 9D                [386]
```

- `POPFW` pops a word from the stack and stores it in the bottom 16 bits of the flags register (or the whole flags register, on processors below a 386).
- `POPFD` pops a doubleword and stores it in the entire flags register.

`POPF` is an alias mnemonic for either `POPFW` or `POPFD`, depending on the current `BITS` setting.

See also `PUSHF` (section B.4.265).

B.4.247 POR: MMX Bitwise OR

```
POR mm1,mm2/m64                 ; 0F EB /r                [PENT,MMX]
POR xmm1,xmm2/m128              ; 66 0F EB /r     [WILLAMETTE,SSE2]
```

`POR` performs a bitwise OR operation between its two operands (i.e. each bit of the result is 1 if and only if at least one of the corresponding bits of the two inputs was 1), and stores the result in the destination (first) operand.

B.4.248 PREFETCH: Prefetch Data Into Caches

```
PREFETCH mem8                   ; 0F 0D /0              [PENT,3DNOW]
PREFETCHW mem8                  ; 0F 0D /1              [PENT,3DNOW]
```

`PREFETCH` and `PREFETCHW` fetch the line of data from memory that contains the specified byte. `PREFETCHW` performs differently on the Athlon to earlier processors.

For more details, see the 3DNow! Technology Manual.

B.4.249 PREFETCHh: Prefetch Data Into Caches

```
PREFETCHNTA m8                  ; 0F 18 /0          [KATMAI]
PREFETCHT0 m8                   ; 0F 18 /1          [KATMAI]
PREFETCHT1 m8                   ; 0F 18 /2          [KATMAI]
PREFETCHT2 m8                   ; 0F 18 /3          [KATMAI]
```

The PREFETCHh instructions fetch the line of data from memory that contains the specified byte. It is placed in the cache according to rules specified by locality hints h:

The hints are:

- T0 (temporal data) – prefetch data into all levels of the cache hierarchy.
- T1 (temporal data with respect to first level cache) – prefetch data into level 2 cache and higher.
- T2 (temporal data with respect to second level cache) – prefetch data into level 2 cache and higher.
- NTA (non–temporal data with respect to all cache levels) – prefetch data into non–temporal cache structure and into a location close to the processor, minimizing cache pollution.

Note that this group of instructions doesn't provide a guarantee that the data will be in the cache when it is needed. For more details, see the Intel IA32 Software Developer Manual, Volume 2.

B.4.250 PSADBW: Packed Sum of Absolute Differences

```
PSADBW mm1,mm2/m64              ; 0F F6 /r          [KATMAI,MMX]
PSADBW xmm1,xmm2/m128           ; 66 0F F6 /r       [WILLAMETTE,SSE2]
```

PSADBW The PSADBW instruction computes the absolute value of the difference of the packed unsigned bytes in the two source operands. These differences are then summed to produce a word result in the lower 16–bit field of the destination register; the rest of the register is cleared. The destination operand is an MMX or an XMM register. The source operand can either be a register or a memory operand.

B.4.251 PSHUFD: Shuffle Packed Doublewords

```
PSHUFD xmm1,xmm2/m128,imm8      ; 66 0F 70 /r ib    [WILLAMETTE,SSE2]
```

PSHUFD shuffles the doublewords in the source (second) operand according to the encoding specified by imm8, and stores the result in the destination (first) operand.

Bits 0 and 1 of imm8 encode the source position of the doubleword to be copied to position 0 in the destination operand. Bits 2 and 3 encode for position 1, bits 4 and 5 encode for position 2, and bits 6 and 7 encode for position 3. For example, an encoding of 10 in bits 0 and 1 of imm8 indicates that the doubleword at bits 64–95 of the source operand will be copied to bits 0–31 of the destination.

B.4.252 PSHUFHW: Shuffle Packed High Words

```
PSHUFHW xmm1,xmm2/m128,imm8     ; F3 0F 70 /r ib    [WILLAMETTE,SSE2]
```

PSHUFW shuffles the words in the high quadword of the source (second) operand according to the encoding specified by imm8, and stores the result in the high quadword of the destination (first) operand.

The operation of this instruction is similar to the PSHUFW instruction, except that the source and destination are the top quadword of a 128–bit operand, instead of being 64–bit operands. The low quadword is copied from the source to the destination without any changes.

B.4.253 PSHUFLW: Shuffle Packed Low Words

```
PSHUFLW xmm1,xmm2/m128,imm8    ; F2 0F 70 /r ib   [WILLAMETTE,SSE2]
```

PSHUFLW shuffles the words in the low quadword of the source (second) operand according to the encoding specified by imm8, and stores the result in the low quadword of the destination (first) operand.

The operation of this instruction is similar to the PSHUFW instruction, except that the source and destination are the low quadword of a 128–bit operand, instead of being 64–bit operands. The high quadword is copied from the source to the destination without any changes.

B.4.254 PSHUFW: Shuffle Packed Words

```
PSHUFW mm1,mm2/m64,imm8        ; 0F 70 /r ib      [KATMAI,MMX]
```

PSHUFW shuffles the words in the source (second) operand according to the encoding specified by imm8, and stores the result in the destination (first) operand.

Bits 0 and 1 of imm8 encode the source position of the word to be copied to position 0 in the destination operand. Bits 2 and 3 encode for position 1, bits 4 and 5 encode for position 2, and bits 6 and 7 encode for position 3. For example, an encoding of 10 in bits 0 and 1 of imm8 indicates that the word at bits 32–47 of the source operand will be copied to bits 0–15 of the destination.

B.4.255 PSLLx: Packed Data Bit Shift Left Logical

```
PSLLW mm1,mm2/m64              ; 0F F1 /r             [PENT,MMX]
PSLLW mm,imm8                  ; 0F 71 /6 ib          [PENT,MMX]

PSLLW xmm1,xmm2/m128           ; 66 0F F1 /r      [WILLAMETTE,SSE2]
PSLLW xmm,imm8                 ; 66 0F 71 /6 ib   [WILLAMETTE,SSE2]

PSLLD mm1,mm2/m64              ; 0F F2 /r             [PENT,MMX]
PSLLD mm,imm8                  ; 0F 72 /6 ib          [PENT,MMX]

PSLLD xmm1,xmm2/m128           ; 66 0F F2 /r      [WILLAMETTE,SSE2]
PSLLD xmm,imm8                 ; 66 0F 72 /6 ib   [WILLAMETTE,SSE2]

PSLLQ mm1,mm2/m64              ; 0F F3 /r             [PENT,MMX]
PSLLQ mm,imm8                  ; 0F 73 /6 ib          [PENT,MMX]

PSLLQ xmm1,xmm2/m128           ; 66 0F F3 /r      [WILLAMETTE,SSE2]
PSLLQ xmm,imm8                 ; 66 0F 73 /6 ib   [WILLAMETTE,SSE2]

PSLLDQ xmm1,imm8               ; 66 0F 73 /7 ib   [WILLAMETTE,SSE2]
```

PSLLx performs logical left shifts of the data elements in the destination (first) operand, moving each bit in the separate elements left by the number of bits specified in the source (second) operand, clearing the low–order bits as they are vacated. PSLLDQ shifts bytes, not bits.

- PSLLW shifts word sized elements.
- PSLLD shifts doubleword sized elements.
- PSLLQ shifts quadword sized elements.
- PSLLDQ shifts double quadword sized elements.

B.4.256 PSRAx: Packed Data Bit Shift Right Arithmetic

```
PSRAW mm1,mm2/m64              ; 0F E1 /r             [PENT,MMX]
PSRAW mm,imm8                  ; 0F 71 /4 ib          [PENT,MMX]
```

```
PSRAW xmm1,xmm2/m128          ; 66 0F E1 /r     [WILLAMETTE,SSE2]
PSRAW xmm,imm8                ; 66 0F 71 /4 ib  [WILLAMETTE,SSE2]

PSRAD mm1,mm2/m64             ; 0F E2 /r             [PENT,MMX]
PSRAD mm,imm8                 ; 0F 72 /4 ib          [PENT,MMX]

PSRAD xmm1,xmm2/m128          ; 66 0F E2 /r     [WILLAMETTE,SSE2]
PSRAD xmm,imm8                ; 66 0F 72 /4 ib  [WILLAMETTE,SSE2]
```

PSRAx performs arithmetic right shifts of the data elements in the destination (first) operand, moving each bit in the separate elements right by the number of bits specified in the source (second) operand, setting the high−order bits to the value of the original sign bit.

- PSRAW shifts word sized elements.
- PSRAD shifts doubleword sized elements.

B.4.257 PSRLx: Packed Data Bit Shift Right Logical

```
PSRLW mm1,mm2/m64             ; 0F D1 /r             [PENT,MMX]
PSRLW mm,imm8                 ; 0F 71 /2 ib          [PENT,MMX]

PSRLW xmm1,xmm2/m128          ; 66 0F D1 /r     [WILLAMETTE,SSE2]
PSRLW xmm,imm8                ; 66 0F 71 /2 ib  [WILLAMETTE,SSE2]

PSRLD mm1,mm2/m64             ; 0F D2 /r             [PENT,MMX]
PSRLD mm,imm8                 ; 0F 72 /2 ib          [PENT,MMX]

PSRLD xmm1,xmm2/m128          ; 66 0F D2 /r     [WILLAMETTE,SSE2]
PSRLD xmm,imm8                ; 66 0F 72 /2 ib  [WILLAMETTE,SSE2]

PSRLQ mm1,mm2/m64             ; 0F D3 /r             [PENT,MMX]
PSRLQ mm,imm8                 ; 0F 73 /2 ib          [PENT,MMX]

PSRLQ xmm1,xmm2/m128          ; 66 0F D3 /r     [WILLAMETTE,SSE2]
PSRLQ xmm,imm8                ; 66 0F 73 /2 ib  [WILLAMETTE,SSE2]

PSRLDQ xmm1,imm8              ; 66 0F 73 /3 ib  [WILLAMETTE,SSE2]
```

PSRLx performs logical right shifts of the data elements in the destination (first) operand, moving each bit in the separate elements right by the number of bits specified in the source (second) operand, clearing the high−order bits as they are vacated. PSRLDQ shifts bytes, not bits.

- PSRLW shifts word sized elements.
- PSRLD shifts doubleword sized elements.
- PSRLQ shifts quadword sized elements.
- PSRLDQ shifts double quadword sized elements.

B.4.258 PSUBx: Subtract Packed Integers

```
PSUBB mm1,mm2/m64             ; 0F F8 /r             [PENT,MMX]
PSUBW mm1,mm2/m64             ; 0F F9 /r             [PENT,MMX]
PSUBD mm1,mm2/m64             ; 0F FA /r             [PENT,MMX]
PSUBQ mm1,mm2/m64             ; 0F FB /r        [WILLAMETTE,SSE2]

PSUBB xmm1,xmm2/m128          ; 66 0F F8 /r     [WILLAMETTE,SSE2]
PSUBW xmm1,xmm2/m128          ; 66 0F F9 /r     [WILLAMETTE,SSE2]
PSUBD xmm1,xmm2/m128          ; 66 0F FA /r     [WILLAMETTE,SSE2]
PSUBQ xmm1,xmm2/m128          ; 66 0F FB /r     [WILLAMETTE,SSE2]
```

PSUBx subtracts packed integers in the source operand from those in the destination operand. It doesn't differentiate between signed and unsigned integers, and doesn't set any of the flags.

- PSUBB operates on byte sized elements.
- PSUBW operates on word sized elements.
- PSUBD operates on doubleword sized elements.
- PSUBQ operates on quadword sized elements.

B.4.259 PSUBSxx, PSUBUSx: Subtract Packed Integers With Saturation

```
PSUBSB mm1,mm2/m64              ; 0F E8 /r              [PENT,MMX]
PSUBSW mm1,mm2/m64              ; 0F E9 /r              [PENT,MMX]

PSUBSB xmm1,xmm2/m128           ; 66 0F E8 /r     [WILLAMETTE,SSE2]
PSUBSW xmm1,xmm2/m128           ; 66 0F E9 /r     [WILLAMETTE,SSE2]

PSUBUSB mm1,mm2/m64             ; 0F D8 /r              [PENT,MMX]
PSUBUSW mm1,mm2/m64             ; 0F D9 /r              [PENT,MMX]

PSUBUSB xmm1,xmm2/m128          ; 66 0F D8 /r     [WILLAMETTE,SSE2]
PSUBUSW xmm1,xmm2/m128          ; 66 0F D9 /r     [WILLAMETTE,SSE2]
```

PSUBSx and PSUBUSx subtracts packed integers in the source operand from those in the destination operand, and use saturation for results that are outside the range supported by the destination operand.

- PSUBSB operates on signed bytes, and uses signed saturation on the results.
- PSUBSW operates on signed words, and uses signed saturation on the results.
- PSUBUSB operates on unsigned bytes, and uses signed saturation on the results.
- PSUBUSW operates on unsigned words, and uses signed saturation on the results.

B.4.260 PSUBSIW: MMX Packed Subtract with Saturation to Implied Destination

```
PSUBSIW mm1,mm2/m64             ; 0F 55 /r              [CYRIX,MMX]
```

PSUBSIW, specific to the Cyrix extensions to the MMX instruction set, performs the same function as PSUBSW, except that the result is not placed in the register specified by the first operand, but instead in the implied destination register, specified as for PADDSIW (section B.4.200).

B.4.261 PSWAPD: Swap Packed Data

```
PSWAPD mm1,mm2/m64              ; 0F 0F /r BB      [PENT,3DNOW]
```

PSWAPD swaps the packed doublewords in the source operand, and stores the result in the destination operand.

In the K6-2 and K6-III processors, this opcode uses the mnemonic PSWAPW, and it swaps the order of words when copying from the source to the destination.

The operation in the K6-2 and K6-III processors is

```
dst[0-15]  = src[48-63];
dst[16-31] = src[32-47];
dst[32-47] = src[16-31];
dst[48-63] = src[0-15].
```

The operation in the K6-x+, ATHLON and later processors is:

```
dst[0-31]  = src[32-63];
dst[32-63] = src[0-31].
```

B.4.262 PUNPCKxxx: Unpack and Interleave Data

```
PUNPCKHBW mm1,mm2/m64          ; 0F 68 /r              [PENT,MMX]
PUNPCKHWD mm1,mm2/m64          ; 0F 69 /r              [PENT,MMX]
PUNPCKHDQ mm1,mm2/m64          ; 0F 6A /r              [PENT,MMX]

PUNPCKHBW xmm1,xmm2/m128       ; 66 0F 68 /r     [WILLAMETTE,SSE2]
PUNPCKHWD xmm1,xmm2/m128       ; 66 0F 69 /r     [WILLAMETTE,SSE2]
PUNPCKHDQ xmm1,xmm2/m128       ; 66 0F 6A /r     [WILLAMETTE,SSE2]
PUNPCKHQDQ xmm1,xmm2/m128      ; 66 0F 6D /r     [WILLAMETTE,SSE2]

PUNPCKLBW mm1,mm2/m32          ; 0F 60 /r              [PENT,MMX]
PUNPCKLWD mm1,mm2/m32          ; 0F 61 /r              [PENT,MMX]
PUNPCKLDQ mm1,mm2/m32          ; 0F 62 /r              [PENT,MMX]

PUNPCKLBW xmm1,xmm2/m128       ; 66 0F 60 /r     [WILLAMETTE,SSE2]
PUNPCKLWD xmm1,xmm2/m128       ; 66 0F 61 /r     [WILLAMETTE,SSE2]
PUNPCKLDQ xmm1,xmm2/m128       ; 66 0F 62 /r     [WILLAMETTE,SSE2]
PUNPCKLQDQ xmm1,xmm2/m128      ; 66 0F 6C /r     [WILLAMETTE,SSE2]
```

`PUNPCKxx` all treat their operands as vectors, and produce a new vector generated by interleaving elements from the two inputs. The `PUNPCKHxx` instructions start by throwing away the bottom half of each input operand, and the `PUNPCKLxx` instructions throw away the top half.

The remaining elements, are then interleaved into the destination, alternating elements from the second (source) operand and the first (destination) operand: so the leftmost part of each element in the result always comes from the second operand, and the rightmost from the destination.

- `PUNPCKxBW` works a byte at a time, producing word sized output elements.
- `PUNPCKxWD` works a word at a time, producing doubleword sized output elements.
- `PUNPCKxDQ` works a doubleword at a time, producing quadword sized output elements.
- `PUNPCKxQDQ` works a quadword at a time, producing double quadword sized output elements.

So, for example, for `MMX` operands, if the first operand held `0x7A6A5A4A3A2A1A0A` and the second held `0x7B6B5B4B3B2B1B0B`, then:

- `PUNPCKHBW` would return `0x7B7A6B6A5B5A4B4A`.
- `PUNPCKHWD` would return `0x7B6B7A6A5B4B5A4A`.
- `PUNPCKHDQ` would return `0x7B6B5B4B7A6A5A4A`.
- `PUNPCKLBW` would return `0x3B3A2B2A1B1A0B0A`.
- `PUNPCKLWD` would return `0x3B2B3A2A1B0B1A0A`.
- `PUNPCKLDQ` would return `0x3B2B1B0B3A2A1A0A`.

B.4.263 PUSH: Push Data on Stack

```
PUSH reg16                     ; o16 50+r              [8086]
PUSH reg32                     ; o32 50+r              [386]

PUSH r/m16                     ; o16 FF /6             [8086]
PUSH r/m32                     ; o32 FF /6             [386]
```

```
PUSH CS                         ; 0E                    [8086]
PUSH DS                         ; 1E                    [8086]
PUSH ES                         ; 06                    [8086]
PUSH SS                         ; 16                    [8086]
PUSH FS                         ; 0F A0                 [386]
PUSH GS                         ; 0F A8                 [386]

PUSH imm8                       ; 6A ib                 [186]
PUSH imm16                      ; o16 68 iw             [186]
PUSH imm32                      ; o32 68 id             [386]
```

`PUSH` decrements the stack pointer (`SP` or `ESP`) by 2 or 4, and then stores the given value at `[SS:SP]` or `[SS:ESP]`.

The address−size attribute of the instruction determines whether `SP` or `ESP` is used as the stack pointer: to deliberately override the default given by the `BITS` setting, you can use an `a16` or `a32` prefix.

The operand−size attribute of the instruction determines whether the stack pointer is decremented by 2 or 4: this means that segment register pushes in `BITS 32` mode will push 4 bytes on the stack, of which the upper two are undefined. If you need to override that, you can use an `o16` or `o32` prefix.

The above opcode listings give two forms for general−purpose register push instructions: for example, `PUSH BX` has the two forms `53` and `FF F3`. NASM will always generate the shorter form when given `PUSH BX`. NDISASM will disassemble both.

Unlike the undocumented and barely supported `POP CS`, `PUSH CS` is a perfectly valid and sensible instruction, supported on all processors.

The instruction `PUSH SP` may be used to distinguish an 8086 from later processors: on an 8086, the value of `SP` stored is the value it has *after* the push instruction, whereas on later processors it is the value *before* the push instruction.

B.4.264 PUSHAx: Push All General−Purpose Registers

```
PUSHA                           ; 60                    [186]
PUSHAD                          ; o32 60                [386]
PUSHAW                          ; o16 60                [186]
```

`PUSHAW` pushes, in succession, `AX`, `CX`, `DX`, `BX`, `SP`, `BP`, `SI` and `DI` on the stack, decrementing the stack pointer by a total of 16.

`PUSHAD` pushes, in succession, `EAX`, `ECX`, `EDX`, `EBX`, `ESP`, `EBP`, `ESI` and `EDI` on the stack, decrementing the stack pointer by a total of 32.

In both cases, the value of `SP` or `ESP` pushed is its *original* value, as it had before the instruction was executed.

`PUSHA` is an alias mnemonic for either `PUSHAW` or `PUSHAD`, depending on the current `BITS` setting.

Note that the registers are pushed in order of their numeric values in opcodes (see section B.2.1).

See also `POPA` (section B.4.245).

B.4.265 PUSHFx: Push Flags Register

```
PUSHF                           ; 9C                    [8086]
PUSHFD                          ; o32 9C                [386]
PUSHFW                          ; o16 9C                [8086]
```

- PUSHFW pushes the bottom 16 bits of the flags register (or the whole flags register, on processors below a 386) onto the stack.
- PUSHFD pushes the entire flags register onto the stack.

PUSHF is an alias mnemonic for either PUSHFW or PUSHFD, depending on the current BITS setting.

See also POPF (section B.4.246).

B.4.266 PXOR: MMX Bitwise XOR

```
PXOR mm1,mm2/m64              ; 0F EF /r              [PENT,MMX]
PXOR xmm1,xmm2/m128           ; 66 0F EF /r     [WILLAMETTE,SSE2]
```

PXOR performs a bitwise XOR operation between its two operands (i.e. each bit of the result is 1 if and only if exactly one of the corresponding bits of the two inputs was 1), and stores the result in the destination (first) operand.

B.4.267 RCL, RCR: Bitwise Rotate through Carry Bit

```
RCL r/m8,1                    ; D0 /2                 [8086]
RCL r/m8,CL                   ; D2 /2                 [8086]
RCL r/m8,imm8                 ; C0 /2 ib              [186]
RCL r/m16,1                   ; o16 D1 /2             [8086]
RCL r/m16,CL                  ; o16 D3 /2             [8086]
RCL r/m16,imm8                ; o16 C1 /2 ib          [186]
RCL r/m32,1                   ; o32 D1 /2             [386]
RCL r/m32,CL                  ; o32 D3 /2             [386]
RCL r/m32,imm8                ; o32 C1 /2 ib          [386]

RCR r/m8,1                    ; D0 /3                 [8086]
RCR r/m8,CL                   ; D2 /3                 [8086]
RCR r/m8,imm8                 ; C0 /3 ib              [186]
RCR r/m16,1                   ; o16 D1 /3             [8086]
RCR r/m16,CL                  ; o16 D3 /3             [8086]
RCR r/m16,imm8                ; o16 C1 /3 ib          [186]
RCR r/m32,1                   ; o32 D1 /3             [386]
RCR r/m32,CL                  ; o32 D3 /3             [386]
RCR r/m32,imm8                ; o32 C1 /3 ib          [386]
```

RCL and RCR perform a 9–bit, 17–bit or 33–bit bitwise rotation operation, involving the given source/destination (first) operand and the carry bit. Thus, for example, in the operation RCL AL,1, a 9–bit rotation is performed in which AL is shifted left by 1, the top bit of AL moves into the carry flag, and the original value of the carry flag is placed in the low bit of AL.

The number of bits to rotate by is given by the second operand. Only the bottom five bits of the rotation count are considered by processors above the 8086.

You can force the longer (286 and upwards, beginning with a C1 byte) form of RCL foo,1 by using a BYTE prefix: RCL foo,BYTE 1. Similarly with RCR.

B.4.268 RCPPS: Packed Single–Precision FP Reciprocal

```
RCPPS xmm1,xmm2/m128          ; 0F 53 /r          [KATMAI,SSE]
```

RCPPS returns an approximation of the reciprocal of the packed single–precision FP values from xmm2/m128. The maximum error for this approximation is: |Error| <= 1.5 x 2^–12

B.4.269 RCPSS: Scalar Single–Precision FP Reciprocal

```
RCPSS xmm1,xmm2/m128          ; F3 0F 53 /r      [KATMAI,SSE]
```

RCPSS returns an approximation of the reciprocal of the lower single–precision FP value from xmm2/m32; the upper three fields are passed through from xmm1. The maximum error for this approximation is: |Error| <= 1.5 x 2^–12

B.4.270 RDMSR: Read Model–Specific Registers

```
RDMSR                         ; 0F 32                [PENT,PRIV]
```

RDMSR reads the processor Model–Specific Register (MSR) whose index is stored in ECX, and stores the result in EDX:EAX. See also WRMSR (section B.4.329).

B.4.271 RDPMC: Read Performance–Monitoring Counters

```
RDPMC                         ; 0F 33                [P6]
```

RDPMC reads the processor performance–monitoring counter whose index is stored in ECX, and stores the result in EDX:EAX.

This instruction is available on P6 and later processors and on MMX class processors.

B.4.272 RDSHR: Read SMM Header Pointer Register

```
RDSHR r/m32                   ; 0F 36 /0         [386,CYRIX,SMM]
```

RDSHR reads the contents of the SMM header pointer register and saves it to the destination operand, which can be either a 32 bit memory location or a 32 bit register.

See also WRSHR (section B.4.330).

B.4.273 RDTSC: Read Time–Stamp Counter

```
RDTSC                         ; 0F 31                [PENT]
```

RDTSC reads the processor's time–stamp counter into EDX:EAX.

B.4.274 RET, RETF, RETN: Return from Procedure Call

```
RET                           ; C3                   [8086]
RET imm16                     ; C2 iw                [8086]

RETF                          ; CB                   [8086]
RETF imm16                    ; CA iw                [8086]

RETN                          ; C3                   [8086]
RETN imm16                    ; C2 iw                [8086]
```

- RET, and its exact synonym RETN, pop IP or EIP from the stack and transfer control to the new address. Optionally, if a numeric second operand is provided, they increment the stack pointer by a further imm16 bytes after popping the return address.
- RETF executes a far return: after popping IP/EIP, it then pops CS, and *then* increments the stack pointer by the optional argument if present.

B.4.275 ROL, ROR: Bitwise Rotate

```
ROL r/m8,1                    ; D0 /0                [8086]
ROL r/m8,CL                   ; D2 /0                [8086]
ROL r/m8,imm8                 ; C0 /0 ib             [186]
ROL r/m16,1                   ; o16 D1 /0            [8086]
ROL r/m16,CL                  ; o16 D3 /0            [8086]
```

```
ROL r/m16,imm8                  ; o16 C1 /0 ib          [186]
ROL r/m32,1                     ; o32 D1 /0             [386]
ROL r/m32,CL                    ; o32 D3 /0             [386]
ROL r/m32,imm8                  ; o32 C1 /0 ib          [386]

ROR r/m8,1                      ; D0 /1                 [8086]
ROR r/m8,CL                     ; D2 /1                 [8086]
ROR r/m8,imm8                   ; C0 /1 ib              [186]
ROR r/m16,1                     ; o16 D1 /1             [8086]
ROR r/m16,CL                    ; o16 D3 /1             [8086]
ROR r/m16,imm8                  ; o16 C1 /1 ib          [186]
ROR r/m32,1                     ; o32 D1 /1             [386]
ROR r/m32,CL                    ; o32 D3 /1             [386]
ROR r/m32,imm8                  ; o32 C1 /1 ib          [386]
```

`ROL` and `ROR` perform a bitwise rotation operation on the given source/destination (first) operand. Thus, for example, in the operation `ROL AL,1`, an 8–bit rotation is performed in which `AL` is shifted left by 1 and the original top bit of `AL` moves round into the low bit.

The number of bits to rotate by is given by the second operand. Only the bottom five bits of the rotation count are considered by processors above the 8086.

You can force the longer (286 and upwards, beginning with a `C1` byte) form of `ROL foo,1` by using a `BYTE` prefix: `ROL foo,BYTE 1`. Similarly with `ROR`.

B.4.276 RSDC: Restore Segment Register and Descriptor

```
RSDC segreg,m80                 ; 0F 79 /r          [486,CYRIX,SMM]
```

`RSDC` restores a segment register (DS, ES, FS, GS, or SS) from mem80, and sets up its descriptor.

B.4.277 RSLDT: Restore Segment Register and Descriptor

```
RSLDT m80                       ; 0F 7B /0          [486,CYRIX,SMM]
```

`RSLDT` restores the Local Descriptor Table (LDTR) from mem80.

B.4.278 RSM: Resume from System–Management Mode

```
RSM                             ; 0F AA                 [PENT]
```

`RSM` returns the processor to its normal operating mode when it was in System–Management Mode.

B.4.279 RSQRTPS: Packed Single–Precision FP Square Root Reciprocal

```
RSQRTPS xmm1,xmm2/m128          ; 0F 52 /r          [KATMAI,SSE]
```

`RSQRTPS` computes the approximate reciprocals of the square roots of the packed single–precision floating–point values in the source and stores the results in xmm1. The maximum error for this approximation is: |Error| <= 1.5 x 2^–12

B.4.280 RSQRTSS: Scalar Single–Precision FP Square Root Reciprocal

```
RSQRTSS xmm1,xmm2/m128          ; F3 0F 52 /r       [KATMAI,SSE]
```

`RSQRTSS` returns an approximation of the reciprocal of the square root of the lowest order single–precision FP value from the source, and stores it in the low doubleword of the destination register. The upper three fields of xmm1 are preserved. The maximum error for this approximation is: |Error| <= 1.5 x 2^–12

B.4.281 RSTS: Restore TSR and Descriptor

```
RSTS m80                        ; 0F 7D /0        [486,CYRIX,SMM]
```

RSTS restores Task State Register (TSR) from mem80.

B.4.282 SAHF: Store AH to Flags

```
SAHF                            ; 9E                    [8086]
```

SAHF sets the low byte of the flags word according to the contents of the AH register.

The operation of SAHF is:

```
 AH --> SF:ZF:0:AF:0:PF:1:CF
```

See also LAHF (section B.4.131).

B.4.283 SAL, SAR: Bitwise Arithmetic Shifts

```
SAL r/m8,1                      ; D0 /4                 [8086]
SAL r/m8,CL                     ; D2 /4                 [8086]
SAL r/m8,imm8                   ; C0 /4 ib              [186]
SAL r/m16,1                     ; o16 D1 /4             [8086]
SAL r/m16,CL                    ; o16 D3 /4             [8086]
SAL r/m16,imm8                  ; o16 C1 /4 ib          [186]
SAL r/m32,1                     ; o32 D1 /4             [386]
SAL r/m32,CL                    ; o32 D3 /4             [386]
SAL r/m32,imm8                  ; o32 C1 /4 ib          [386]

SAR r/m8,1                      ; D0 /7                 [8086]
SAR r/m8,CL                     ; D2 /7                 [8086]
SAR r/m8,imm8                   ; C0 /7 ib              [186]
SAR r/m16,1                     ; o16 D1 /7             [8086]
SAR r/m16,CL                    ; o16 D3 /7             [8086]
SAR r/m16,imm8                  ; o16 C1 /7 ib          [186]
SAR r/m32,1                     ; o32 D1 /7             [386]
SAR r/m32,CL                    ; o32 D3 /7             [386]
SAR r/m32,imm8                  ; o32 C1 /7 ib          [386]
```

SAL and SAR perform an arithmetic shift operation on the given source/destination (first) operand. The vacated bits are filled with zero for SAL, and with copies of the original high bit of the source operand for SAR.

SAL is a synonym for SHL (see section B.4.290). NASM will assemble either one to the same code, but NDISASM will always disassemble that code as SHL.

The number of bits to shift by is given by the second operand. Only the bottom five bits of the shift count are considered by processors above the 8086.

You can force the longer (286 and upwards, beginning with a C1 byte) form of SAL foo,1 by using a BYTE prefix: SAL foo,BYTE 1. Similarly with SAR.

B.4.284 SALC: Set AL from Carry Flag

```
SALC                            ; D6                [8086,UNDOC]
```

SALC is an early undocumented instruction similar in concept to SETcc (section B.4.287). Its function is to set AL to zero if the carry flag is clear, or to 0xFF if it is set.

B.4.285 SBB: Subtract with Borrow

```
SBB r/m8,reg8                   ; 18 /r                 [8086]
SBB r/m16,reg16                 ; o16 19 /r             [8086]
SBB r/m32,reg32                 ; o32 19 /r             [386]

SBB reg8,r/m8                   ; 1A /r                 [8086]
SBB reg16,r/m16                 ; o16 1B /r             [8086]
SBB reg32,r/m32                 ; o32 1B /r             [386]

SBB r/m8,imm8                   ; 80 /3 ib              [8086]
SBB r/m16,imm16                 ; o16 81 /3 iw          [8086]
SBB r/m32,imm32                 ; o32 81 /3 id          [386]

SBB r/m16,imm8                  ; o16 83 /3 ib          [8086]
SBB r/m32,imm8                  ; o32 83 /3 ib          [386]

SBB AL,imm8                     ; 1C ib                 [8086]
SBB AX,imm16                    ; o16 1D iw             [8086]
SBB EAX,imm32                   ; o32 1D id             [386]
```

SBB performs integer subtraction: it subtracts its second operand, plus the value of the carry flag, from its first, and leaves the result in its destination (first) operand. The flags are set according to the result of the operation: in particular, the carry flag is affected and can be used by a subsequent SBB instruction.

In the forms with an 8−bit immediate second operand and a longer first operand, the second operand is considered to be signed, and is sign−extended to the length of the first operand. In these cases, the BYTE qualifier is necessary to force NASM to generate this form of the instruction.

To subtract one number from another without also subtracting the contents of the carry flag, use SUB (section B.4.305).

B.4.286 SCASB, SCASW, SCASD: Scan String

```
SCASB                           ; AE                    [8086]
SCASW                           ; o16 AF                [8086]
SCASD                           ; o32 AF                [386]
```

SCASB compares the byte in AL with the byte at [ES:DI] or [ES:EDI], and sets the flags accordingly. It then increments or decrements (depending on the direction flag: increments if the flag is clear, decrements if it is set) DI (or EDI).

The register used is DI if the address size is 16 bits, and EDI if it is 32 bits. If you need to use an address size not equal to the current BITS setting, you can use an explicit a16 or a32 prefix.

Segment override prefixes have no effect for this instruction: the use of ES for the load from [DI] or [EDI] cannot be overridden.

SCASW and SCASD work in the same way, but they compare a word to AX or a doubleword to EAX instead of a byte to AL, and increment or decrement the addressing registers by 2 or 4 instead of 1.

The REPE and REPNE prefixes (equivalently, REPZ and REPNZ) may be used to repeat the instruction up to CX (or ECX – again, the address size chooses which) times until the first unequal or equal byte is found.

B.4.287 SETcc: Set Register from Condition

```
SETcc r/m8                      ; 0F 90+cc /2           [386]
```

SETcc sets the given 8−bit operand to zero if its condition is not satisfied, and to 1 if it is.

B.4.288 SFENCE: Store Fence

```
SFENCE                  ; 0F AE /7              [KATMAI]
```

SFENCE performs a serialising operation on all writes to memory that were issued before the SFENCE instruction. This guarantees that all memory writes before the SFENCE instruction are visible before any writes after the SFENCE instruction.

SFENCE is ordered respective to other SFENCE instruction, MFENCE, any memory write and any other serialising instruction (such as CPUID).

Weakly ordered memory types can be used to achieve higher processor performance through such techniques as out−of−order issue, write−combining, and write−collapsing. The degree to which a consumer of data recognizes or knows that the data is weakly ordered varies among applications and may be unknown to the producer of this data. The SFENCE instruction provides a performance−efficient way of insuring store ordering between routines that produce weakly−ordered results and routines that consume this data.

SFENCE uses the following ModRM encoding:

```
Mod (7:6)         = 11B
Reg/Opcode (5:3)  = 111B
R/M (2:0)         = 000B
```

All other ModRM encodings are defined to be reserved, and use of these encodings risks incompatibility with future processors.

See also LFENCE (section B.4.137) and MFENCE (section B.4.151).

B.4.289 SGDT, SIDT, SLDT: Store Descriptor Table Pointers

```
SGDT mem                ; 0F 01 /0              [286,PRIV]
SIDT mem                ; 0F 01 /1              [286,PRIV]
SLDT r/m16              ; 0F 00 /0              [286,PRIV]
```

SGDT and SIDT both take a 6−byte memory area as an operand: they store the contents of the GDTR (global descriptor table register) or IDTR (interrupt descriptor table register) into that area as a 32−bit linear address and a 16−bit size limit from that area (in that order). These are the only instructions which directly use *linear* addresses, rather than segment/offset pairs.

SLDT stores the segment selector corresponding to the LDT (local descriptor table) into the given operand.

See also LGDT, LIDT and LLDT (section B.4.138).

B.4.290 SHL, SHR: Bitwise Logical Shifts

```
SHL r/m8,1              ; D0 /4                 [8086]
SHL r/m8,CL             ; D2 /4                 [8086]
SHL r/m8,imm8           ; C0 /4 ib              [186]
SHL r/m16,1             ; o16 D1 /4             [8086]
SHL r/m16,CL            ; o16 D3 /4             [8086]
SHL r/m16,imm8          ; o16 C1 /4 ib          [186]
SHL r/m32,1             ; o32 D1 /4             [386]
SHL r/m32,CL            ; o32 D3 /4             [386]
SHL r/m32,imm8          ; o32 C1 /4 ib          [386]

SHR r/m8,1              ; D0 /5                 [8086]
SHR r/m8,CL             ; D2 /5                 [8086]
SHR r/m8,imm8           ; C0 /5 ib              [186]
```

```
SHR r/m16,1                     ; o16 D1 /5             [8086]
SHR r/m16,CL                    ; o16 D3 /5             [8086]
SHR r/m16,imm8                  ; o16 C1 /5 ib          [186]
SHR r/m32,1                     ; o32 D1 /5             [386]
SHR r/m32,CL                    ; o32 D3 /5             [386]
SHR r/m32,imm8                  ; o32 C1 /5 ib          [386]
```

`SHL` and `SHR` perform a logical shift operation on the given source/destination (first) operand. The vacated bits are filled with zero.

A synonym for `SHL` is `SAL` (see section B.4.283). NASM will assemble either one to the same code, but NDISASM will always disassemble that code as `SHL`.

The number of bits to shift by is given by the second operand. Only the bottom five bits of the shift count are considered by processors above the 8086.

You can force the longer (286 and upwards, beginning with a `C1` byte) form of `SHL foo,1` by using a `BYTE` prefix: `SHL foo,BYTE 1`. Similarly with `SHR`.

B.4.291 SHLD, SHRD: Bitwise Double–Precision Shifts

```
SHLD r/m16,reg16,imm8           ; o16 0F A4 /r ib       [386]
SHLD r/m16,reg32,imm8           ; o32 0F A4 /r ib       [386]
SHLD r/m16,reg16,CL             ; o16 0F A5 /r          [386]
SHLD r/m16,reg32,CL             ; o32 0F A5 /r          [386]

SHRD r/m16,reg16,imm8           ; o16 0F AC /r ib       [386]
SHRD r/m32,reg32,imm8           ; o32 0F AC /r ib       [386]
SHRD r/m16,reg16,CL             ; o16 0F AD /r          [386]
SHRD r/m32,reg32,CL             ; o32 0F AD /r          [386]
```

- `SHLD` performs a double–precision left shift. It notionally places its second operand to the right of its first, then shifts the entire bit string thus generated to the left by a number of bits specified in the third operand. It then updates only the *first* operand according to the result of this. The second operand is not modified.
- `SHRD` performs the corresponding right shift: it notionally places the second operand to the *left* of the first, shifts the whole bit string right, and updates only the first operand.

For example, if `EAX` holds `0x01234567` and `EBX` holds `0x89ABCDEF`, then the instruction `SHLD EAX,EBX,4` would update `EAX` to hold `0x12345678`. Under the same conditions, `SHRD EAX,EBX,4` would update `EAX` to hold `0xF0123456`.

The number of bits to shift by is given by the third operand. Only the bottom five bits of the shift count are considered.

B.4.292 SHUFPD: Shuffle Packed Double–Precision FP Values

```
SHUFPD xmm1,xmm2/m128,imm8      ; 66 0F C6 /r ib  [WILLAMETTE,SSE2]
```

`SHUFPD` moves one of the packed double–precision FP values from the destination operand into the low quadword of the destination operand; the upper quadword is generated by moving one of the double–precision FP values from the source operand into the destination. The select (third) operand selects which of the values are moved to the destination register.

The select operand is an 8–bit immediate: bit 0 selects which value is moved from the destination operand to the result (where 0 selects the low quadword and 1 selects the high quadword) and bit 1 selects which value is moved from the source operand to the result. Bits 2 through 7 of the shuffle operand are reserved.

B.4.293 SHUFPS: Shuffle Packed Single–Precision FP Values

```
SHUFPS xmm1,xmm2/m128,imm8     ; 0F C6 /r ib      [KATMAI,SSE]
```

SHUFPS moves two of the packed single–precision FP values from the destination operand into the low quadword of the destination operand; the upper quadword is generated by moving two of the single–precision FP values from the source operand into the destination. The select (third) operand selects which of the values are moved to the destination register.

The select operand is an 8–bit immediate: bits 0 and 1 select the value to be moved from the destination operand the low doubleword of the result, bits 2 and 3 select the value to be moved from the destination operand the second doubleword of the result, bits 4 and 5 select the value to be moved from the source operand the third doubleword of the result, and bits 6 and 7 select the value to be moved from the source operand to the high doubleword of the result.

B.4.294 SMI: System Management Interrupt

```
SMI                            ; F1               [386,UNDOC]
```

SMI puts some AMD processors into SMM mode. It is available on some 386 and 486 processors, and is only available when DR7 bit 12 is set, otherwise it generates an Int 1.

B.4.295 SMINT, SMINTOLD: Software SMM Entry (CYRIX)

```
SMINT                          ; 0F 38            [PENT,CYRIX]
SMINTOLD                       ; 0F 7E            [486,CYRIX]
```

SMINT puts the processor into SMM mode. The CPU state information is saved in the SMM memory header, and then execution begins at the SMM base address.

SMINTOLD is the same as SMINT, but was the opcode used on the 486.

This pair of opcodes are specific to the Cyrix and compatible range of processors (Cyrix, IBM, Via).

B.4.296 SMSW: Store Machine Status Word

```
SMSW r/m16                     ; 0F 01 /4         [286,PRIV]
```

SMSW stores the bottom half of the CR0 control register (or the Machine Status Word, on 286 processors) into the destination operand. See also LMSW (section B.4.139).

For 32–bit code, this would store all of CR0 in the specified register (or the bottom 16 bits if the destination is a memory location), without needing an operand size override byte.

B.4.297 SQRTPD: Packed Double–Precision FP Square Root

```
SQRTPD xmm1,xmm2/m128          ; 66 0F 51 /r      [WILLAMETTE,SSE2]
```

SQRTPD calculates the square root of the packed double–precision FP value from the source operand, and stores the double–precision results in the destination register.

B.4.298 SQRTPS: Packed Single–Precision FP Square Root

```
SQRTPS xmm1,xmm2/m128          ; 0F 51 /r         [KATMAI,SSE]
```

SQRTPS calculates the square root of the packed single–precision FP value from the source operand, and stores the single–precision results in the destination register.

B.4.299 SQRTSD: Scalar Double–Precision FP Square Root

```
SQRTSD xmm1,xmm2/m128          ; F2 0F 51 /r      [WILLAMETTE,SSE2]
```

SQRTSD calculates the square root of the low−order double−precision FP value from the source operand, and stores the double−precision result in the destination register. The high−quadword remains unchanged.

B.4.300 SQRTSS: Scalar Single−Precision FP Square Root

```
SQRTSS xmm1,xmm2/m128          ; F3 0F 51 /r      [KATMAI,SSE]
```

SQRTSS calculates the square root of the low−order single−precision FP value from the source operand, and stores the single−precision result in the destination register. The three high doublewords remain unchanged.

B.4.301 STC, STD, STI: Set Flags

```
STC                            ; F9               [8086]
STD                            ; FD               [8086]
STI                            ; FB               [8086]
```

These instructions set various flags. STC sets the carry flag; STD sets the direction flag; and STI sets the interrupt flag (thus enabling interrupts).

To clear the carry, direction, or interrupt flags, use the CLC, CLD and CLI instructions (section B.4.20). To invert the carry flag, use CMC (section B.4.22).

B.4.302 STMXCSR: Store Streaming SIMD Extension Control/Status

```
STMXCSR m32                    ; 0F AE /3         [KATMAI,SSE]
```

STMXCSR stores the contents of the MXCSR control/status register to the specified memory location. MXCSR is used to enable masked/unmasked exception handling, to set rounding modes, to set flush−to−zero mode, and to view exception status flags. The reserved bits in the MXCSR register are stored as 0s.

For details of the MXCSR register, see the Intel processor docs.

See also LDMXCSR (section B.4.133).

B.4.303 STOSB, STOSW, STOSD: Store Byte to String

```
STOSB                          ; AA               [8086]
STOSW                          ; o16 AB           [8086]
STOSD                          ; o32 AB           [386]
```

STOSB stores the byte in AL at [ES:DI] or [ES:EDI], and sets the flags accordingly. It then increments or decrements (depending on the direction flag: increments if the flag is clear, decrements if it is set) DI (or EDI).

The register used is DI if the address size is 16 bits, and EDI if it is 32 bits. If you need to use an address size not equal to the current BITS setting, you can use an explicit a16 or a32 prefix.

Segment override prefixes have no effect for this instruction: the use of ES for the store to [DI] or [EDI] cannot be overridden.

STOSW and STOSD work in the same way, but they store the word in AX or the doubleword in EAX instead of the byte in AL, and increment or decrement the addressing registers by 2 or 4 instead of 1.

The REP prefix may be used to repeat the instruction CX (or ECX – again, the address size chooses which) times.

B.4.304 STR: Store Task Register

```
STR r/m16                      ; 0F 00 /1         [286,PRIV]
```

STR stores the segment selector corresponding to the contents of the Task Register into its operand. When the operand size is 32 bit and the destination is a register, the upper 16−bits are cleared to 0s. When the destination operand is a memory location, 16 bits are written regardless of the operand size.

B.4.305 SUB: Subtract Integers

```
SUB r/m8,reg8                   ; 28 /r                 [8086]
SUB r/m16,reg16                 ; o16 29 /r             [8086]
SUB r/m32,reg32                 ; o32 29 /r             [386]

SUB reg8,r/m8                   ; 2A /r                 [8086]
SUB reg16,r/m16                 ; o16 2B /r             [8086]
SUB reg32,r/m32                 ; o32 2B /r             [386]

SUB r/m8,imm8                   ; 80 /5 ib              [8086]
SUB r/m16,imm16                 ; o16 81 /5 iw          [8086]
SUB r/m32,imm32                 ; o32 81 /5 id          [386]

SUB r/m16,imm8                  ; o16 83 /5 ib          [8086]
SUB r/m32,imm8                  ; o32 83 /5 ib          [386]

SUB AL,imm8                     ; 2C ib                 [8086]
SUB AX,imm16                    ; o16 2D iw             [8086]
SUB EAX,imm32                   ; o32 2D id             [386]
```

SUB performs integer subtraction: it subtracts its second operand from its first, and leaves the result in its destination (first) operand. The flags are set according to the result of the operation: in particular, the carry flag is affected and can be used by a subsequent SBB instruction (section B.4.285).

In the forms with an 8−bit immediate second operand and a longer first operand, the second operand is considered to be signed, and is sign−extended to the length of the first operand. In these cases, the BYTE qualifier is necessary to force NASM to generate this form of the instruction.

B.4.306 SUBPD: Packed Double−Precision FP Subtract

```
SUBPD xmm1,xmm2/m128            ; 66 0F 5C /r       [WILLAMETTE,SSE2]
```

SUBPD subtracts the packed double−precision FP values of the source operand from those of the destination operand, and stores the result in the destination operation.

B.4.307 SUBPS: Packed Single−Precision FP Subtract

```
SUBPS xmm1,xmm2/m128            ; 0F 5C /r          [KATMAI,SSE]
```

SUBPS subtracts the packed single−precision FP values of the source operand from those of the destination operand, and stores the result in the destination operation.

B.4.308 SUBSD: Scalar Single−FP Subtract

```
SUBSD xmm1,xmm2/m128            ; F2 0F 5C /r       [WILLAMETTE,SSE2]
```

SUBSD subtracts the low−order double−precision FP value of the source operand from that of the destination operand, and stores the result in the destination operation. The high quadword is unchanged.

B.4.309 SUBSS: Scalar Single−FP Subtract

```
SUBSS xmm1,xmm2/m128            ; F3 0F 5C /r       [KATMAI,SSE]
```

SUBSS subtracts the low−order single−precision FP value of the source operand from that of the destination operand, and stores the result in the destination operation. The three high doublewords are unchanged.

B.4.310 SVDC: Save Segment Register and Descriptor

```
SVDC m80,segreg                 ; 0F 78 /r          [486,CYRIX,SMM]
```

SVDC saves a segment register (DS, ES, FS, GS, or SS) and its descriptor to mem80.

B.4.311 SVLDT: Save LDTR and Descriptor

```
SVLDT m80                       ; 0F 7A /0          [486,CYRIX,SMM]
```

SVLDT saves the Local Descriptor Table (LDTR) to mem80.

B.4.312 SVTS: Save TSR and Descriptor

```
SVTS m80                        ; 0F 7C /0          [486,CYRIX,SMM]
```

SVTS saves the Task State Register (TSR) to mem80.

B.4.313 SYSCALL: Call Operating System

```
SYSCALL                         ; 0F 05                 [P6,AMD]
```

SYSCALL provides a fast method of transferring control to a fixed entry point in an operating system.

- The EIP register is copied into the ECX register.
- Bits [31−0] of the 64−bit SYSCALL/SYSRET Target Address Register (STAR) are copied into the EIP register.
- Bits [47−32] of the STAR register specify the selector that is copied into the CS register.
- Bits [47−32]+1000b of the STAR register specify the selector that is copied into the SS register.

The CS and SS registers should not be modified by the operating system between the execution of the SYSCALL instruction and its corresponding SYSRET instruction.

For more information, see the SYSCALL and SYSRET Instruction Specification (AMD document number 21086.pdf).

B.4.314 SYSENTER: Fast System Call

```
SYSENTER                        ; 0F 34                 [P6]
```

SYSENTER executes a fast call to a level 0 system procedure or routine. Before using this instruction, various MSRs need to be set up:

- SYSENTER_CS_MSR contains the 32−bit segment selector for the privilege level 0 code segment. (This value is also used to compute the segment selector of the privilege level 0 stack segment.)
- SYSENTER_EIP_MSR contains the 32−bit offset into the privilege level 0 code segment to the first instruction of the selected operating procedure or routine.
- SYSENTER_ESP_MSR contains the 32−bit stack pointer for the privilege level 0 stack.

SYSENTER performs the following sequence of operations:

- Loads the segment selector from the SYSENTER_CS_MSR into the CS register.
- Loads the instruction pointer from the SYSENTER_EIP_MSR into the EIP register.

- Adds 8 to the value in SYSENTER_CS_MSR and loads it into the SS register.
- Loads the stack pointer from the SYSENTER_ESP_MSR into the ESP register.
- Switches to privilege level 0.
- Clears the VM flag in the EFLAGS register, if the flag is set.
- Begins executing the selected system procedure.

In particular, note that this instruction des not save the values of CS or (E)IP. If you need to return to the calling code, you need to write your code to cater for this.

For more information, see the Intel Architecture Software Developer's Manual, Volume 2.

B.4.315 SYSEXIT: Fast Return From System Call

```
SYSEXIT                         ; 0F 35                 [P6,PRIV]
```

SYSEXIT executes a fast return to privilege level 3 user code. This instruction is a companion instruction to the SYSENTER instruction, and can only be executed by privilege level 0 code. Various registers need to be set up before calling this instruction:

- SYSENTER_CS_MSR contains the 32−bit segment selector for the privilege level 0 code segment in which the processor is currently executing. (This value is used to compute the segment selectors for the privilege level 3 code and stack segments.)
- EDX contains the 32−bit offset into the privilege level 3 code segment to the first instruction to be executed in the user code.
- ECX contains the 32−bit stack pointer for the privilege level 3 stack.

SYSEXIT performs the following sequence of operations:

- Adds 16 to the value in SYSENTER_CS_MSR and loads the sum into the CS selector register.
- Loads the instruction pointer from the EDX register into the EIP register.
- Adds 24 to the value in SYSENTER_CS_MSR and loads the sum into the SS selector register.
- Loads the stack pointer from the ECX register into the ESP register.
- Switches to privilege level 3.
- Begins executing the user code at the EIP address.

For more information on the use of the SYSENTER and SYSEXIT instructions, see the Intel Architecture Software Developer's Manual, Volume 2.

B.4.316 SYSRET: Return From Operating System

```
SYSRET                          ; 0F 07                 [P6,AMD,PRIV]
```

SYSRET is the return instruction used in conjunction with the SYSCALL instruction to provide fast entry/exit to an operating system.

- The ECX register, which points to the next sequential instruction after the corresponding SYSCALL instruction, is copied into the EIP register.
- Bits [63−48] of the STAR register specify the selector that is copied into the CS register.
- Bits [63−48]+1000b of the STAR register specify the selector that is copied into the SS register.
- Bits [1−0] of the SS register are set to 11b (RPL of 3) regardless of the value of bits [49−48] of the STAR register.

The CS and SS registers should not be modified by the operating system between the execution of the SYSCALL instruction and its corresponding SYSRET instruction.

For more information, see the SYSCALL and SYSRET Instruction Specification (AMD document number 21086.pdf).

B.4.317 TEST: Test Bits (notional bitwise AND)

```
TEST r/m8,reg8                  ; 84 /r                 [8086]
TEST r/m16,reg16                ; o16 85 /r             [8086]
TEST r/m32,reg32                ; o32 85 /r             [386]

TEST r/m8,imm8                  ; F6 /0 ib              [8086]
TEST r/m16,imm16                ; o16 F7 /0 iw          [8086]
TEST r/m32,imm32                ; o32 F7 /0 id          [386]

TEST AL,imm8                    ; A8 ib                 [8086]
TEST AX,imm16                   ; o16 A9 iw             [8086]
TEST EAX,imm32                  ; o32 A9 id             [386]
```

TEST performs a 'mental' bitwise AND of its two operands, and affects the flags as if the operation had taken place, but does not store the result of the operation anywhere.

B.4.318 UCOMISD: Unordered Scalar Double−Precision FP compare and set EFLAGS

```
UCOMISD xmm1,xmm2/m128          ; 66 0F 2E /r       [WILLAMETTE,SSE2]
```

UCOMISD compares the low−order double−precision FP numbers in the two operands, and sets the ZF, PF and CF bits in the EFLAGS register. In addition, the OF, SF and AF bits in the EFLAGS register are zeroed out. The unordered predicate (ZF, PF and CF all set) is returned if either source operand is a NaN (qNaN or sNaN).

B.4.319 UCOMISS: Unordered Scalar Single−Precision FP compare and set EFLAGS

```
UCOMISS xmm1,xmm2/m128          ; 0F 2E /r          [KATMAI,SSE]
```

UCOMISS compares the low−order single−precision FP numbers in the two operands, and sets the ZF, PF and CF bits in the EFLAGS register. In addition, the OF, SF and AF bits in the EFLAGS register are zeroed out. The unordered predicate (ZF, PF and CF all set) is returned if either source operand is a NaN (qNaN or sNaN).

B.4.320 UD0, UD1, UD2: Undefined Instruction

```
UD0                             ; 0F FF                 [186,UNDOC]
UD1                             ; 0F B9                 [186,UNDOC]
UD2                             ; 0F 0B                 [186]
```

UDx can be used to generate an invalid opcode exception, for testing purposes.

UD0 is specifically documented by AMD as being reserved for this purpose.

UD1 is documented by Intel as being available for this purpose.

UD2 is specifically documented by Intel as being reserved for this purpose. Intel document this as the preferred method of generating an invalid opcode exception.

All these opcodes can be used to generate invalid opcode exceptions on all currently available processors.

B.4.321 UMOV: User Move Data

```
UMOV r/m8,reg8                  ; 0F 10 /r              [386,UNDOC]
UMOV r/m16,reg16                ; o16 0F 11 /r          [386,UNDOC]
UMOV r/m32,reg32                ; o32 0F 11 /r          [386,UNDOC]

UMOV reg8,r/m8                  ; 0F 12 /r              [386,UNDOC]
UMOV reg16,r/m16                ; o16 0F 13 /r          [386,UNDOC]
UMOV reg32,r/m32                ; o32 0F 13 /r          [386,UNDOC]
```

This undocumented instruction is used by in–circuit emulators to access user memory (as opposed to host memory). It is used just like an ordinary memory/register or register/register MOV instruction, but accesses user space.

This instruction is only available on some AMD and IBM 386 and 486 processors.

B.4.322 UNPCKHPD: Unpack and Interleave High Packed Double–Precision FP Values

```
UNPCKHPD xmm1,xmm2/m128         ; 66 0F 15 /r       [WILLAMETTE,SSE2]
```

UNPCKHPD performs an interleaved unpack of the high–order data elements of the source and destination operands, saving the result in xmm1. It ignores the lower half of the sources.

The operation of this instruction is:

```
dst[63-0]   := dst[127-64];
dst[127-64] := src[127-64].
```

B.4.323 UNPCKHPS: Unpack and Interleave High Packed Single–Precision FP Values

```
UNPCKHPS xmm1,xmm2/m128         ; 0F 15 /r          [KATMAI,SSE]
```

UNPCKHPS performs an interleaved unpack of the high–order data elements of the source and destination operands, saving the result in xmm1. It ignores the lower half of the sources.

The operation of this instruction is:

```
dst[31-0]   := dst[95-64];
dst[63-32]  := src[95-64];
dst[95-64]  := dst[127-96];
dst[127-96] := src[127-96].
```

B.4.324 UNPCKLPD: Unpack and Interleave Low Packed Double–Precision FP Data

```
UNPCKLPD xmm1,xmm2/m128         ; 66 0F 14 /r       [WILLAMETTE,SSE2]
```

UNPCKLPD performs an interleaved unpack of the low–order data elements of the source and destination operands, saving the result in xmm1. It ignores the lower half of the sources.

The operation of this instruction is:

```
dst[63-0]   := dst[63-0];
dst[127-64] := src[63-0].
```

B.4.325 UNPCKLPS: Unpack and Interleave Low Packed Single–Precision FP Data

```
UNPCKLPS xmm1,xmm2/m128         ; 0F 14 /r          [KATMAI,SSE]
```

UNPCKLPS performs an interleaved unpack of the low–order data elements of the source and destination operands, saving the result in xmm1. It ignores the lower half of the sources.

The operation of this instruction is:

```
dst[31-0]   := dst[31-0];
dst[63-32]  := src[31-0];
dst[95-64]  := dst[63-32];
dst[127-96] := src[63-32].
```

B.4.326 VERR, VERW: Verify Segment Readability/Writability

```
VERR r/m16                      ; 0F 00 /4              [286,PRIV]
VERW r/m16                      ; 0F 00 /5              [286,PRIV]
```

- VERR sets the zero flag if the segment specified by the selector in its operand can be read from at the current privilege level. Otherwise it is cleared.
- VERW sets the zero flag if the segment can be written.

B.4.327 WAIT: Wait for Floating–Point Processor

```
WAIT                            ; 9B                    [8086]
FWAIT                           ; 9B                    [8086]
```

WAIT, on 8086 systems with a separate 8087 FPU, waits for the FPU to have finished any operation it is engaged in before continuing main processor operations, so that (for example) an FPU store to main memory can be guaranteed to have completed before the CPU tries to read the result back out.

On higher processors, WAIT is unnecessary for this purpose, and it has the alternative purpose of ensuring that any pending unmasked FPU exceptions have happened before execution continues.

B.4.328 WBINVD: Write Back and Invalidate Cache

```
WBINVD                          ; 0F 09                 [486]
```

WBINVD invalidates and empties the processor's internal caches, and causes the processor to instruct external caches to do the same. It writes the contents of the caches back to memory first, so no data is lost. To flush the caches quickly without bothering to write the data back first, use INVD (section B.4.125).

B.4.329 WRMSR: Write Model–Specific Registers

```
WRMSR                           ; 0F 30                 [PENT]
```

WRMSR writes the value in EDX:EAX to the processor Model–Specific Register (MSR) whose index is stored in ECX. See also RDMSR (section B.4.270).

B.4.330 WRSHR: Write SMM Header Pointer Register

```
WRSHR r/m32                     ; 0F 37 /0          [386,CYRIX,SMM]
```

WRSHR loads the contents of either a 32–bit memory location or a 32–bit register into the SMM header pointer register.

See also RDSHR (section B.4.272).

B.4.331 XADD: Exchange and Add

```
XADD r/m8,reg8                  ; 0F C0 /r              [486]
XADD r/m16,reg16                ; o16 0F C1 /r          [486]
XADD r/m32,reg32                ; o32 0F C1 /r          [486]
```

XADD exchanges the values in its two operands, and then adds them together and writes the result into the destination (first) operand. This instruction can be used with a LOCK prefix for multi–processor synchronisation purposes.

B.4.332 XBTS: Extract Bit String

```
XBTS reg16,r/m16                ; o16 0F A6 /r          [386,UNDOC]
XBTS reg32,r/m32                ; o32 0F A6 /r          [386,UNDOC]
```

The implied operation of this instruction is:

```
XBTS r/m16,reg16,AX,CL
XBTS r/m32,reg32,EAX,CL
```

Writes a bit string from the source operand to the destination. CL indicates the number of bits to be copied, and (E)AX indicates the low order bit offset in the source. The bits are written to the low order bits of the destination register. For example, if CL is set to 4 and AX (for 16–bit code) is set to 5, bits 5–8 of src will be copied to bits 0–3 of dst. This instruction is very poorly documented, and I have been unable to find any official source of documentation on it.

XBTS is supported only on the early Intel 386s, and conflicts with the opcodes for CMPXCHG486 (on early Intel 486s). NASM supports it only for completeness. Its counterpart is IBTS (see section B.4.116).

B.4.333 XCHG: Exchange

```
XCHG reg8,r/m8                  ; 86 /r                 [8086]
XCHG reg16,r/m8                 ; o16 87 /r             [8086]
XCHG reg32,r/m32                ; o32 87 /r             [386]

XCHG r/m8,reg8                  ; 86 /r                 [8086]
XCHG r/m16,reg16                ; o16 87 /r             [8086]
XCHG r/m32,reg32                ; o32 87 /r             [386]

XCHG AX,reg16                   ; o16 90+r              [8086]
XCHG EAX,reg32                  ; o32 90+r              [386]
XCHG reg16,AX                   ; o16 90+r              [8086]
XCHG reg32,EAX                  ; o32 90+r              [386]
```

XCHG exchanges the values in its two operands. It can be used with a LOCK prefix for purposes of multi–processor synchronisation.

XCHG AX,AX or XCHG EAX,EAX (depending on the BITS setting) generates the opcode 90h, and so is a synonym for NOP (section B.4.190).

B.4.334 XLATB: Translate Byte in Lookup Table

```
XLAT                            ; D7                    [8086]
XLATB                           ; D7                    [8086]
```

XLATB adds the value in AL, treated as an unsigned byte, to BX or EBX, and loads the byte from the resulting address (in the segment specified by DS) back into AL.

The base register used is BX if the address size is 16 bits, and EBX if it is 32 bits. If you need to use an address size not equal to the current BITS setting, you can use an explicit a16 or a32 prefix.

The segment register used to load from [BX+AL] or [EBX+AL] can be overridden by using a segment register name as a prefix (for example, es xlatb).

B.4.335 XOR: Bitwise Exclusive OR

```
XOR r/m8,reg8                   ; 30 /r                 [8086]
XOR r/m16,reg16                 ; o16 31 /r             [8086]
XOR r/m32,reg32                 ; o32 31 /r             [386]
```

```
XOR reg8,r/m8                   ; 32 /r                 [8086]
XOR reg16,r/m16                 ; o16 33 /r             [8086]
XOR reg32,r/m32                 ; o32 33 /r             [386]

XOR r/m8,imm8                   ; 80 /6 ib              [8086]
XOR r/m16,imm16                 ; o16 81 /6 iw          [8086]
XOR r/m32,imm32                 ; o32 81 /6 id          [386]

XOR r/m16,imm8                  ; o16 83 /6 ib          [8086]
XOR r/m32,imm8                  ; o32 83 /6 ib          [386]

XOR AL,imm8                     ; 34 ib                 [8086]
XOR AX,imm16                    ; o16 35 iw             [8086]
XOR EAX,imm32                   ; o32 35 id             [386]
```

XOR performs a bitwise XOR operation between its two operands (i.e. each bit of the result is 1 if and only if exactly one of the corresponding bits of the two inputs was 1), and stores the result in the destination (first) operand.

In the forms with an 8–bit immediate second operand and a longer first operand, the second operand is considered to be signed, and is sign–extended to the length of the first operand. In these cases, the BYTE qualifier is necessary to force NASM to generate this form of the instruction.

The MMX instruction PXOR (see section B.4.266) performs the same operation on the 64–bit MMX registers.

B.4.336 XORPD: Bitwise Logical XOR of Double–Precision FP Values

```
XORPD xmm1,xmm2/m128            ; 66 0F 57 /r       [WILLAMETTE,SSE2]
```

XORPD returns a bit–wise logical XOR between the source and destination operands, storing the result in the destination operand.

B.4.337 XORPS: Bitwise Logical XOR of Single–Precision FP Values

```
XORPS xmm1,xmm2/m128            ; 0F 57 /r          [KATMAI,SSE]
```

XORPS returns a bit–wise logical XOR between the source and destination operands, storing the result in the destination operand.

Index

!= operator 52
$$ token 36, 78
$
  Here token 36
  prefix 31, 34, 80
% operator 36
%! 63
%$ and %$$ prefixes 55, 56
%% operator 36, 46
%+ 43
%+1 and %-1 syntax 50
%0 parameter count 47, 48
& operator 36
&& operator 52
* operator 36
+ modifier 47
+ operator
  binary 36
  unary 36
− operator
  binary 36
  unary 36
..@ symbol prefix 39, 46
/ operator 36
// operator 36
< operator 52
<< operator 36
<= operator 52
<> operator 52
= operator 52
== operator 52
> operator 52
>= operator 52
>> operator 36
? MASM syntax 32
^ operator 36
^^ operator 52
| operator 36
|| operator 52
~ operator 36
−a option 27, 108
a16 102, 125, 148, 153, 160, 163, 174, 181, 186, 190, 197
a32 102, 125, 148, 153, 160, 163, 174, 181, 186, 190, 197
a86 20, 29, 30
AAA 116
AAD 116
AAM 116
AAS 116
ABSOLUTE 66, 73
ADC 116
ADD 117
addition 36
ADDPD 118
ADDPS 118
addressing, mixed−size 101
address−size prefixes 31
ADDSD 118
ADDSS 118
algebra 34
ALIGN 60, 70, 72
ALIGNB 60
alignment
  in bin sections 71
  in elf sections 77
  in obj sections 72
  in win32 sections 76
  of elf common variables 79
ALINK 82
alink.sourceforge.net 82
alloc 77
alt.lang.asm 20, 21
ambiguity 30
AND 118
ANDNPD 119
ANDNPS 119
ANDPD 119
ANDPS 120
a.out
  BSD version 79
  Linux version 79
aout 23, 79
aoutb 79, 97
%arg 61
arg 90, 96
ARPL 120
as86 20, 23, 79
assembler directives 65
assembly passes 38
assembly−time options 26
%assign 43
ASSUME 30
AT 60

Autoconf 21
autoexec.bat 21
auto-sync 108
-b 107
bin 23, 24, 70
    multisection 71
binary 34
binary files 32
16–bit mode, versus 32–bit mode 65
bit shift 36
BITS 65, 70
bitwise AND 36
bitwise OR 36
bitwise XOR 36
block IFs 56
boot loader 70
boot sector 104
Borland
    Pascal 91
    Win32 compilers 71
BOUND 120
braces
    after % sign 49
    around macro parameters 45
BSD 97
BSF 120
BSR 120
.bss 77, 79, 80
BSWAP 120
BT 121
BTC 121
BTR 121
BTS 121
bugs 105
bugtracker 105
BYTE 104
C calling convention 87, 94
C symbol names 85
CALL 121
CALL FAR 37
case sensitivity 29, 41, 42, 43, 45, 52, 73
CBW 122
CDQ 122
changing sections 65
character constant 32, 35
circular references 41
CLASS 72
CLC 122
CLD 122
%clear 58
CLFLUSH 122
CLI 122
CLTS 122
c16.mac 90, 93
c32.mac 96
CMC 123
CMOVcc 123
CMP 123
CMPccPD 123
CMPccPS 124
CMPccSD 125
CMPccSS 126
CMPEQPD 123
CMPEQPS 124
CMPEQSD 125
CMPEQSS 126
CMPLEPD 123
CMPLEPS 124
CMPLESD 125
CMPLESS 126
CMPLTPD 123
CMPLTPS 124
CMPLTSD 125
CMPLTSS 126
CMPNEQPD 123
CMPNEQPS 124
CMPNEQSD 125
CMPNEQSS 126
CMPNLEPD 123
CMPNLEPS 124
CMPNLESD 125
CMPNLESS 126
CMPNLTPD 123
CMPNLTPS 124
CMPNLTSD 125
CMPNLTSS 126
CMPORDPD 123
CMPORDPS 124
CMPORDSD 125
CMPORDSS 126
CMPSB 125
CMPSD 125
CMPSW 125
CMPUNORDPD 123
CMPUNORDPS 124
CMPUNORDSD 125
CMPUNORDSS 126
CMPXCHG 126
CMPXCHG486 126
CMPXCHG8B 127
coff 23, 77
colon 31
.COM 70, 84
COMISD 127
COMISS 127
command–line 23, 70

commas in macro parameters 47
COMMON 68, 72
elf extensions to 79
obj extensions to 75
Common Object File Format 77
common variables 68
alignment in elf 79
element size 75
comp.lang.asm.x86 20, 21
comp.os.linux.announce 21
comp.os.msdos.programmer 85
concatenating macro parameters 49
condition codes 112
condition codes as macro parameters 50
condition predicates 124, 125, 126
conditional assembly 50
conditional jump 149
conditional jumps 104
conditional–return macro 50
configure 21
constants 34
context stack 55, 56
context–local labels 55
context–local single–line macros 56
control registers 112
counting macro parameters 48
CPU 68
CPUID 35, 128
creating contexts 55
critical expression 32, 33, 38, 44, 67
CVTDQ2PD 128
CVTDQ2PS 128
CVTPD2DQ 128
CVTPD2PI 129
CVTPD2PS 129
CVTPI2PD 129
CVTPI2PS 129
CVTPS2DQ 129
CVTPS2PD 130
CVTPS2PI 130
CVTSD2SI 130
CVTSD2SS 130
CVTSI2SD 130
CVTSI2SS 131
CVTSS2SD 131
CVTSS2SI 131
CVTTPD2DQ 131
CVTTPD2PI 131
CVTTPS2DQ 132
CVTTPS2PI 132
CVTTSD2SI 132
CVTTSS2SI 132
CWD 122
CWDE 122
-D option 26
-d option 26
DAA 132
DAS 132
.data 77, 79, 80
_DATA 87
data 78, 80
data structure 89, 96
DB 32, 35
dbg 81
DD 32, 35
debug information 24
debug information format 24
debug registers 112
DEC 133
declaring structures 59
default macro parameters 47
default name 70
default–WRT mechanism 75
%define 26, 41
defining sections 65
design goals 29
DevPac 32, 39
disabling listing expansion 50
DIV 133
division 36
DIVPD 133
DIVPS 134
DIVSD 134
DIVSS 134
DJGPP 77, 94
djlink 82
DLL symbols
exporting 74
importing 74
DOS 21, 25
DOS archive 21
DOS source archive 21
DQ 32, 35
.drectve 76
DT 32, 35
DUP 30, 33
DW 32, 35
DWORD 32
-E option 25
-e option 26, 109
effective addresses 31, 33, 38, 114
element size, in common variables 75
ELF 23, 77
shared libraries 77
16–bit code and 79
%elif 50, 52

%elifctx 51
%elifdef 51
%elifid 53
%elifidn 52
%elifidni 52
%elifmacro 51
%elifnctx 51
%elifndef 51
%elifnid 53
%elifnidn 52
%elifnidni 52
%elifnmacro 51
%elifnnum 53
%elifnstr 53
%elifnum 53
%elifstr 53
%else 50
e−mail 21
EMMS 134
endproc 90, 96
%endrep 54
ENDSTRUC 59, 67
ENTER 134
environment 28
EQU 32, 33, 38
%error 53
error messages 25
error reporting format 25
EVEN 60
.EXE 71, 82
EXE_begin 83
EXE2BIN 85
exebin.mac 83
exec 77
Executable and Linkable Format 77
EXE_end 83
EXE_stack 83
%exitrep 54
EXPORT 74
export 80
exporting symbols 68
expressions 26, 36
extension 23, 70
EXTERN 67
    obj extensions to 75
    rdf extensions to 81
−F option 24
−f option 24, 70
FABS 135
FADD 135
FADDP 135
far call 30, 122
far common variables 75
far jump 150
far pointer 37
FARCODE 90, 93
FBLD 135
FBSTP 135
FCHS 136
FCLEX 136
FCMOVcc 136
FCOM 137
FCOMI 137
FCOMIP 137
FCOMP 137
FCOMPP 137
FCOS 137
FDECSTP 137
FDIV 138
FDIVP 138
FDIVR 138
FDIVRP 138
FEMMS 138
FFREE 138
FIADD 139
FICOM 139
FICOMP 139
FIDIV 139
FIDIVR 139
FILD 139
__FILE__ 58
FIMUL 139
FINCSTP 139
FINIT 140
FIST 139
FISTP 139
FISUB 140
FLAT 73
flat memory model 94
flat−form binary 70
FLD 140
FLDCW 140
FLDENV 141
FLDxx 140
floating−point 30, 31, 32, 35
    constants 35
    registers 112
FMUL 141
FMULP 141
FNINIT 140
FNOP 141
follows= 71
format−specific directives 65
forward references 38
FPATAN 141
FPREM 141

FPREM1 141
FPTAN 141
frame pointer 87, 91, 94
FreeBSD 79, 97
FreeLink 82
FRNDINT 142
FRSTOR 142
FSAVE 142
FSCALE 142
FSETPM 142
FSIN 142
FSINCOS 142
FSQRT 142
FST 143
FSTCW 143
FSTENV 143
FSTP 143
FSTSW 143
FSUB 143
FSUBP 143
FSUBR 143
FSUBRP 143
ftp.kernel.org 21
ftp.simtel.net 82
FTST 144
FUCOMxx 144
function 78, 80
functions
    C calling convention 87, 94
    Pascal calling convention 91
FXAM 145
FXCH 145
FxDISI 137
FxENI 137
F2XM1 135
FXRSTOR 145
FXSAVE 145
FXTRACT 145
FYL2X 146
FYL2XP1 146
-g option 24
gas 20
gcc 20
general purpose register 110
GLOBAL 68
    aoutb extensions to 78
    elf extensions to 78
    rdf extensions to 80
global offset table 97
_GLOBAL_OFFSET_TABLE_ 78
gnu-elf-extensions 28
..got 78
GOT relocations 99
GOT 78, 97
..gotoff 78
GOTOFF relocations 98
..gotpc 78
GOTPC relocations 98
graphics 32
greedy macro parameters 46
GROUP 73
groups 37
-h 107
hex 34
HLT 146
hybrid syntaxes 30
-I option 25
-i option 25, 108
%iassign 43
ibiblio.org 21
IBTS 146
ICEBP 148
%idefine 41
IDIV 146
IEND 60
%if 50, 52
%ifctx 51, 56
%ifdef 51
%ifid 52
%ifidn 52
%ifidni 52
ifmacro 51
%ifnctx 51
%ifndef 51
%ifnid 53
%ifnidn 52
%ifnidni 52
%ifnmacro 51
%ifnnum 53
%ifnstr 53
%ifnum 52
%ifstr 52
%imacro 44
immediate operand 110
IMPORT 74
import library 74
importing symbols 67
IMUL 147
IN 147
INC 147
INCBIN 32, 35
incbin 25
%include 25, 26, 54
include search path 26
including other files 54
inefficient code 104

infinite loop 36
informational section 76
INSB 148
INSD 148
INSTALL 22
installing 21
instances of structures 60
INSW 148
INT 148
INT01 148
INT1 148
INT3 148
integer overflow 36
Intel number formats 36
INTO 149
INVD 149
INVLPG 149
IRET 149
IRETD 149
IRETW 149
ISTRUC 60
iterating over macro parameters 48
Jcc 149
Jcc NEAR 104
JCXZ 149
JECXZ 149
JMP 150
JMP DWORD 101
jumps, mixed−size 101
-k 109
-l option 24
label prefix 39
LAHF 150
LAR 150
ld86 79
LDMXCSR 151
LDS 151
LEA 151
LEAVE 151
LES 151
LFENCE 152
LFS 151
LGDT 152
LGS 151
LIBRARY 80
licence 20
LIDT 152
%line 63
__LINE__ 58
linker, free 82
Linux
    a.out 79
    as86 79
    ELF 77
listing file 24
little−endian 35
LLDT 152
LMSW 152
LOADALL 152
LOADALL286 152
%local 62
local labels 39
LODSB 153
LODSD 153
LODSW 153
logical AND 52
logical OR 52
logical XOR 52
LOOP 153
LOOPE 153
LOOPNE 153
LOOPNZ 153
LOOPZ 153
LSL 153
LSS 151
LTR 154
-M option 24
%macro 44
macro library 25
macro processor 41
macro−local labels 46
macro-params 28
macros 33
macro-selfref 28
make 21
makefile dependencies 24
makefiles 21
Makefile.unx 22
man pages 21
map files 71
MASKMOVDQU 154
MASKMOVQ 154
MASM 20
MASM 29, 33, 71
MAXPD 154
MAXPS 154
MAXSD 154
MAXSS 154
memory models 30, 86
memory operand 32
memory references 29, 33, 110
MFENCE 155
Microsoft OMF 71

Minix 79
MINPD 155
MINPS 155
MINSD 155
MINSS 155
misc subdirectory 83, 90, 96
mixed–language program 85
mixed–size addressing 101
mixed–size instruction 101
MMX registers 112
ModR/M byte 111, 114
MODULE 80
modulo operators 36
MOV 156
MOVAPD 156
MOVAPS 157
MOVD 157
MOVDQA 157
MOVDQ2Q 157
MOVDQU 157
MOVHLPS 157
MOVHPD 158
MOVHPS 158
MOVLHPS 158
MOVLPD 158
MOVLPS 159
MOVMSKPD 159
MOVMSKPS 159
MOVNTDQ 159
MOVNTI 159
MOVNTPD 159
MOVNTPS 160
MOVNTQ 160
MOVQ 160
MOVQ2DQ 160
MOVSB 160
MOVSD 160
MOVSS 161
MOVSW 160
MOVSX 161
MOVUPD 161
MOVUPS 161
MOVZX 161
MS–DOS 70
MS–DOS device drivers 85
MUL 161
MULPD 162
MULPS 162
MULSD 162
MULSS 162
multi–line macros 28, 44
multipass optimization 27
multiple section names 70
multiplication 36
multipush macro 48
Multisection 71
nasm.1 21
NASM version 58
nasm version id 58
nasm version string 58
__NASMDEFSEG 72
nasm-devel 21
nasm.exe 21
nasm -f <format> -y 24
nasm -hf 24
__NASM_MAJOR__ 58
__NASM_MINOR__ 58
nasm.out 23
___NASM_PATCHLEVEL__ 58
__NASM_SUBMINOR__ 58
__NASM_VER__ 58
__NASM_VERSION_ID__ 58
nasmw.exe 21
nasmXXXs.zip 21
nasm-X.XX.tar.gz 21
nasmXXX.zip 21
ndisasm.1 21
ndisasm 107
ndisasm.exe 21
ndisasmw.exe 21
near call 30, 122
near common variables 75
near jump 150
NEG 162
NetBSD 79, 97
new releases 21
noalloc 77
nobits 71, 77
noexec 77
.nolist 50
NOP 162
NOT 162
‘nowait’ 30
nowrite 77
number-overflow 28
numeric constants 32, 34
-o option 23, 107
o16 103, 175, 181
o32 103, 175, 181
.OBJ 82
obj 23, 71
object 78, 80
octal 34
OF_DBG 81
OF_DEFAULT 24
OFFSET 29

OMF 71
omitted parameters 47
-On option 27
one's complement 36
OpenBSD 79, 97
operands 31
operand-size prefixes 31
operating system 70
    writing 101
operators 36
OR 162
ORG 70, 84, 85, 104
ORPD 163
orphan-labels 28, 31
ORPS 163
OS/2 71, 73
other preprocessor directives 63
OUT 163
out of range, jumps 104
output file format 24
output formats 70
OUTSB 163
OUTSD 163
OUTSW 163
overlapping segments 37
OVERLAY 73
overloading
    multi-line macros 45
    single-line macros 42
-P option 26
-p option 26, 55
PACKSSDW 164
PACKSSWB 164
PACKUSWB 164
PADDB 164
PADDD 164
PADDQ 165
PADDSB 165
PADDSIW 165
PADDSW 165
PADDUSB 165
PADDUSW 165
PADDW 164
PAND 165
PANDN 165
paradox 38
PASCAL 93
Pascal calling convention 91
passes, assembly 38
PATH 21
PAUSE 166
PAVEB 166
PAVGB 166
PAVGUSB 166
PAVGW 166
PCMPxx 166
PDISTIB 167
period 39
Perl 21
perverse 26
PEXTRW 167
PFACC 168
PFADD 168
PFCMPEQ 168
PFCMPGE 168
PFCMPGT 168
PFCMPxx 168
PF2ID 168
PF2IW 168, 171
PFMAX 169
PFMIN 169
PFMUL 169
PFNACC 169
PFPNACC 169
PFRCP 169
PFRCPIT1 170
PFRCPIT2 170
PFRSQIT1 170
PFRSQRT 170
PFSUB 170
PFSUBR 170
PharLap 72
PIC 77, 79, 97
PI2FD 171
PINSRW 171
..plt 78
PLT relocations 78, 99, 100
plt relocations 100
PMACHRIW 171
PMADDWD 171
PMAGW 172
PMAXSW 172
PMAXUB 172
PMINSW 172
PMINUB 172
PMOVMSKB 172
PMULHRIW 173
PMULHRWA 173
PMULHRWC 173
PMULHUW 173
PMULHW 173
PMULLW 173
PMULUDQ 174
PMVccZB 174
%pop 55
POP 174

POPAx 175
POPFx 175
POR 175
position−independent code 77, 79, 97
--postfix 28
precedence 36
pre−defining macros 26, 42
preferred 37
PREFETCH 175
PREFETCHh 176
PREFETCHNTA 176
PREFETCHT0 176
PREFETCHT1 176
PREFETCHT2 176
--prefix 28
pre−including files 26
preprocess−only mode 27
preprocessor 26, 27, 33, 36, 41
preprocessor expressions 26
preprocessor loops 54
preprocessor variables 43
primitive directives 65
PRIVATE 72
proc 80, 90, 96
procedure linkage table 78, 99, 100
processor mode 65
progbits 71, 77
program entry point 74, 82
program origin 70
PSADBW 176
pseudo−instructions 32
PSHUFD 176
PSHUFHW 176
PSHUFLW 177
PSHUFW 177
PSLLx 177
PSRAx 177
PSRLx 178
PSUBSIW 179
PSUBSxx 179
PSUBUSx 179
PSUBx 178
PSWAPD 179
PSWAPW 179
PUBLIC 68, 72
PUNPCKxxx 180
pure binary 70
%push 55
PUSH 180
PUSHAx 181
PUSHFx 181
PXOR 182
quick start 29
QWORD 32
-r 107
RCL 182
RCPPS 182
RCPSS 183
RCR 182
rdf 23, 80
RDMSR 183
rdoff subdirectory 22, 80
RDPMC 183
RDSHR 183
RDTSC 183
redirecting errors 25
register push 181
relational operators 52
Relocatable Dynamic Object File Format 80
relocations, PIC−specific 77
removing contexts 55
renaming contexts 56
%rep 33, 54
repeating 33, 54
%repl 56
reporting bugs 105
RESB 30, 32, 38
RESD 32
RESQ 32
REST 32
restricted memory references 110
RESW 32
RET 183
RETF 183
RETN 183
ROL 183
ROR 183
%rotate 48
rotating macro parameters 48
RPL 120
RSDC 184
RSLDT 184
RSM 184
RSQRTPS 184
RSQRTSS 184
RSTS 185
-s option 25, 108
SAHF 185
SAL 185
SALC 185
SAR 185
SBB 186
SCASB 186
SCASD 186
SCASW 186

searching for include files 54
__SECT__ 66, 67
SECTION 65
elf extensions to 77
win32 extensions to 76
section alignment
in bin 71
in elf 77
in obj 72
in win32 76
section, bin extensions to 70
SEG 36, 37, 72
SEGMENT 65
elf extensions to 72
segment address 36, 37
segment alignment
in bin 71
in obj 72
segment names, Borland Pascal 92
segment override 30, 31
segment registers 112
segments 37
groups of 73
separator character 29
SETcc 186
SFENCE 187
SGDT 187
shared libraries 79, 97
shared library 78
shift command 48
SHL 187
SHLD 188
SHR 187
SHRD 188
SHUFPD 188
SHUFPS 189
SIB byte 111, 114
SIDT 187
signed division 36
signed modulo 36
single−line macros 41
size, of symbols 78
SLDT 187
SMI 189
SMINT 189
SMINTOLD 189
SMSW 189
Solaris x86 77
-soname 100
sound 32
source code 21
source−listing file 24
SQRTPD 189
SQRTPS 189
SQRTSD 189
SQRTSS 190
square brackets 29, 33
sse condition predicates 112
STACK 72
stack frame 134
%stacksize 62
standard macros 58
standardised section names 66, 76, 77, 79, 80
..start 74, 82
start= 71
status flags 113
STC 190
STD 190
stderr 25
stdout 25
STI 190
STMXCSR 190
STOSB 190
STOSD 190
STOSW 190
STR 190
STRICT 37
string constant 32
string handling in macros 44
string length 44
%strlen 44
STRUC 59, 67, 89, 96
stub preprocessor 27
SUB 191
SUBPD 191
SUBPS 191
SUBSD 191
SUBSS 191
%substr 44
sub−strings 44
subtraction 36
suppressible warning 28
suppressing preprocessing 27
SVDC 192
SVLDT 192
SVTS 192
switching between sections 65
..sym 78
symbol sizes, specifying 78
symbol types, specifying 78
symbols
exporting from DLLs 74
importing from DLLs 74
synchronisation 108
.SYS 70, 85

SYSCALL 192
SYSENTER 192
SYSEXIT 193
SYSRET 193
-t 27
TASM 20, 27
tasm 29, 71
tasm compatible preprocessor directives 61
TBYTE 30
TEST 194
test subdirectory 82
test registers 112
testing
  arbitrary numeric expressions 52
  context stack 51
  exact text identity 52
  multi-line macro existence 51
  single-line macro existence 51
  token types 52
.text 77, 79, 80
_TEXT 87
TIMES 32, 33, 38, 104, 105
TLINK 85
trailing colon 31
two-pass assembler 38
TWORD 30, 32
type, of symbols 78
-U option 26
-u option 26, 107
UCOMISD 194
UCOMISS 194
UD0 194
UD1 194
UD2 194
UMOV 195
unary operators 36
%undef 26, 43
undefining macros 26
underscore, in C symbols 85
uninitialised 32
uninitialised storage 30
Unix 21
  SCO 77
  source archive 21
  System V 77
UnixWare 77
UNPCKHPD 195
UNPCKHPS 195
UNPCKLPD 195
UNPCKLPS 195
unrolled loops 33
unsigned division 36
unsigned modulo 36
UPPERCASE 29, 73
USE16 65, 73
USE32 65, 73
user-defined errors 53
user-level assembler directives 58
user-level directives 65
-v option 28
VAL 82
valid characters 31
variable types 30
VERR 196
version 28
version number of NASM 58
VERW 196
vfollows= 71
Visual C++ 76
vstart= 71
-w option 28
WAIT 196
warnings 28
[warning +warning-name] 28
[warning -warning-name] 28
WBINVD 196
Win32 21, 23, 71, 76, 94
Windows 82
Windows 95 21
Windows NT 21
write 77
writing operating systems 101
WRMSR 196
WRSHR 196
WRT 37, 72, 77, 79
WRT ..got 99
WRT ..gotoff 98
WRT ..gotpc 98
WRT ..plt 100
WRT ..sym 100
WWW page 20
www.cpan.org 21
www.delorie.com 82
www.pcorner.com 82
-X option 25
XADD 196
XBTS 197
XCHG 197
%xdefine 42
x2ftp.oulu.fi 82
%xidefine 42
XLATB 197
XOR 197
XORPD 198
XORPS 198

−y option 28