

flat assembler 1.58

Programmer's Manual

Tomasz Grysztar

Contents

1	Introduction	5
1.1	Compiler overview	5
1.1.1	System requirements	5
1.1.2	Compiler usage	6
1.1.3	Executing compiler from command line	7
1.1.4	Command line compiler messages	8
1.1.5	Output formats	9
1.2	Assembly syntax	9
1.2.1	Instruction syntax	9
1.2.2	Data definitions	10
1.2.3	Constants and labels	12
1.2.4	Numerical expressions	13
1.2.5	Jumps and calls	14
1.2.6	Size settings	15
2	Instruction set	17
2.1	The x86 architecture instructions	17
2.1.1	Data movement instructions	17
2.1.2	Type conversion instructions	19
2.1.3	Binary arithmetic instructions	20
2.1.4	Decimal arithmetic instructions	22
2.1.5	Logical instructions	23
2.1.6	Control transfer instructions	25
2.1.7	I/O instructions	28
2.1.8	Strings operations	29
2.1.9	Flag control instructions	31
2.1.10	Conditional operations	32
2.1.11	Miscellaneous instructions	32
2.1.12	System instructions	34
2.1.13	FPU instructions	36
2.1.14	MMX instructions	41

2.1.15	SSE instructions	42
2.1.16	SSE2 instructions	48
2.1.17	SSE3 instructions	51
2.1.18	AMD 3DNow! instructions	52
2.2	Control directives	54
2.2.1	Repeating blocks of instructions	54
2.2.2	Conditional assembly	55
2.2.3	Other directives	56
2.3	Preprocessor directives	58
2.3.1	Including source files	59
2.3.2	Symbolic constants	59
2.3.3	Macroinstructions	60
2.3.4	Structures	68
2.4	Formatter directives	69
2.4.1	MZ executable	70
2.4.2	Portable Executable	70
2.4.3	Common Object File Format	71
2.4.4	Executable and Linkable Format	72

Chapter 1

Introduction

This chapter contains all the most important information you need to begin using the flat assembler. If you are experienced assembly language programmer, you should read at least this chapter before using this compiler.

1.1 Compiler overview

Flat assembler is a fast assembly language compiler for the x86 architecture processors, which does multiple passes to optimize the size of generated machine code. It is self-compilable and versions for different operating systems are provided. They are designed to be used from the system command line and they should not differ in behavior.

This document describes also the IDE version designed for the Windows system, which uses the graphical interface instead of console and has the integrated editor. But from the view of compilation it has exactly the same functionality as all the console versions, and so later parts (beginning from 1.2) of this document are common with other releases. The executable of the IDE version is called `fasmw.exe`, while `fasm.exe` is the command line version.

1.1.1 System requirements

All versions require the x86 architecture 32-bit processor (at least 80386), although they can produce programs for the x86 architecture 16-bit processors, too. Windows console version requires any Win32 operating system, while Windows GUI version requires the Win32 GUI system version 4.0 or higher, so it should run on all systems compatible with Windows 95.

The example source provided with this version require you have environment variable `INCLUDE` set to the path of the `include` directory, which is the part of flat assembler package. If such variable already exists in your system and contains paths used by some other program, it's enough to add the new path to it (the different paths should be separated with semicolons). If you don't want to define such variable in the system, or don't know how to do it, you can set it for the flat assembler IDE only by editing the `fasmw.ini` file in its directory (this file is created by `fasmw.exe` when it's executed, but you can also create it by yourself). In this case you should add the `Include` value into the `Environment` section. For example, when you have unpacked the flat assembler files into the `c:\fasmw` directory, you should put the following two lines into your `c:\fasmw\fasmw.ini` file:

```
[Environment]
Include = c:\fasmw\include
```

If you don't define the `INCLUDE` environment variable properly, you will have to manually provide the full path to the Win32 includes in every program you want to compile.

1.1.2 Compiler usage

To start working with flat assembler, simply double click on the icon of `fasmw.exe` file, or drag the icon of your source file onto it. You can also later open new source files with the *Open* command from the *File* menu, or by dragging the files into the editor window. You can have multiple source files opened at one time, each one is represented by one tab button at the bottom of the editor window. To select file for editing, click on the corresponding tab with left mouse button. Compiler by default operates on the file you are currently editing, but you can force it to always operate on some particular file by clicking the appropriate tab with right mouse button and selecting the *Assign* command. Only single file can be assigned to compiler at one time.

When your source file is ready, you can execute the compiler with *Compile* command from the *Run* menu. When the compilation is successful, compiler will display the summary of compilation process; otherwise it will display the information about error that occurred. Compilation summary includes the information of how many passes was done, how much time it took, and how many bytes were written into destination file. It also contains a text field called *Display*, in which will appear any messages from the `display` directives in source (see 2.2.3). Error summary consists at least of the error message and a text field *Display*, which has the same purpose as above. If error is related to some specific line of source code, the summary contains also a text

field *Instruction*, which contains the preprocessed form of instruction that caused an error if the error occurred after the preprocessor stage (otherwise it's empty) and the *Source* list, which shows location of all the source lines related to this error, when you select a line from this list, it will be at the same time selected in the editor window (if file which contains that line is not loaded, it will be automatically added).

The *Run* command also executes the compiler, and in case of successful compilation it runs the compiled program if only it is one of the formats that can be run in Windows environment, otherwise you'll get a message that such type of file cannot be executed. If an error occurs, compiler displays information about it in the same form as if the *Compile* command was used.

If the compiler runs out of memory, you can increase the memory allocation in the *Compiler setup* dialog, which you can start from the *Options* menu. You can specify there the amount of kilobytes that the compiler should use, and also the priority of the compiler's thread.

1.1.3 Executing compiler from command line

To perform compilation from the command line you need to execute the **fasm.exe** executable, providing two parameters – first should be name of source file, second should be name of destination file. If no second parameter is given, the name for output file will be guessed automatically. After displaying short information about the program name and version, compiler will read the data from source file and compile it. When the compilation is successful, compiler will write the generated code to the destination file and display the summary of compilation process; otherwise it will display the information about error that occurred.

The source file should be a text file, and can be created in any text editor. Line breaks are accepted in both DOS and Unix standards, tabulators are treated as spaces.

In the command line you can also include **-m** option followed by a number, which specifies how many kilobytes of memory flat assembler should maximally use. In case of DOS version this options limits only the usage of extended memory. The **-p** option followed by a number can be used to specify the limit for number of passes the assembler performs. If code cannot be generated within specified amount of passes, the assembly will be terminated with an error message. The maximum value of this setting is 65536, while the default limit, used when no such option is included in command line, is 100. It is also possible to limit the number of passes the assembler performs, with the **-p** option followed by a number specifying the maximum number of passes.

There are no command line options that would affect the output of compiler, flat assembler requires only the source code to include the information it really needs. For example, to specify output format you specify it by using the `format` directive at the beginning of source.

1.1.4 Command line compiler messages

As it is stated above, after the successful compilation, the compiler displays the compilation summary. It includes the information of how many passes was done, how much time it took, and how many bytes were written into the destination file. The following is an example of the compilation summary:

```
flat assembler  version 1.58
38 passes, 5.3 seconds, 77824 bytes.
```

In case of error during the compilation process, the program will display an error message. For example, when compiler can't find the input file, it will display the following message:

```
flat assembler  version 1.58
error: source file not found.
```

If the error is connected with a specific part of source code, the source line that caused the error will be also displayed. Also placement of this line in the source is given to help you finding this error, for example:

```
flat assembler  version 1.58
example.asm [3]:
        mov     ax,1
error: illegal instruction.
```

It means that in the third line of the `example.asm` file compiler has encountered an unrecognized instruction. When the line that caused error contains a macroinstruction, also the line in macroinstruction definition that generated the erroneous instruction is displayed:

```
flat assembler  version 1.58
example.asm [6]:
        stoschar 7
example.asm [3] stoschar [1]:
        mov     al,char
error: illegal instruction.
```

It means that the macroinstruction in the sixth line of the `example.asm` file generated an unrecognized instruction with the first line of its definition.

1.1.5 Output formats

By default, when there is no `format` directive in source file, flat assembler simply puts generated instruction codes into output, creating this way flat binary file. By default it generates 16-bit code, but you can always turn it into the 16-bit or 32-bit mode by using `use16` or `use32` directive. Some of the output formats switch into 32-bit mode, when selected – more information about formats which you can choose can be found in 2.4.

The extension of destination file is chosen automatically by compiler, depending on the selected output format.

All output code is always in the order in which it was entered into the source file.

1.2 Assembly syntax

The information provided below is intended mainly for the assembler programmers that have been using some other assembly compilers before. If you are beginner, you should look for the assembly programming tutorials.

Flat assembler by default uses the Intel syntax for the assembly instructions, although you can customize it using the preprocessor capabilities (macroinstructions and symbolic constants). It also has its own set of the directives – the instructions for compiler.

All symbols defined inside the sources are case-sensitive.

Operator	Bits	Bytes
<code>byte</code>	8	1
<code>word</code>	16	2
<code>dword</code>	32	4
<code>fword</code>	48	6
<code>pword</code>	48	6
<code>qword</code>	64	8
<code>tword</code>	80	10
<code>dqword</code>	128	16

Table 1.1: Size operators.

1.2.1 Instruction syntax

Instructions in assembly language are separated by line breaks, and one instruction is expected to fill the one line of text. If a line contains a semicolon,

except for the semicolons in quoted strings, the rest of this line is the comment and compiler ignores it. If a line contains `\` characters, the next line is attached at this point. After the `\` character, the line should not contain anything but comments, which are started with a semicolon.

Every instruction consists of the mnemonic and the various number of operands, separated with commas. The operand can be register, immediate value or a data addressed in memory, it can also be preceded by size operator to define or override its size (table 1.1). Names of available registers you can find in table 1.2, their sizes cannot be overridden. Immediate value can be specified by any numerical expression.

When operand is a data in memory, the address of that data (also any numerical expression, but it may contain registers) should be enclosed in square brackets or preceded by `ptr` operator. For example instruction `mov eax, 3` will put the immediate value 3 into the `eax` register, instruction `mov eax, [7]` will put the 32-bit value from the address 7 into `eax` and the instruction `mov byte [7], 3` will put the immediate value 3 into the byte at address 7, it can also be written as `mov byte ptr 7, 3`. To specify which segment register should be used for addressing, segment register name followed with a colon should be put just before the address value (inside the square brackets or after the `ptr` operator).

Type	Bits									
General	8	al	cl	dl	bl	ah	ch	dh	bh	
	16	ax	cx	dx	bx	sp	bp	si	di	
	32	eax	ecx	edx	ebx	esp	ebp	esi	edi	
Segment	16	es	cs	ss	ds	fs	gs			
Control	32	cr0		cr2	cr3	cr4				
Debug	32	dr0	dr1	dr2	dr3				dr6	dr7
FPU	80	st0	st1	st2	st3	st4	st5	st6	st7	
MMX	64	mm0	mm1	mm2	mm3	mm4	mm5	mm6	mm7	
SSE	128	xmm0	xmm1	xmm2	xmm3	xmm4	xmm5	xmm6	xmm7	

Table 1.2: Registers.

1.2.2 Data definitions

To define data or reserve a space for it, use one of the directives listed in table 1.3. The data definition directive should be followed by one or more of numerical expressions, separated with commas. These expressions define the

values for data cells of size depending on which directive is used. For example `db 1,2,3` will define the three bytes of values 1, 2 and 3 respectively.

The `db` and `du` directives also accept the quoted string values of any length, which will be converted into chain of bytes when `db` is used and into chain of words with zeroed high byte when `du` is used. For example `db 'abc'` will define the three bytes of values 61, 62 and 63.

The `dp` directive and its synonym `df` accept the values consisting of two numerical expressions separated with colon, the first value will become the high word and the second value will become the low double word of the far pointer value. Also `dd` accepts such pointers consisting of two word values separated with colon. The `dt` directive accepts only floating point values and creates data in FPU double extended precision format.

The `file` is a special directive and its syntax is different. This directive includes a chain of bytes from file and it should be followed by the quoted file name, then optionally numerical expression specifying offset in file preceded by the colon, then – also optionally – comma and numerical expression specifying count of bytes to include (if no count is specified, all data up to the end of file is included).

Size (bytes)	Define data	Reserve data
1	<code>db</code> <code>file</code>	<code>rb</code>
2	<code>dw</code> <code>du</code>	<code>rw</code>
4	<code>dd</code>	<code>rd</code>
6	<code>dp</code> <code>df</code>	<code>rp</code> <code>rf</code>
8	<code>dq</code>	<code>rq</code>
10	<code>dt</code>	<code>rt</code>

Table 1.3: Data directives.

The data reservation directive should be followed by only one numerical expression, and this value defines how many cells of the specified size should be reserved. All data definition directives also accept the `?` value, which means that this cell should not be initialized to any value and the effect is the same as by using the data reservation directive. The uninitialized data may not be included in the output file, so its values should be always considered unknown.

1.2.3 Constants and labels

In the numerical expressions you can also use constants or labels instead of numbers. To define the constant or label you should use the specific directives. Each label can be defined only once and it is accessible from the any place of source (even before it was defined). Constant can be redefined many times, but in this case it is accessible only after it was defined, and is always equal to the value from last definition before the place where it's used. When a constant is defined only once in source, it is – like the label – accessible from anywhere.

The definition of constant consists of name of the constant followed by the `=` character and numerical expression, which after calculation will become the value of constant. This value is always calculated at the time the constant is defined. For example you can define `count` constant by using the directive `count = 17`, and then use it in the assembly instructions, like `mov cx, count` – which will become `mov cx, 17` during the compilation process.

There are different ways to define labels. The simplest is to follow the name of label by the colon, this directive can even be followed by the other instruction in the same line. It defines the label whose value is equal to offset of the point where it's defined. This method is usually used to label the places in code. The other way is to follow the name of label (without a colon) by some data directive. It defines the label with value equal to offset of the beginning of defined data, and remembered as a label for data with cell size as specified for that data directive in table 1.3.

The label can be treated as constant of value equal to offset of labeled code or data. For example when you define data using the labeled directive `char db 224`, to put the offset of this data into `bx` register you should use `mov bx, char` instruction, and to put the value of byte addressed by `char` label to `dl` register, you should use `mov dl, [char]` (or `mov dl, ptr char`). But when you try to assemble `mov ax, [char]`, it will cause an error, because `fasm` compares the sizes of operands, which should be equal. You can force assembling that instruction by using size override: `mov ax, word [char]`, but remember that this instruction will read the two bytes beginning at `char` address, while it was defined as a one byte.

The last and the most flexible way to define labels is to use `label` directive. This directive should be followed by the name of label, then optionally size operator and then – also optionally `at` operator and the numerical expression defining the address at which this label should be defined. For example `label wchar word at char` will define a new label for the 16-bit data at the address of `char`. Now the instruction `mov ax, [wchar]` will be after compilation the same as `mov ax, word [char]`. If no address is specified,

`label` directive defines the label at current offset. Thus `mov [wchar],57568` will copy two bytes while `mov [char],224` will copy one byte to the same address.

The label whose name begins with dot is treated as local label, and its name is attached to the name of last global label (with name beginning with anything but dot) to make the full name of this label. So you can use the short name (beginning with dot) of this label anywhere before the next global label is defined, and in the other places you have to use the full name. Label beginning with two dots are the exception - they are like global, but they don't become the new prefix for local labels.

The `@@` name means anonymous label, you can have defined many of them in the source. Symbol `@b` (or equivalent `@r`) references the nearest preceding anonymous label, symbol `@f` references the nearest following anonymous label. These special symbol are case-insensitive.

The `load` directive allows to define constant with a binary value loaded from the already assembled code. This directive should be followed by the name of the constant, then optionally size operator, then `from` operator and a numerical expression specifying a valid address in currently generated code space. The size operator has unusual meaning in this case – it states how many bytes (up to 8) have to be loaded to form the binary value of constant. If no size operator is specified, one byte is loaded (thus value is in range from 0 to 255). The loaded data cannot exceed current offset.

1.2.4 Numerical expressions

In the above examples all the numerical expressions were the simple numbers, constants or labels. But they can be more complex, by using the arithmetical or logical operators for calculations at compile time. All these operators with their priority values are listed in table 1.4. The operations with higher priority value will be calculated first, you can of course change this behavior by putting some parts of expression into parenthesis. The `+`, `-`, `*` and `/` are standard arithmetical operations, `mod` calculates the remainder from division. The `and`, `or`, `xor`, `shl`, `shr` and `not` perform the same logical operations as assembly instructions of those names. The `rva` is specific to PE output format and performs the conversion of an address into the RVA.

The numbers in the expression are by default treated as a decimal, binary numbers should have the `b` letter attached at the end, octal number end with `o` letter, hexadecimal numbers should begin with `0x` characters (like in C language) or with the `$` character (like in Pascal language) or they should end with `h` letter. Also quoted string, when encountered in expression, will be converted into number – the first character will become the least significant

Priority	Operators
0	<code>+</code> <code>-</code>
1	<code>*</code> <code>/</code>
2	<code>mod</code>
3	<code>and</code> <code>or</code> <code>xor</code>
4	<code>shl</code> <code>shr</code>
5	<code>not</code>
6	<code>rva</code>

Table 1.4: Arithmetical and logical operators by priority.

byte of number.

The numerical expression used as an address value can also contain any of general registers used for addressing, they can be added and multiplied by appropriate values, as it is allowed for x86 architecture instructions.

There are also some special symbols that can be used inside the numerical expression. First is `$`, which is always equal to the value of current offset. Second is `%`, which is the number of current repeat in parts of code that are repeated using some special directives (see 2.2). There's also `%t` symbol, which is always equal to the current time stamp.

Any numerical expression can also consist of single floating point value (flat assembler does not allow any floating point operations at compilation time) in the scientific notation, they can end with the `f` letter to be recognized, otherwise they should contain at least one of the `.` or `E` characters. So `1.0`, `1E0` and `1f` define the same floating point value, while simple `1` defines an integer value.

1.2.5 Jumps and calls

The operand of any jump or call instruction can be preceded not only by the size operator, but also by one of the operators specifying type of the jump: `near` or `far`. For example, when assembler is in 16-bit mode, instruction `jmp dword [0]` will become the far jump and when assembler is in 32-bit mode, it will become the near jump. To force this instruction to be treated

differently, use the `jmp near dword [0]` or `jmp far dword [0]` form.

When operand of near jump is the immediate value, assembler will generate the shortest variant of this jump instruction if possible (but won't create 32-bit instruction in 16-bit mode nor 16-bit instruction in 32-bit mode, unless there is a size operator stating it). By specifying the size operator you can force it to always generate long variant (for example `jmp word 0` in 16-bit mode and `jmp dword 0` in 32-bit mode) or to always generate short variant and terminate with an error when it's impossible (for example `jmp byte 0`).

1.2.6 Size settings

When instruction uses some memory addressing, by default the shorter 8-bit form is generated if only address value fits in range, but it can be overridden using the `word` or `dword` operator before the address inside the square brackets (or after the `ptr` operator). Such size operator placement can also be used to force address size other than default for the given mode.

Instructions `adc`, `add`, `and`, `cmp`, `or`, `sbb`, `sub` and `xor` with first operand being 16-bit or 32-bit are by default generated in shortened 8-bit form when the second operand is immediate value fitting in the range for signed 8-bit values. It also can be overridden by putting the `word` or `dword` operator before the immediate value. The similar rules applies to the `imul` instruction with the last operand being immediate value.

Immediate value as an operand for `push` instruction without a size operator is by default treated as a word value if assembler is in 16-bit mode and as a double word value if assembler is in 32-bit mode, shorter 8-bit form of this instruction is used if possible, `word` or `dword` size operator forces the `push` instruction to be generated in longer form for specified size. `pushw` and `pushd` mnemonics force assembler to generate 16-bit or 32-bit code without forcing it to use the longer form of instruction.

Chapter 2

Instruction set

This chapter provides the detailed information about the instructions and directives supported by flat assembler. Directives for defining constants and labels were already discussed in 1.2.3, all other directives will be described later in this chapter.

2.1 The x86 architecture instructions

In this section you can find both the information about the syntax and purpose the assembly language instructions. If you need more technical information, look for the Intel Architecture Software Developer's Manual.

Assembly instructions consist of the mnemonic (instruction's name) and from zero to three operands. If there are two or more operands, usually first is the destination operand and second is the source operand. Each operand can be register, memory or immediate value (see 1.2 for details about syntax of operands). After the description of each instruction there are examples of different combinations of operands, if the instruction has any.

Some instructions act as prefixes and can be followed by other instruction in the same line, and there can be more than one prefix in a line. Each name of the segment register is also a mnemonic of instruction prefix, although it is recommended to use segment overrides inside the square brackets instead of these prefixes.

2.1.1 Data movement instructions

`mov` transfers a byte, word or double word from the source operand to the destination operand. It can transfer data between general registers, from the general register to memory, or from memory to general register, but it

cannot move from memory to memory. It can also transfer an immediate value to general register or memory, segment register to general register or memory, general register or memory to segment register, control or debug register to general register and general register to control or debug register. The `mov` can be assembled only if the size of source operand and size of destination operand are the same. Below are the examples for each of the allowed combinations:

```

mov bx,ax      ; general register to general register
mov [char],al  ; general register to memory
mov bl,[char]  ; memory to general register
mov dl,32      ; immediate value to general register
mov [char],32  ; immediate value to memory
mov ax,ds      ; segment register to general register
mov [bx],ds    ; segment register to memory
mov ds,ax      ; general register to segment register
mov ds,[bx]    ; memory to segment register
mov eax,cr0    ; control register to general register
mov cr3,ebx    ; general register to control register

```

`xchg` swaps the contents of two operands. It can swap two byte operands, two word operands or two double word operands. Order of operands is not important. The operands may be two general registers, or general register with memory. For example:

```

xchg ax,bx     ; swap two general registers
xchg al,[char] ; swap register with memory

```

`push` decrements the stack frame pointer (`esp` register), then transfers the operand to the top of stack indicated by `esp`. The operand can be memory, general register, segment register or immediate value of word or double word size. If operand is an immediate value and no size is specified, it is by default treated as a word value if assembler is in 16-bit mode and as a double word value if assembler is in 32-bit mode. `pushw` and `pushd` mnemonics are variants of this instruction that store the values of word or double word size respectively. If more operands follow in the same line (separated only with spaces, not commas), compiler will assemble chain of the `push` instructions with these operands. The examples are with single operands:

```

push ax        ; store general register
push es        ; store segment register
pushw [bx]     ; store memory
push 1000h     ; store immediate value

```

pusha saves the contents of the eight general register on the stack. This instruction has no operands. There are two version of this instruction, one 16-bit and one 32-bit, assembler automatically generates the right version for current mode, but it can be overridden by using **pushaw** or **pushad** mnemonic to always get the 16-bit or 32-bit version. The 16-bit version of this instruction pushes general registers on the stack in the following order: **ax**, **cx**, **dx**, **bx**, the initial value of **sp** before **ax** was pushed, **bp**, **si** and **di**. The 32-bit version pushes equivalent 32-bit general registers in the same order.

pop transfers the word or double word at the current top of stack to the destination operand, and then increments **esp** to point to the new top of stack. The operand can be memory, general register or segment register. **popw** and **popd** mnemonics are variants of this instruction for restoring the values of word or double word size respectively. If more operands separated with spaces follow in the same line, compiler will assemble chain of the **pop** instructions with these operands.

```
pop bx          ; restore general register
pop ds          ; restore segment register
popw [si]       ; restore memory
```

popa restores the registers saved on the stack by **pusha** instruction, except for the saved value of **sp** (or **esp**), which is ignored. This instruction has no operands. To force assembling 16-bit or 32-bit version of this instruction use **popaw** or **popad** mnemonic.

2.1.2 Type conversion instructions

The type conversion instructions convert bytes into words, words into double words, and double words into quad words. These conversions can be done using the sign extension or zero extension. The sign extension fills the extra bits of the larger item with the value of the sign bit of the smaller item, the zero extension simply fills them with zeros.

cwq and **cdq** double the size of value **ax** or **eax** register respectively and store the extra bits into the **dx** or **edx** register. The conversion is done using the sign extension. These instructions have no operands.

cbw extends the sign of the byte in **al** throughout **ax**, and **cwde** extends the sign of the word in **ax** throughout **eax**. These instructions also have no operands.

movsx converts a byte to word or double word and a word to double word using the sign extension. **movzx** does the same, but it uses the zero extension. The source operand can be general register or memory, while the destination operand must be a general register. For example:

```

movsx ax,al      ; byte register to word register
movsx edx,dl     ; byte register to double word register
movsx eax,ax     ; word register to double word register
movsx ax,byte [bx] ; byte memory to word register
movsx edx,byte [bx] ; byte memory to double word register
movsx eax,word [bx] ; word memory to double word register

```

2.1.3 Binary arithmetic instructions

add replaces the destination operand with the sum of the source and destination operands and sets CF if overflow has occurred. The operands may be bytes, words or double words. The destination operand can be general register or memory, the source operand can be general register or immediate value, it can also be memory if the destination operand is register.

```

add ax,bx      ; add register to register
add ax,[si]    ; add memory to register
add [di],al    ; add register to memory
add al,48      ; add immediate value to register
add [char],48  ; add immediate value to memory

```

adc sums the operands, adds one if CF is set, and replaces the destination operand with the result. Rules for the operands are the same as for the **add** instruction. An **add** followed by multiple **adc** instructions can be used to add numbers longer than 32 bits.

inc adds one to the operand, it does not affect CF. The operand can be a general register or memory, and the size of the operand can be byte, word or double word.

```

inc ax      ; increment register by one
inc byte [bx] ; increment memory by one

```

sub subtracts the source operand from the destination operand and replaces the destination operand with the result. If a borrow is required, the CF is set. Rules for the operands are the same as for the **add** instruction.

sbb subtracts the source operand from the destination operand, subtracts one if CF is set, and stores the result to the destination operand. Rules for the operands are the same as for the **add** instruction. A **sub** followed by multiple **sbb** instructions may be used to subtract numbers longer than 32 bits.

dec subtracts one from the operand, it does not affect CF. Rules for the operand are the same as for the **inc** instruction.

cmp subtracts the source operand from the destination operand. It updates the flags as the **sub** instruction, but does not alter the source and destination operands. Rules for the operands are the same as for the **sub** instruction.

neg subtracts a signed integer operand from zero. The effect of this instruction is to reverse the sign of the operand from positive to negative or from negative to positive. Rules for the operand are the same as for the **inc** instruction.

xadd exchanges the destination operand with the source operand, then loads the sum of the two values into the destination operand. Rules for the operands are the same as for the **add** instruction.

All the above binary arithmetic instructions update SF, ZF, PF and OF flags. SF is always set to the same value as the result's sign bit, ZF is set when all the bits of result are zero, PF is set when low order eight bits of result contain an even number of set bits, OF is set if result is too large for a positive number or too small for a negative number (excluding sign bit) to fit in destination operand.

mul performs an unsigned multiplication of the operand and the accumulator. If the operand is a byte, the processor multiplies it by the contents of **al** and returns the 16-bit result to **ah** and **al**. If the operand is a word, the processor multiplies it by the contents of **ax** and returns the 32-bit result to **dx** and **ax**. If the operand is a double word, the processor multiplies it by the contents of **eax** and returns the 64-bit result in **edx** and **eax**. **mul** sets CF and OF when the upper half of the result is nonzero, otherwise they are cleared. Rules for the operand are the same as for the **inc** instruction.

imul performs a signed multiplication operation. This instruction has three variations. First has one operand and behaves in the same way as the **mul** instruction. Second has two operands, in this case destination operand is multiplied by the source operand and the result replaces the destination operand. Destination operand must be a general register, it can be word or double word, source operand can be general register, memory or immediate value. Third form has three operands, the destination operand must be a general register, word or double word in size, source operand can be general register or memory, and third operand must be an immediate value. The source operand is multiplied by the immediate value and the result is stored in the destination register. All the three forms calculate the product to twice the size of operands and set CF and OF when the upper half of the result is nonzero, but second and third form truncate the product to the size of operands. So second and third forms can be also used for unsigned operands because, whether the operands are signed or unsigned, the lower half of the product is the same. Below are the examples for all three forms:

```
imul bl          ; accumulator by register
imul word [si]   ; accumulator by memory
imul bx,cx       ; register by register
imul bx,[si]     ; register by memory
imul bx,10       ; register by immediate value
imul ax,bx,10    ; register by immediate value to register
imul ax,[si],10  ; memory by immediate value to register
```

div performs an unsigned division of the accumulator by the operand. The dividend (the accumulator) is twice the size of the divisor (the operand), the quotient and remainder have the same size as the divisor. If divisor is byte, the dividend is taken from **ax** register, the quotient is stored in **al** and the remainder is stored in **ah**. If divisor is word, the upper half of dividend is taken from **dx**, the lower half of dividend is taken from **ax**, the quotient is stored in **ax** and the remainder is stored in **dx**. If divisor is double word, the upper half of dividend is taken from **edx**, the lower half of dividend is taken from **eax**, the quotient is stored in **eax** and the remainder is stored in **edx**. Rules for the operand are the same as for the **mul** instruction.

idiv performs a signed division of the accumulator by the operand. It uses the same registers as the **div** instruction, and the rules for the operand are the same.

2.1.4 Decimal arithmetic instructions

Decimal arithmetic is performed by combining the binary arithmetic instructions (already described in the prior section) with the decimal arithmetic instructions. The decimal arithmetic instructions are used to adjust the results of a previous binary arithmetic operation to produce a valid packed or unpacked decimal result, or to adjust the inputs to a subsequent binary arithmetic operation so the operation will produce a valid packed or unpacked decimal result.

daa adjusts the result of adding two valid packed decimal operands in **al**. **daa** must always follow the addition of two pairs of packed decimal numbers (one digit in each half-byte) to obtain a pair of valid packed decimal digits as results. The carry flag is set if carry was needed. This instruction has no operands.

das adjusts the result of subtracting two valid packed decimal operands in **al**. **das** must always follow the subtraction of one pair of packed decimal numbers (one digit in each half-byte) from another to obtain a pair of valid packed decimal digits as results. The carry flag is set if a borrow was needed. This instruction has no operands.

aaa changes the contents of register **al** to a valid unpacked decimal number, and zeroes the top four bits. **aaa** must always follow the addition of two unpacked decimal operands in **al**. The carry flag is set and **ah** is incremented if a carry is necessary. This instruction has no operands.

aas changes the contents of register **al** to a valid unpacked decimal number, and zeroes the top four bits. **aas** must always follow the subtraction of one unpacked decimal operand from another in **al**. The carry flag is set and **ah** decremented if a borrow is necessary. This instruction has no operands.

aam corrects the result of a multiplication of two valid unpacked decimal numbers. **aam** must always follow the multiplication of two decimal numbers to produce a valid decimal result. The high order digit is left in **ah**, the low order digit in **al**. The generalized version of this instruction allows adjustment of the contents of the **ax** to create two unpacked digits of any number base. The standard version of this instruction has no operands, the generalized version has one operand – an immediate value specifying the number base for the created digits.

aad modifies the numerator in **ah** and **al** to prepare for the division of two valid unpacked decimal operands so that the quotient produced by the division will be a valid unpacked decimal number. **ah** should contain the high order digit and **al** the low order digit. This instruction adjusts the value and places the result in **al**, while **ah** will contain zero. The generalized version of this instruction allows adjustment of two unpacked digits of any number base. Rules for the operand are the same as for the **aam** instruction.

2.1.5 Logical instructions

not inverts the bits in the specified operand to form a one's complement of the operand. It has no effect on the flags. Rules for the operand are the same as for the **inc** instruction.

and, **or** and **xor** instructions perform the standard logical operations. They update the SF, ZF and PF flags. Rules for the operands are the same as for the **add** instruction.

bt, **bts**, **btr** and **btc** instructions operate on a single bit which can be in memory or in a general register. The location of the bit is specified as an offset from the low order end of the operand. The value of the offset is taken from the second operand, it either may be an immediate byte or a general register. These instructions first assign the value of the selected bit to CF. **bt** instruction does nothing more, **bts** sets the selected bit to 1, **btr** resets the selected bit to 0, **btc** changes the bit to its complement. The first operand can be word or double word.

```

bt    ax,15          ; test bit in register
bts   word [bx],15   ; test and set bit in memory
btr   ax,cx          ; test and reset bit in register
btc   word [bx],cx   ; test and complement bit in memory

```

bsf and **bsr** instructions scan a word or double word for first set bit and store the index of this bit into destination operand, which must be general register. The bit string being scanned is specified by source operand, it may be either general register or memory. The ZF flag is set if the entire string is zero (no set bits are found); otherwise it is cleared. If no set bit is found, the value of the destination register is undefined. **bsf** from low order to high order (starting from bit index zero). **bsr** scans from high order to low order (starting from bit index 15 of a word or index 31 of a double word).

```

bsf   ax,bx          ; scan register forward
bsr   ax,[si]        ; scan memory reverse

```

shl shifts the destination operand left by the number of bits specified in the second operand. The destination operand can be byte, word, or double word general register or memory. The second operand can be an immediate value or the **cl** register. The processor shifts zeros in from the right (low order) side of the operand as bits exit from the left side. The last bit that exited is stored in CF. **sal** is a synonym for **shl**.

```

shl   al,1           ; shift register left by one bit
shl   byte [bx],1    ; shift memory left by one bit
shl   ax,cl          ; shift register left by count from cl
shl   word [bx],cl   ; shift memory left by count from cl

```

shr and **sar** shift the destination operand right by the number of bits specified in the second operand. Rules for operands are the same as for the **shl** instruction. **shr** shifts zeros in from the left side of the operand as bits exit from the right side. The last bit that exited is stored in CF. **sar** preserves the sign of the operand by shifting in zeros on the left side if the value is positive or by shifting in ones if the value is negative.

shld shifts bits of the destination operand to the left by the number of bits specified in third operand, while shifting high order bits from the source operand into the destination operand on the right. The source operand remains unmodified. The destination operand can be a word or double word general register or memory, the source operand must be a general register, third operand can be an immediate value or the **cl** register.

```

shld  ax,bx,1        ; shift register left by one bit

```



```
shld [di],bx,1    ; shift memory left by one bit
shld ax,bx,cl     ; shift register left by count from cl
shld [di],bx,cl   ; shift memory left by count from cl
```

shrd shifts bits of the destination operand to the right, while shifting low order bits from the source operand into the destination operand on the left. The source operand remains unmodified. Rules for operands are the same as for the **shld** instruction.

rol and **rcl** rotate the byte, word or double word destination operand left by the number of bits specified in the second operand. For each rotation specified, the high order bit that exits from the left of the operand returns at the right to become the new low order bit. **rcl** additionally puts in CF each high order bit that exits from the left side of the operand before it returns to the operand as the low order bit on the next rotation cycle. Rules for operands are the same as for the **shl** instruction.

ror and **rcr** rotate the byte, word or double word destination operand right by the number of bits specified in the second operand. For each rotation specified, the low order bit that exits from the right of the operand returns at the left to become the new high order bit. **rcr** additionally puts in CF each low order bit that exits from the right side of the operand before it returns to the operand as the high order bit on the next rotation cycle. Rules for operands are the same as for the **shl** instruction.

test performs the same action as the **and** instruction, but it does not alter the destination operand, only updates flags. Rules for the operands are the same as for the **and** instruction.

bswap reverses the byte order of a 32-bit general register: bits 0 through 7 are swapped with bits 24 through 31, and bits 8 through 15 are swapped with bits 16 through 23. This instruction is provided for converting little-endian values to big-endian format and vice versa.

```
bswap edx        ; swap bytes in register
```

2.1.6 Control transfer instructions

jmp unconditionally transfers control to the target location. The destination address can be specified directly within the instruction or indirectly through a register or memory, the acceptable size of this address depends on whether the jump is near or far (it can be specified by preceding the operand with **near** or **far** operator) and whether the instruction is 16-bit or 32-bit. Operand for near jump should be **word** size for 16-bit instruction or the **dword** size for 32-bit instruction. Operand for far jump should be **dword** size for

16-bit instruction or `pword` size for 32-bit instruction. A direct `jmp` instruction includes the destination address as part of the instruction, the operand specifying address should be the numerical expression for near jump, or two numerical expressions separated with colon for far jump, the first specifies selector of segment, the second is the offset within segment. An indirect `jmp` instruction obtains the destination address indirectly through a register or a pointer variable, the operand should be general register or memory. See also 1.2.5 for more details.

```
jmp 100h          ; direct near jump
jmp 0FFFFh:0      ; direct far jump
jmp ax            ; indirect near jump
jmp pword [ebx]   ; indirect far jump
```

`call` transfers control to the procedure, saving on the stack the address of the instruction following the `call` for later use by a `ret` (return) instruction. Rules for the operands are the same as for the `jmp` instruction, but the `call` has no short variant of direct instruction and thus it not optimized.

`ret`, `retn` and `retf` instructions terminate the execution of a procedure and transfers control back to the program that originally invoked the procedure using the address that was stored on the stack by the `call` instruction. `ret` is the equivalent for `retn`, which returns from the procedure that was executed using the near call, while `retf` returns from the procedure that was executed using the far call. These instructions default to the size of address appropriate for the current code setting, but the size of address can be forced to 16-bit by using the `retw`, `retnw` and `retfw` mnemonics, and to 32-bit by using the `rettd`, `retnd` and `retfd` mnemonics. All these instructions may optionally specify an immediate operand, by adding this constant to the stack pointer, they effectively remove any arguments that the calling program pushed on the stack before the execution of the `call` instruction.

`iret` returns control to an interrupted procedure. It differs from `ret` in that it also pops the flags from the stack into the flags register. The flags are stored on the stack by the interrupt mechanism. It defaults to the size of return address appropriate for the current code setting, but it can be forced to use 16-bit or 32-bit address by using the `iretw` or `iretd` mnemonic.

The conditional transfer instructions are jumps that may or may not transfer control, depending on the state of the CPU flags when the instruction executes. The mnemonics for conditional jumps may be obtained by attaching the condition mnemonic (see table 2.1) to the `j` mnemonic, for example `jc` instruction will transfer the control when the CF flag is set. The conditional jumps can be near and direct only, and can be optimized (see 1.2.5), the operand should be an immediate value specifying target address.

Mnemonic	Condition tested	Description
o	$OF = 1$	overflow
no	$OF = 0$	not overflow
c b nae	$CF = 1$	carry below not above nor equal
nc ae nb	$CF = 0$	not carry above or equal not below
e z	$ZF = 1$	equal zero
ne nz	$ZF = 0$	not equal not zero
be na	$CF \text{ or } ZF = 1$	below or equal not above
a nbe	$CF \text{ or } ZF = 0$	above not below nor equal
s	$SF = 1$	sign
ns	$SF = 0$	not sign
p pe	$PF = 1$	parity parity even
np po	$PF = 0$	not parity parity odd
l nge	$SF \text{ xor } OF = 1$	less not greater nor equal
ge nl	$SF \text{ xor } OF = 0$	greater or equal not less
le ng	$(SF \text{ xor } OF) \text{ or } ZF = 1$	less or equal not greater
g nle	$(SF \text{ xor } OF) \text{ or } ZF = 0$	greater not less nor equal

Table 2.1: Conditions.

The `loop` instructions are conditional jumps that use a value placed in `cx` (or `ecx`) to specify the number of repetitions of a software loop. All `loop` instructions automatically decrement `cx` (or `ecx`) and terminate the loop (don't transfer the control) when `cx` (or `ecx`) is zero. It uses `cx` or `ecx` whether the current code setting is 16-bit or 32-bit, but it can be forced to use `cx` with the `loopw` mnemonic or to use `ecx` with the `looped` mnemonic. `loope` and `loopz` are the synonyms for the same instruction, which acts as the standard `loop`, but also terminates the loop when ZF flag is set. `loopew` and `loopzw` mnemonics force them to use `cx` register while `looped` and `loopzd` force them to use `ecx` register. `loopne` and `loopnz` are the synonyms for the same instructions, which acts as the standard `loop`, but also terminate the loop when ZF flag is not set. `loopnew` and `loopnzw` mnemonics force them to use `cx` register while `loopned` and `loopnzd` force them to use `ecx` register. Every `loop` instruction needs an operand being an immediate value specifying target address, it can be only short jump (in the range of 128 bytes back and 127 bytes forward from the address of instruction following the `loop` instruction).

`jcxz` branches to the label specified in the instruction if it finds a value of zero in `cx`, `jecxz` does the same, but checks the value of `ecx` instead of `cx`. Rules for the operands are the same as for the `loop` instruction.

`int` activates the interrupt service routine that corresponds to the number specified as an operand to the instruction, the number should be in range from 0 to 255. The interrupt service routine terminates with an `iret` instruction that returns control to the instruction that follows `int`. `int3` mnemonic codes the short (one byte) trap that invokes the interrupt 3. `into` instruction invokes the interrupt 4 if the OF flag is set.

`bound` verifies that the signed value contained in the specified register lies within specified limits. An interrupt 5 occurs if the value contained in the register is less than the lower bound or greater than the upper bound. It needs two operands, the first operand specifies the register being tested, the second operand should be memory address for the two signed limit values. The operands can be `word` or `dword` in size.

```
bound ax,[bx]      ; check word for bounds
bound eax,[esi]    ; check double word for bounds
```

2.1.7 I/O instructions

`in` transfers a byte, word, or double word from an input port to `al`, `ax`, or `eax`. I/O ports can be addressed either directly, with the immediate byte value coded in instruction, or indirectly via the `dx` register. The destination

operand should be `al`, `ax`, or `eax` register. The source operand should be an immediate value in range from 0 to 255, or `dx` register.

```
in al,20h      ; input byte from port 20h
in ax,dx       ; input word from port addressed by dx
```

`out` transfers a byte, word, or double word to an output port from `al`, `ax`, or `eax`. The program can specify the number of the port using the same methods as the `in` instruction. The destination operand should be an immediate value in range from 0 to 255, or `dx` register. The source operand should be `al`, `ax`, or `eax` register.

```
out 20h,ax     ; output word to port 20h
out dx,al      ; output byte to port addressed by dx
```

2.1.8 Strings operations

The string operations operate on one element of a string. A string element may be a byte, a word, or a double word. The string elements are addressed by `si` and `di` (or `esi` and `edi`) registers. After every string operation `si` and/or `di` (or `esi` and/or `edi`) are automatically updated to point to the next element of the string. If DF (direction flag) is zero, the index registers are incremented, if DF is one, they are decremented. The amount of the increment or decrement is 1, 2, or 4 depending on the size of the string element. Every string operation instruction has short forms which have no operands and use `si` and/or `di` when the code type is 16-bit, and `esi` and/or `edi` when the code type is 32-bit. `si` and `esi` by default address data in the segment selected by `ds`, `di` and `edi` always address data in the segment selected by `es`. Short form is obtained by attaching to the mnemonic of string operation letter specifying the size of string element, it should be `b` for byte element, `w` for word element, and `d` for double word element. Full form of string operation needs operands providing the size operator and the memory addresses, which can be `si` or `esi` with any segment prefix, `di` or `edi` always with `es` segment prefix.

`movs` transfers the string element pointed to by `si` (or `esi`) to the location pointed to by `di` (or `edi`). Size of operands can be `byte`, `word` or `dword`. The destination operand should be memory addressed by `di` or `edi`, the source operand should be memory addressed by `si` or `esi` with any segment prefix.

```
movs byte [di],[si]      ; transfer byte
movs word [es:di],[ss:si] ; transfer word
movsd                      ; transfer double word
```

cmps subtracts the destination string element from the source string element and updates the flags AF, SF, PF, CF and OF, but it does not change any of the compared elements. If the string elements are equal, ZF is set, otherwise it is cleared. The first operand for this instruction should be the source string element addressed by **si** or **esi** with any segment prefix, the second operand should be the destination string element addressed by **di** or **edi**.

```

cmpsb                ; compare bytes
cmps word [ds:si],[es:di] ; compare words
cmps dword [fs:esi],[edi] ; compare double words

```

scas subtracts the destination string element from **al**, **ax**, or **eax** (depending on the size of string element) and updates the flags AF, SF, ZF, PF, CF and OF. If the values are equal, ZF is set, otherwise it is cleared. The operand should be the destination string element addressed by **di** or **edi**.

```

scas byte [es:di]      ; scan byte
scasw                ; scan word
scas dword [es:edi]    ; scan double word

```

lods places the source string element into **al**, **ax**, or **eax**. The operand should be the source string element addressed by **si** or **esi** with any segment prefix.

```

lods byte [ds:si]      ; load byte
lods word [cs:si]      ; load word
lods dword             ; load double word

```

stos places the value of **al**, **ax**, or **eax** into the destination string element. Rules for the operand are the same as for the **scas** instruction.

ins transfers a byte, word, or double word from an input port addressed by **dx** register to the destination string element. The destination operand should be memory addressed by **di** or **edi**, the source operand should be the **dx** register.

```

insb                ; input byte
ins word [es:di],dx ; input word
ins dword [edi],dx  ; input double word

```

outs transfers the source string element to an output port addressed by **dx** register. The destination operand should be the **dx** register and the source operand should be memory addressed by **si** or **esi** with any segment prefix.

```

outs dx,byte [si]           ; output byte
outsw                     ; output word
outs dx,dword [gs:esi]     ; output double word

```

The repeat prefixes **rep**, **repe/repz**, and **repne/repnz** specify repeated string operation. When a string operation instruction has a repeat prefix, the operation is executed repeatedly, each time using a different element of the string. The repetition terminates when one of the conditions specified by the prefix is satisfied. All three prefixes automatically decrease **cx** or **ecx** register (depending whether string operation instruction uses the 16-bit or 32-bit addressing) after each operation and repeat the associated operation until **cx** or **ecx** is zero. **repe/repz** and **repne/repnz** are used exclusively with the **scas** and **cmps** instructions (described below). When these prefixes are used, repetition of the next instruction depends on the zero flag (ZF) also, **repe** and **repz** terminate the execution when the ZF is zero, **repne** and **repnz** terminate the execution when the ZF is set.

```

rep  movsd                ; transfer multiple double words
repe cmpsb                ; compare bytes until not equal

```

2.1.9 Flag control instructions

The flag control instructions provide a method for directly changing the state of bits in the flag register. All instructions described in this section have no operands.

stc sets the CF (carry flag) to 1, **clic** zeroes the CF, **cmc** changes the CF to its complement. **std** sets the DF (direction flag) to 1, **cld** zeroes the DF, **sti** sets the IF (interrupt flag) to 1 and therefore enables the interrupts, **cli** zeroes the IF and therefore disables the interrupts.

lahf copies SF, ZF, AF, PF, and CF to bits 7, 6, 4, 2, and 0 of the **ah** register. The contents of the remaining bits are undefined. The flags remain unaffected.

sahf transfers bits 7, 6, 4, 2, and 0 from the **ah** register into SF, ZF, AF, PF, and CF.

pushf decrements **esp** by two or four and stores the low word or double word of flags register at the top of stack, size of stored data depends on the current code setting. **pushfw** variant forces storing the word and **pushfd** forces storing the double word.

popf transfers specific bits from the word or double word at the top of stack, then increments **esp** by two or four, this value depends on the current code setting. **popfw** variant forces restoring from the word and **popfd** forces restoring from the double word.

2.1.10 Conditional operations

The instructions obtained by attaching the condition mnemonic (see table 2.1) to the **set** mnemonic set a byte to one if the condition is true and set the byte to zero otherwise. The operand should be an 8-bit general register or the byte in memory.

```
setne al          ; set al if zero flag cleared
seto byte [bx]    ; set byte if overflow
```

salc instruction sets the all bits of **al** register when the carry flag is set and zeroes the **al** register otherwise. This instruction has no arguments.

The instructions obtained by attaching the condition mnemonic to the **cmov** mnemonic transfer the word or double word from the general register or memory to the general register only when the condition is true. The destination operand should be general register, the source operand can be general register or memory.

```
cmovz ax,bx       ; move when zero flag set
cmovnc eax,[ebx]  ; move when carry flag cleared
```

cmpxchg compares the value in the **al**, **ax**, or **eax** register with the destination operand. If the two values are equal, the source operand is loaded into the destination operand. Otherwise, the destination operand is loaded into the **al**, **ax**, or **eax** register. The destination operand may be a general register or memory, the source operand must be a general register.

```
cmpxchg dl,bl     ; compare and exchange with register
cmpxchg [bx],dx   ; compare and exchange with memory
```

cmpxchg8b compares the 64-bit value in **edx** and **eax** registers with the destination operand. If the values are equal, the 64-bit value in **ecx** and **ebx** registers is stored in the destination operand. Otherwise, the value in the destination operand is loaded into **edx** and **eax** registers. The destination operand should be a quad word in memory.

```
cmpxchg8b [bx]    ; compare and exchange 8 bytes
```

2.1.11 Miscellaneous instructions

nop instruction occupies one byte but affects nothing but the instruction pointer. This instruction has no operands and doesn't perform any operation.

ud2 instruction generates an invalid opcode exception. This instruction is provided for software testing to explicitly generate an invalid opcode. This instruction has no operands.

xlat replaces a byte in the **al** register with a byte indexed by its value in a translation table addressed by **bx** or **ebx**. The operand should be a byte memory addressed by **bx** or **ebx** with any segment prefix. This instruction has also a short form **xlatb** which has no operands and uses the **bx** or **ebx** address in the segment selected by **ds** depending on the current code setting.

lds transfers a pointer variable from the source operand to **ds** and the destination register. The source operand must be a memory operand, and the destination operand must be a general register. The **ds** register receives the segment selector of the pointer while the destination register receives the offset part of the pointer. **les**, **lfs**, **lgs** and **lss** operate identically to **lds** except that rather than **ds** register the **es**, **fs**, **gs** and **ss** is used respectively.

```
lds bx,[si]      ; load pointer to ds:bx
```

lea transfers the offset of the source operand (rather than its value) to the destination operand. The source operand must be a memory operand, and the destination operand must be a general register.

```
lea dx,[bx+si+1] ; load effective address to dx
```

cuid returns processor identification and feature information in the **eax**, **ebx**, **ecx**, and **edx** registers. The information returned is selected by entering a value in the **eax** register before the instruction is executed. This instruction has no operands.

pause instruction delays the execution of the next instruction an implementation specific amount of time. It can be used to improve the performance of spin wait loops. This instruction has no operands.

enter creates a stack frame that may be used to implement the scope rules of block-structured high-level languages. A **leave** instruction at the end of a procedure complements an **enter** at the beginning of the procedure to simplify stack management and to control access to variables for nested procedures. The **enter** instruction includes two parameters. The first parameter specifies the number of bytes of dynamic storage to be allocated on the stack for the routine being entered. The second parameter corresponds to the lexical nesting level of the routine, it can be in range from 0 to 31. The specified lexical level determines how many sets of stack frame pointers the CPU copies into the new stack frame from the preceding frame. This list of stack frame pointers is sometimes called the display. The first word (or double word when code is 32-bit) of the display is a pointer to the last stack frame. This pointer enables a **leave** instruction to reverse the action of the previous **enter** instruction by effectively discarding the last stack frame. After **enter** creates the new display for a procedure, it allocates the dynamic

storage space for that procedure by decrementing `esp` by the number of bytes specified in the first parameter. To enable a procedure to address its display, `enter` leaves `bp` (or `ebp`) pointing to the beginning of the new stack frame. If the lexical level is zero, `enter` pushes `bp` (or `ebp`), copies `sp` to `bp` (or `esp` to `ebp`) and then subtracts the first operand from `esp`. For nesting levels greater than zero, the processor pushes additional frame pointers on the stack before adjusting the stack pointer.

```
enter 2048,0      ; enter and allocate 2048 bytes on stack
```

2.1.12 System instructions

`lmsw` loads the operand into the machine status word (bits 0 through 15 of `cr0` register), while `smsw` stores the machine status word into the destination operand. The operand can be a 16-bit or 32-bit general register or the word in memory.

```
lmsw ax          ; load machine status from register
smsw [bx]        ; store machine status to memory
```

`lgdt` and `lidt` instructions load the values in operand into the global descriptor table register or the interrupt descriptor table register respectively. `sgdt` and `sidt` store the contents of the global descriptor table register or the interrupt descriptor table register in the destination operand. The operand should be a 6 bytes in memory.

```
lgdt [ebx]       ; load global descriptor table
```

`lldt` loads the operand into the segment selector field of the local descriptor table register and `sldt` stores the segment selector from the local descriptor table register in the operand. `ltr` loads the operand into the segment selector field of the task register and `str` stores the segment selector from the task register in the operand. Rules for operand are the same as for the `lmsw` and `smsw` instructions.

`lar` loads the access rights from the segment descriptor specified by the selector in source operand into the destination operand and sets the ZF flag. The destination operand can be a 16-bit or 32-bit general register. The source operand should be a 16-bit general register or memory.

```
lar ax,[bx]      ; load access rights into word
lar eax,dx       ; load access rights into double word
```

lsl loads the segment limit from the segment descriptor specified by the selector in source operand into the destination operand and sets the ZF flag. Rules for operand are the same as for the **lar** instruction.

verr and **verw** verify whether the code or data segment specified with the operand is readable or writable from the current privilege level. The operand should be a word, it can be general register or memory. If the segment is accessible and readable (for **verr**) or writable (for **verw**) the ZF flag is set, otherwise it's cleared. Rules for operand are the same as for the **lldt** instruction.

arpl compares the RPL (requestor's privilege level) fields of two segment selectors. The first operand contains one segment selector and the second operand contains the other. If the RPL field of the destination operand is less than the RPL field of the source operand, the ZF flag is set and the RPL field of the destination operand is increased to match that of the source operand. Otherwise, the ZF flag is cleared and no change is made to the destination operand. The destination operand can be a word general register or memory, the source operand must be a general register.

```
arpl bx,ax      ; adjust RPL of selector in register
arpl [bx],ax    ; adjust RPL of selector in memory
```

clts clears the TS (task switched) flag in the **cr0** register. This instruction has no operands.

lock prefix causes the processor's bus-lock signal to be asserted during execution of the accompanying instruction. In a multiprocessor environment, the bus-lock signal insures that the processor has exclusive use of any shared memory while the signal is asserted. The **lock** prefix can be prepended only to the following instructions and only to those forms of the instructions where the destination operand is a memory operand: **add**, **adc**, **and**, **btc**, **btr**, **bts**, **cmpxchg**, **cmpxchg8b**, **dec**, **inc**, **neg**, **not**, **or**, **sbb**, **sub**, **xor**, **xadd** and **xchg**. If the **lock** prefix is used with one of these instructions and the source operand is a memory operand, an undefined opcode exception may be generated. An undefined opcode exception will also be generated if the **lock** prefix is used with any instruction not in the above list. The **xchg** instruction always asserts the bus-lock signal regardless of the presence or absence of the **lock** prefix.

hlt stops instruction execution and places the processor in a halted state. An enabled interrupt, a debug exception, the **BINIT**, **INIT** or the **RESET** signal will resume execution. This instruction has no operands.

invlpg invalidates (flushes) the TLB (translation lookaside buffer) entry specified with the operand, which should be a memory. The processor de-

termines the page that contains that address and flushes the TLB entry for that page.

rdmsr loads the contents of a 64-bit MSR (model specific register) of the address specified in the **ecx** register into registers **edx** and **eax**. **wrmsr** writes the contents of registers **edx** and **eax** into the 64-bit MSR of the address specified in the **ecx** register. **rdtsc** loads the current value of the processor's time stamp counter from the 64-bit MSR into the **edx** and **eax** registers. The processor increments the time stamp counter MSR every clock cycle and resets it to 0 whenever the processor is reset. **rdpmc** loads the contents of the 40-bit performance monitoring counter specified in the **ecx** register into registers **edx** and **eax**. These instructions have no operands.

wbinvd writes back all modified cache lines in the processor's internal cache to main memory and invalidates (flushes) the internal caches. The instruction then issues a special function bus cycle that directs external caches to also write back modified data and another bus cycle to indicate that the external caches should be invalidated. This instruction has no operands.

rsm return program control from the system management mode to the program that was interrupted when the processor received an SMM interrupt. This instruction has no operands.

sysenter executes a fast call to a level 0 system procedure, **sysexit** executes a fast return to level 3 user code. The addresses used by these instructions are stored in MSRs. These instructions have no operands.

2.1.13 FPU instructions

The FPU (Floating-Point Unit) instructions operate on the floating-point values in three formats: single precision (32-bit), double precision (64-bit) and double extended precision (80-bit). The FPU registers form the stack and each of them holds the double extended precision floating-point value. When some values are pushed onto the stack or are removed from the top, the FPU registers are shifted, so **st0** is always the value on the top of FPU stack, **st1** is the first value below the top, etc. The **st0** name has also the synonym **st**.

fld pushes the floating-point value onto the FPU register stack. The operand can be 32-bit, 64-bit or 80-bit memory location or the FPU register, it's value is then loaded onto the top of FPU register stack (the **st0** register) and is automatically converted into the double extended precision format.

```
fld dword [bx]    ; load single prevision value from memory
fld st2           ; push value of st2 onto register stack
```

`fld1`, `fldz`, `fldl2t`, `fldl2e`, `fldpi`, `fldlg2` and `fldln2` load the commonly used constants onto the FPU register stack. The loaded constants are +1.0, +0.0, $\log_2 10$, $\log_2 e$, π , $\log_{10} 2$ and $\ln 2$ respectively. These instructions have no operands.

`fild` convert the signed integer source operand into double extended precision floating-point format and pushes the result onto the FPU register stack. The source operand can be a 16-bit, 32-bit or 64-bit memory location.

```
fild qword [bx] ; load 64-bit integer from memory
```

`fst` copies the value of `st0` register to the destination operand, which can be 32-bit or 64-bit memory location or another FPU register. `fstp` performs the same operation as `fst` and then pops the register stack, getting rid of `st0`. `fstp` accepts the same operands as the `fst` instruction and can also store value in the 80-bit memory.

```
fst st3          ; copy value of st0 into st3 register
fstp tword [bx]  ; store value in memory and pop stack
```

`fist` converts the value in `st0` to a signed integer and stores the result in the destination operand. The operand can be 16-bit or 32-bit memory location. `fistp` performs the same operation and then pops the register stack, it accepts the same operands as the `fist` instruction and can also store integer value in the 64-bit memory, so it has the same rules for operands as `fild` instruction.

`fbld` converts the packed BCD integer into double extended precision floating-point format and pushes this value onto the FPU stack. `fbstp` converts the value in `st0` to an 18-digit packed BCD integer, stores the result in the destination operand, and pops the register stack. The operand should be an 80-bit memory location.

`fadd` adds the destination and source operand and stores the sum in the destination location. The destination operand is always an FPU register, if the source is a memory location, the destination is `st0` register and only source operand should be specified. If both operands are FPU registers, at least one of them should be `st0` register. An operand in memory can be a 32-bit or 64-bit value.

```
fadd qword [bx] ; add double precision value to st0
fadd st2,st0    ; add st0 to st2
```

`faddp` adds the destination and source operand, stores the sum in the destination location and then pops the register stack. The destination operand must be an FPU register and the source operand must be the `st0`. When no operands are specified, `st1` is used as a destination operand.

```

faddp          ; add st0 to st1 and pop the stack
faddp st2,st0  ; add st0 to st2 and pop the stack

```

`fiadd` instruction converts an integer source operand into double extended precision floating-point value and adds it to the destination operand. The operand should be a 16-bit or 32-bit memory location.

```

fiadd word [bx] ; add word integer to st0

```

`fsub`, `fsubr`, `fmul`, `fdiv`, `fdivr` instruction are similar to `fadd`, have the same rules for operands and differ only in the performed computation. `fsub` subtracts the source operand from the destination operand, `fsubr` subtract the destination operand from the source operand, `fmul` multiplies the destination and source operands, `fdiv` divides the destination operand by the source operand and `fdivr` divides the source operand by the destination operand. `fsubp`, `fsubrp`, `fmulp`, `fdivp`, `fdivrp` perform the same operations and pop the register stack, the rules for operand are the same as for the `faddp` instruction. `fisub`, `fisubr`, `fimul`, `fidiv`, `fidivr` perform these operations after converting the integer source operand into floating-point value, they have the same rules for operands as `fiadd` instruction.

`fsqrt` computes the square root of the value in `st0` register, `fsin` computes the sine of that value, `fcos` computes the cosine of that value, `fchs` complements its sign bit, `fabs` clears its sign to create the absolute value, `frndint` rounds it to the nearest integral value, depending on the current rounding mode. `f2xm1` computes the exponential value of 2 to the power of `st0` and subtracts the 1.0 from it, the value of `st0` must lie in the range -1.0 to $+1.0$. All these instruction store the result in `st0` and have no operands.

`fsincos` computes both the sine and the cosine of the value in `st0` register, stores the sine in `st0` and pushes the cosine on the top of FPU register stack. `fptan` computes the tangent of the value in `st0`, stores the result in `st0` and pushes a 1.0 onto the FPU register stack. `fpatan` computes the arctangent of the value in `st1` divided by the value in `st0`, stores the result in `st1` and pops the FPU register stack. `fyl2x` computes the binary logarithm of `st0`, multiplies it by `st1`, stores the result in `st1` and pop the FPU register stack; `fyl2xp1` performs the same operation but it adds 1.0 to `st0` before computing the logarithm. `fprem` computes the remainder obtained from dividing the value in `st0` by the value in `st1`, and stores the result in `st0`. `fprem1` performs the same operation as `fprem`, but it computes the remainder in the way specified by IEEE Standard 754. `fscale` truncates the value in `st1` and increases the exponent of `st0` by this value. `fxtract` separates the value in `st0` into its exponent and significand, stores the exponent

in `st0` and pushes the significand onto the register stack. `fnop` performs no operation. These instructions have no operands.

`fxch` exchanges the contents of `st0` and another FPU register. The operand should be an FPU register, if no operand is specified, the contents of `st0` and `st1` are exchanged.

`fcom` and `fcomp` compare the contents of `st0` and the source operand and set flags in the FPU status word according to the results. `fcomp` additionally pops the register stack after performing the comparison. The operand can be a single or double precision value in memory or the FPU register. When no operand is specified, `st1` is used as a source operand.

```
fcom           ; compare st0 with st1
fcomp st2      ; compare st0 with st2 and pop stack
```

`fcompp` compares the contents of `st0` and `st1`, sets flags in the FPU status word according to the results and pops the register stack twice. This instruction has no operands.

`fucom`, `fucomp` and `fucompp` performs an unordered comparison of two FPU registers. Rules for operands are the same as for the `fcom`, `fcomp` and `fcompp`, but the source operand must be an FPU register.

`ficom` and `ficomp` compare the value in `st0` with an integer source operand and set the flags in the FPU status word according to the results. `ficomp` additionally pops the register stack after performing the comparison. The integer value is converted to double extended precision floating-point format before the comparison is made. The operand should be a 16-bit or 32-bit memory location.

```
ficom word [bx] ; compare st0 with 16-bit integer
```

`fcomi`, `fcomip`, `fucomi`, `fucomip` perform the comparison of `st0` with another FPU register and set the ZF, PF and CF flags according to the results. `fcomip` and `fucomip` additionally pop the register stack after performing the comparison. The instructions obtained by attaching the FPU condition mnemonic (see table 2.2) to the `fcmov` mnemonic transfer the specified FPU register into `st0` register if the given test condition is true. These instructions allow two different syntaxes, one with single operand specifying the source FPU register, and one with two operands, in that case destination operand should be `st0` register and the second operand specifies the source FPU register.

```
fcomi st2      ; compare st0 with st2 and set flags
fcmovb st0,st2 ; transfer st2 to st0 if below
```

Mnemonic	Condition tested	Description
b	$CF = 1$	below
e	$ZF = 1$	equal
be	$CF \text{ or } ZF = 1$	below or equal
u	$PF = 1$	unordered
nb	$CF = 0$	not below
ne	$ZF = 0$	not equal
nbe	$CF \text{ and } ZF = 0$	not below nor equal
nu	$PF = 0$	not unordered

Table 2.2: FPU conditions.

fstst compares the value in **st0** with 0.0 and sets the flags in the FPU status word according to the results. **fxam** examines the contents of the **st0** and sets the flags in FPU status word to indicate the class of value in the register. These instructions have no operands.

fstsw and **fnstsw** store the current value of the FPU status word in the destination location. The destination operand can be either a 16-bit memory or the **ax** register. **fstsw** checks for pending unmasked FPU exceptions before storing the status word, **fnstsw** does not.

fstcw and **fnstcw** store the current value of the FPU control word at the specified destination in memory. **fstcw** checks for pending unmasked FPU exceptions before storing the control word, **fnstcw** does not. **fldcw** loads the operand into the FPU control word. The operand should be a 16-bit memory location.

fstenv and **fnstenv** store the current FPU operating environment at the memory location specified with the destination operand, and then mask all FPU exceptions. **fstenv** checks for pending unmasked FPU exceptions before proceeding, **fnstenv** does not. **fldenv** loads the complete operating environment from memory into the FPU. **fsave** and **fnsave** store the current FPU state (operating environment and register stack) at the specified destination in memory and reinitializes the FPU. **fsave** check for pending unmasked FPU exceptions before proceeding, **fnsave** does not. **frstor** loads the FPU state from the specified memory location. All these instructions need an operand being a memory location.

finit and **fninit** set the FPU operating environment into its default state. **finit** checks for pending unmasked FPU exception before proceeding, **fninit** does not. **fclex** and **fnclx** clear the FPU exception flags in the FPU status word. **fclex** checks for pending unmasked FPU exception before

proceeding, `fnclex` does not. `wait` and `fwait` are synonyms for the same instruction, which causes the processor to check for pending unmasked FPU exceptions and handle them before proceeding. These instructions have no operands.

`ffree` sets the tag associated with specified FPU register to empty. The operand should be an FPU register.

`fincstp` and `fdecstp` rotate the FPU stack by one by adding or subtracting one to the pointer of the top of stack. These instructions have no operands.

2.1.14 MMX instructions

The MMX instructions operate on the packed integer types and use the MMX registers, which are the low 64-bit parts of the 80-bit FPU registers. Because of this MMX instructions cannot be used at the same time as FPU instructions. They can operate on packed bytes (eight 8-bit integers), packed words (four 16-bit integers) or packed double words (two 32-bit integers), use of packed formats allows to perform operations on multiple data at one time.

`movq` copies a quad word from the source operand to the destination operand. At least one of the operands must be a MMX register, the second one can be also a MMX register or 64-bit memory location.

```
movq mm0,mm1      ; move quad word from register to register
movq mm2,[ebx]    ; move quad word from memory to register
```

`movd` copies a double word from the source operand to the destination operand. One of the operands must be a MMX register, the second one can be a general register or 32-bit memory location. Only low double word of MMX register is used.

All general MMX operations have two operands, the destination operand should be a MMX register, the source operand can be a MMX register or 64-bit memory location. Operation is performed on the corresponding data elements of the source and destination operand and stored in the data elements of the destination operand. `paddb`, `paddw` and `paddq` perform the addition of packed bytes, packed words, or packed double words. `psubb`, `psubw` and `psubq` perform the subtraction of appropriate types. `paddsb`, `paddsw`, `psubsb` and `psubsw` perform the addition or subtraction of packed bytes or packed words with the signed saturation. `paddusb`, `paddusw`, `psubusb`, `psubusw` are analogous, but with unsigned saturation. `pmulhw` and `pmullw` performs a signed multiply of the packed words and store the high or low words of the results in the destination operand. `pmaddwd` performs a multiply

of the packed words and adds the four intermediate double word products in pairs to produce result as a packed double words. **pand**, **por** and **pxor** perform the logical operations on the quad words, **pandn** performs also a logical negation of the destination operand before the operation. **pcmpeqb**, **pcmpeqw** and **pcmpeqd** compare for equality of packed bytes, packed words or packed double words. If a pair of data elements is equal, the corresponding data element in the destination operand is filled with bits of value 1, otherwise it's set to 0. **pcmpgtb**, **pcmpgtw** and **pcmpgtd** perform the similar operation, but they check whether the data elements in the destination operand are greater than the corresponding data elements in the source operand. **packsswb** converts packed signed words into packed signed bytes, **packssdw** converts packed signed double words into packed signed words, using saturation to handle overflow conditions. **packuswb** converts packed signed words into packed unsigned bytes. Converted data elements from the source operand are stored in the low part of the destination operand, while converted data elements from the destination operand are stored in the high part. **punpckhbw**, **punpckhwd** and **punpckhdq** interleaves the data elements from the high parts of the source and destination operands and stores the result into the destination operand. **punpcklbw**, **punpcklwd** and **punpckldq** perform the same operation, but the low parts of the source and destination operand are used.

```
paddsb mm0,[esi] ; add packed bytes with signed saturation
pcmpeqw mm3,mm7 ; compare packed words for equality
```

psllw, **pslld** and **psllq** perform logical shift left of the packed words, packed double words or a single quad word in the destination operand by the amount specified in the source operand. **psrlw**, **psrld** and **psrlq** perform logical shift right of the packed words, packed double words or a single quad word. **psraw** and **psrad** perform arithmetic shift of the packed words or double words. The destination operand should be a MMX register, while source operand can be a MMX register, 64-bit memory location, or 8-bit immediate value.

```
psllw mm2,mm4    ; shift words left logically
psrad mm4,[ebx]  ; shift double words right arithmetically
```

emms makes the FPU registers usable for the FPU instructions, it must be used before using the FPU instructions if any MMX instructions were used.

2.1.15 SSE instructions

The SSE extension adds more MMX instructions and also introduces the operations on packed single precision floating point values. The 128-bit

packed single precision format consists of four single precision floating point values. The 128-bit SSE registers are designed for the purpose of operations on this data type.

movaps and **movups** transfer a double quad word operand containing packed single precision values from source operand to destination operand. At least one of the operands have to be a SSE register, the second one can be also a SSE register or 128-bit memory location. Memory operands for **movaps** instruction must be aligned on boundary of 16 bytes, operands for **movups** instruction don't have to be aligned.

```
movups xmm0,[ebx]    ; move unaligned double quad word
```

movlps moves packed two single precision values between the memory and the low quad word of SSE register. **movhps** moved packed two single precision values between the memory and the high quad word of SSE register. One of the operands must be a SSE register, and the other operand must be a 64-bit memory location.

```
movlps xmm0,[ebx]    ; move memory to low quad word of xmm0
movhps [esi],xmm7     ; move high quad word of xmm7 to memory
```

movlhps moves packed two single precision values from the low quad word of source register to the high quad word of destination register. **movhlps** moves two packed single precision values from the high quad word of source register to the low quad word of destination register. Both operands have to be a SSE registers.

movmskps transfers the most significant bit of each of the four single precision values in the SSE register into low four bits of a general register. The source operand must be a SSE register, the destination operand must be a general register.

movss transfers a single precision value between source and destination operand (only the low double word is transferred). At least one of the operands have to be a SSE register, the second one can be also a SSE register or 32-bit memory location.

```
movss [edi],xmm3     ; move low double word of xmm3 to memory
```

Each of the SSE arithmetic operations has two variants. When the mnemonic ends with **ps**, the source operand can be a 128-bit memory location or a SSE register, the destination operand must be a SSE register and the operation is performed on packed four single precision values, for each pair of the corresponding data elements separately, the result is stored in the destination register. When the mnemonic ends with **ss**, the source operand

can be a 32-bit memory location or a SSE register, the destination operand must be a SSE register and the operation is performed on single precision values, only low double words of SSE registers are used in this case, the result is stored in the low double word of destination register. **addps** and **addss** add the values, **subps** and **subss** subtract the source value from destination value, **mulps** and **mulss** multiply the values, **divps** and **divss** divide the destination value by the source value, **rcpps** and **rcpss** compute the approximate reciprocal of the source value, **sqrtps** and **sqrtss** compute the square root of the source value, **rsqrtps** and **rsqrtss** compute the approximate reciprocal of square root of the source value, **maxps** and **maxss** compare the source and destination values and return the greater one, **minps** and **minss** compare the source and destination values and return the lesser one.

```
mulss xmm0,[ebx]    ; multiply single precision values
addps xmm3,xmm7     ; add packed single precision values
```

andps, **andnps**, **orps** and **xorps** perform the logical operations on packed single precision values. The source operand can be a 128-bit memory location or a SSE register, the destination operand must be a SSE register.

cmppps compares packed single precision values and returns a mask result into the destination operand, which must be a SSE register. The source operand can be a 128-bit memory location or SSE register, the third operand must be an immediate operand selecting code of one of the eight compare conditions (table 2.3). **cmpss** performs the same operation on single precision values, only low double word of destination register is affected, in this case source operand can be a 32-bit memory location or SSE register. These two instructions have also variants with only two operands and the condition encoded within mnemonic. Their mnemonics are obtained by attaching the mnemonic from table 2.3 to the **cmp** mnemonic and then attaching the **ps** or **ss** at the end.

```
cmppps xmm2,xmm4,0  ; compare packed single precision values
cmpltss xmm0,[ebx]  ; compare single precision values
```

comiss and **ucomiss** compare the single precision values and set the ZF, PF and CF flags to show the result. The destination operand must be a SSE register, the source operand can be a 32-bit memory location or SSE register.

shufps moves any two of the four single precision values from the destination operand into the low quad word of the destination operand, and any two of the four values from the source operand into the high quad word of the destination operand. The destination operand must be a SSE register,

Code	Mnemonic	Description
0	eq	equal
1	lt	less than
2	le	less than or equal
3	unord	unordered
4	neq	not equal
5	nlt	not less than
6	nle	not less than nor equal
7	ord	ordered

Table 2.3: SSE conditions.

the source operand can be a 128-bit memory location or SSE register, the third operand must be an 8-bit immediate value selecting which values will be moved into the destination operand. Bits 0 and 1 select the value to be moved from destination operand to the low double word of the result, bits 2 and 3 select the value to be moved from the destination operand to the second double word, bits 4 and 5 select the value to be moved from the source operand to the third double word, and bits 6 and 7 select the value to be moved from the source operand to the high double word of the result.

```
shufps xmm0,xmm0,10010011b ; shuffle double words
```

unpckhps performs an interleaved unpack of the values from the high parts of the source and destination operands and stores the result in the destination operand, which must be a SSE register. The source operand can be a 128-bit memory location or a SSE register. **unpcklps** performs an interleaved unpack of the values from the low parts of the source and destination operand and stores the result in the destination operand, the rules for operands are the same.

cvtpi2ps converts packed two double word integers into the the packed two single precision floating point values and stores the result in the low quad word of the destination operand, which should be a SSE register. The source operand can be a 64-bit memory location or MMX register.

```
cvtpi2ps xmm0,mm0 ; integers to single precision values
```

cvtsi2ss converts a double word integer into a single precision floating point value and stores the result in the low double word of the destination operand, which should be a SSE register. The source operand can be a 32-bit memory location or 32-bit general register.

`cvtsi2ss xmm0,eax ; integer to single precision value`

`cvtps2pi` converts packed two single precision floating point values into packed two double word integers and stores the result in the destination operand, which should be a MMX register. The source operand can be a 64-bit memory location or SSE register, only low quad word of SSE register is used. `cvttps2pi` performs the similar operation, except that truncation is used to round a source values to integers, rules for the operands are the same.

`cvtps2pi mm0,xmm0 ; single precision values to integers`

`cvtss2si` convert a single precision floating point value into a double word integer and stores the result in the destination operand, which should be a 32-bit general register. The source operand can be a 32-bit memory location or SSE register, only low double word of SSE register is used. `cvttss2si` performs the similar operation, except that truncation is used to round a source value to integer, rules for the operands are the same.

`cvtss2si eax,xmm0 ; single precision value to integer`

`pextrw` copies the word in the source operand specified by the third operand to the destination operand. The source operand must be a MMX register, the destination operand must be a 32-bit general register (but only the low word of it is affected), the third operand must an 8-bit immediate value.

`pextrw eax,mm0,1 ; extract word into eax`

`pinsrw` inserts a word from the source operand in the destination operand at the location specified with the third operand, which must be an 8-bit immediate value. The destination operand must be a MMX register, the source operand can be a 16-bit memory location or 32-bit general register (only low word of the register is used).

`pinsrw mm1,ebx,2 ; insert word from ebx`

`pavgb` and `pavgw` compute average of packed bytes or words. `pmaxub` return the maximum values of packed unsigned bytes, `pminub` returns the minimum values of packed unsigned bytes, `pmaxsw` returns the maximum values of packed signed words, `pminsw` returns the minimum values of packed signed words. `pmulhuw` performs a unsigned multiply of the packed words and stores the high words of the results in the destination operand. `psadbw` computes the absolute differences of packed unsigned bytes, sums the differences, and

stores the sum in the low word of destination operand. All these instructions follow the same rules for operands as the general MMX operations described in previous section.

pmovmskb creates a mask made of the most significant bit of each byte in the source operand and stores the result in the low byte of destination operand. The source operand must be a MMX register, the destination operand must be a 32-bit general register.

pshufw inserts words from the source operand in the destination operand from the locations specified with the third operand. The destination operand must be a MMX register, the source operand can be a 64-bit memory location or MMX register, third operand must be an 8-bit immediate value selecting which values will be moved into destination operand, in the similar way as the third operand of the **shufps** instruction.

movntq moves the quad word from the source operand to memory using a non-temporal hint to minimize cache pollution. The source operand should be a MMX register, the destination operand should be a 64-bit memory location. **movntps** stores packed single precision values from the SSE register to memory using a non-temporal hint. The source operand should be a SSE register, the destination operand should be a 128-bit memory location. **maskmovq** stores selected bytes from the first operand into a 64-bit memory location using a non-temporal hint. Both operands should be MMX registers, the second operand selects which bytes from the source operand are written to memory. The memory location is pointed to by DI (or EDI) register in the segment selected by DS.

prefetcht0, **prefetcht1**, **prefetcht2** and **prefetchnta** fetch the line of data from memory that contains byte specified with the operand to a specified location in hierarchy. The operand should be an 8-bit memory location.

sfence performs a serializing operation on all instructions stored to memory that were issued prior to it. This instruction has no operands.

ldmxcsr loads the 32-bit memory operand into the MXCSR register. **stmxcsr** stores the contents of MXCSR into a 32-bit memory operand.

fxsave saves the current state of the FPU, MXCSR register, and all the FPU and SSE registers to a 512-byte memory location specified in the destination operand. **fxrstor** reloads data previously stored with **fxsave** instruction from the specified 512-byte memory location. The memory operand for both those instructions must be aligned on 16-byte boundary, it should declare operand of no specified size.

2.1.16 SSE2 instructions

The SSE2 extension introduces the operations on packed double precision floating point values, extends the syntax of MMX instructions, and adds also some new instructions.

`movapd` and `movupd` transfer a double quad word operand containing packed double precision values from source operand to destination operand. These instructions are analogous to `movaps` and `movups` and have the same rules for operands.

`movlpd` moves double precision value between the memory and the low quad word of SSE register. `movhpd` moved double precision value between the memory and the high quad word of SSE register. These instructions are analogous to `movlps` and `movhps` and have the same rules for operands.

`movmskpd` transfers the most significant bit of each of the two double precision values in the SSE register into low two bits of a general register. This instruction is analogous to `movmskps` and has the same rules for operands.

`movsd` transfers a double precision value between source and destination operand (only the low quad word is transferred). At least one of the operands have to be a SSE register, the second one can be also a SSE register or 64-bit memory location.

Arithmetic operations on double precision values are: `addpd`, `addsd`, `subpd`, `subsd`, `mulpd`, `mulsd`, `divpd`, `divsd`, `sqrtpd`, `sqrtsd`, `maxpd`, `maxsd`, `minpd`, `minsd`, and they are analogous to arithmetic operations on single precision values described in previous section. When the mnemonic ends with `pd` instead of `ps`, the operation is performed on packed two double precision values, but rules for operands are the same. When the mnemonic ends with `sd` instead of `ss`, the source operand can be a 64-bit memory location or a SSE register, the destination operand must be a SSE register and the operation is performed on double precision values, only low quad words of SSE registers are used in this case.

`andpd`, `andnpd`, `orpd` and `xorpd` perform the logical operations on packed double precision values. They are analogous to SSE logical operations on single precision values and have the same rules for operands.

`cmppd` compares packed double precision values and returns and returns a mask result into the destination operand. This instruction is analogous to `cmpps` and has the same rules for operands. `cmpsd` performs the same operation on double precision values, only low quad word of destination register is affected, in this case source operand can be a 64-bit memory or SSE register. Variant with only two operands are obtained by attaching the condition mnemonic from table 2.3 to the `cmp` mnemonic and then attaching the `pd` or `sd` at the end.

comisd and **ucomisd** compare the double precision values and set the ZF, PF and CF flags to show the result. The destination operand must be a SSE register, the source operand can be a 128-bit memory location or SSE register.

shufpd moves any of the two double precision values from the destination operand into the low quad word of the destination operand, and any of the two values from the source operand into the high quad word of the destination operand. This instruction is analogous to **shufps** and has the same rules for operand. Bit 0 of the third operand selects the value to be moved from the destination operand, bit 1 selects the value to be moved from the source operand, the rest of bits are reserved and must be zeroed.

unpckhpd performs an unpack of the high quad words from the source and destination operands, **unpcklpd** performs an unpack of the low quad words from the source and destination operands. They are analogous to **unpckhps** and **unpcklps**, and have the same rules for operands.

cvtps2pd converts the packed two single precision floating point values to two packed double precision floating point values, the destination operand must be a SSE register, the source operand can be a 64-bit memory location or SSE register. **cvtpd2ps** converts the packed two double precision floating point values to packed two single precision floating point values, the destination operand must be a SSE register, the source operand can be a 128-bit memory location or SSE register. **cvtss2sd** converts the single precision floating point value to double precision floating point value, the destination operand must be a SSE register, the source operand can be a 32-bit memory location or SSE register. **cvtsd2ss** converts the double precision floating point value to single precision floating point value, the destination operand must be a SSE register, the source operand can be 64-bit memory location or SSE register.

cvtpi2pd converts packed two double word integers into the the packed double precision floating point values, the destination operand must be a SSE register, the source operand can be a 64-bit memory location or MMX register. **cvtsi2sd** converts a double word integer into a double precision floating point value, the destination operand must be a SSE register, the source operand can be a 32-bit memory location or 32-bit general register. **cvtpd2pi** converts packed double precision floating point values into packed two double word integers, the destination operand should be a MMX register, the source operand can be a 128-bit memory location or SSE register. **cvttpd2pi** performs the similar operation, except that truncation is used to round a source values to integers, rules for operands are the same. **cvtsd2si** converts a double precision floating point value into a double word integer, the destination operand should be a 32-bit general register, the source operand

can be a 64-bit memory location or SSE register. `cvttssd2si` performs the similar operation, except that truncation is used to round a source value to integer, rules for operands are the same.

`cvtps2dq` and `cvttps2dq` convert packed single precision floating point values to packed four double word integers, storing them in the destination operand. `cvtpd2dq` and `cvttpd2dq` convert packed double precision floating point values to packed two double word integers, storing the result in the low quad word of the destination operand. `cvtdq2ps` converts packed four double word integers to packed single precision floating point values. `cvtdq2pd` converts packed two double word integers from the low quad word of the source operand to packed double precision floating point values. For all these instruction destination operand must be a SSE register, the source operand can be a 128-bit memory location or SSE register.

`movdqa` and `movdqu` transfer a double quad word operand containing packed integers from source operand to destination operand. At least one of the operands have to be a SSE register, the second one can be also a SSE register or 128-bit memory location. Memory operands for `movdqa` instruction must be aligned on boundary of 16 bytes, operands for `movdqu` instruction don't have to be aligned.

`movq2dq` moves the contents of the MMX source register to the low quad word of destination SSE register. `movdq2q` moves the low quad word from the source SSE register to the destination MMX register.

```
movq2dq xmm0,mm1    ; move from MMX register to SSE register
movdq2q mm0,xmm1    ; move from SSE register to MMX register
```

All MMX instructions operating on the 64-bit packed integers (those with mnemonics starting with `p\verb`) are extended to operate on 128-bit packed integers located in SSE registers. Additional syntax for these instructions needs an SSE register where MMX register was needed, and the 128-bit memory location or SSE register where 64-bit memory location of MMX register were needed. The exception is `pshufw` instruction, which doesn't allow extended syntax, but has two new variants: `pshufhw` and `pshuflw`, which allow only the extended syntax, and perform the same operation as `pshufw` on the high or low quad words of operands respectively. Also the new instruction `pshufd` is introduced, which performs the same operation as `pshufw`, but on the double words instead of words, it allows only the extended syntax.

```
psubb xmm0,[esi]    ; subtract 16 packed bytes
pextrw eax,xmm0,7    ; extract highest word into eax
```

paddq performs the addition of packed quad words, **psubq** performs the subtraction of packed quad words, **pmuludq** performs an unsigned multiply of low double words from each corresponding quad words and returns the results in packed quad words. These instructions follow the same rules for operands as the general MMX operations described in 2.1.14.

pslldq and **psrldq** perform logical shift left or right of the double quad word in the destination operand by the amount of bits specified in the source operand. The destination operand should be a SSE register, source operand should be an 8-bit immediate value.

punpckhqdq interleaves the high quad word of the source operand and the high quad word of the destination operand and writes them to the destination SSE register. **punpcklqdq** interleaves the low quad word of the source operand and the low quad word of the destination operand and writes them to the destination SSE register. The source operand can be a 128-bit memory location or SSE register.

movntdq stores packed integer data from the SSE register to memory using non-temporal hint. The source operand should be a SSE register, the destination operand should be a 128-bit memory location. **movntpd** stores packed double precision values from the SSE register to memory using a non-temporal hint. Rules for operand are the same. **movnti** stores integer from a general register to memory using a non-temporal hint. The source operand should be a 32-bit general register, the destination operand should be a 32-bit memory location. **maskmovdqu** stores selected bytes from the first operand into a 128-bit memory location using a non-temporal hint. Both operands should be a SSE registers, the second operand selects which bytes from the source operand are written to memory. The memory location is pointed by DI (or EDI) register in the segment selected by DS and does not need to be aligned.

clflush writes and invalidates the cache line associated with the address of byte specified with the operand, which should be a 8-bit memory location.

lfence performs a serializing operation on all instruction loading from memory that were issued prior to it. **mfence** performs a serializing operation on all instruction accessing memory that were issued prior to it, and so it combines the functions of **sfence** (described in previous section) and **lfence** instructions. These instructions have no operands.

2.1.17 SSE3 instructions

Prescott technology introduces some new instructions to improve the performance of SSE and SSE2 – this extension is called SSE3.

fisttp behaves like the **fistp** instruction and accepts the same operands,

the only difference is that it always used truncation, irrespective of the rounding mode.

movshdup loads into destination operand the 128-bit value obtained from the source value of the same size by filling the each quad word with the two duplicates of the value in its high double word. **movsldup** performs the same action, except it duplicates the values of low double words. The destination operand should be SSE register, the source operand can be SSE register or 128-bit memory location.

movddup loads the 64-bit source value and duplicates it into high and low quad word of the destination operand. The destination operand should be SSE register, the source operand can be SSE register or 64-bit memory location.

lddqu is functionally equivalent to **movdqu** instruction with memory as source operand, but it may improve performance when the source operand crosses a cacheline boundary. The destination operand has to be SSE register, the source operand must be 128-bit memory location.

addsubps performs single precision addition of second and fourth pairs and single precision substraction of the first and third pairs of floating point values in the operands. **addsubpd** performs double precision addition of the second pair and double precision substraction of the first pair of floating point values in the operand. **haddps** performs the addition of two single precision values within the each quad word of source and destination operands, and stores the results of such horizontal addition of values from destination operand into low quad word of destination operand, and the results from the source operand into high quad word of destination operand. **haddpd** performs the addition of two double precision values within each operand, and stores the result from destination operand into low quad word of destination operand, and the result from source operand into high quad word of destination operand. All these instruction need the destination operand to be SSE register, source operand can be SSE register or 128-bit memory location.

monitor sets up an address range for monitoring of write-back stores. It need its three operands to be EAX, ECX and EDX register in that order. **mwait** waits for a write-back store to the address range set up by the **monitor** instruction. It uses two operands with additional parameters, first being the EAX and second the ECX register.

2.1.18 AMD 3DNow! instructions

The 3DNow! extension adds a new MMX instructions to those described in 2.1.14, and introduces operation on the 64-bit packed floating point values, each consisting of two single precision floating point values.

These instructions follow the same rules as the general MMX operations, the destination operand should be a MMX register, the source operand can be a MMX register or 64-bit memory location. **pavgusb** computes the rounded averages of packed unsigned bytes. **pmulhrw** performs a signed multiply of the packed words, round the high word of each double word results and stores them in the destination operand. **pi2fd** converts packed double word integers into packed floating point values. **pf2id** converts packed floating point values into packed double word integers using truncation. **pi2fw** converts packed word integers into packed floating point values, only low words of each double word in source operand are used. **pf2iw** converts packed floating point values to packed word integers, results are extended to double words using the sign extension. **pfadd** adds packed floating point values. **pfsb** and **pfsbr** subtracts packed floating point values, the first one subtracts source values from destination values, the second one subtracts destination values from the source values. **pfmul** multiplies packed floating point values. **pfacc** adds the low and high floating point values of the destination operand, storing the result in the low double word of destination, and adds the low and high floating point values of the source operand, storing the result in the high double word of destination. **pfnacc** subtracts the high floating point value of the destination operand from the low, storing the result in the low double word of destination, and subtracts the high floating point value of the source operand from the low, storing the result in the high double word of destination. **pfpnacc** subtracts the high floating point value of the destination operand from the low, storing the result in the low double word of destination, and adds the low and high floating point values of the source operand, storing the result in the high double word of destination. **pfmax** and **pfmin** compute the maximum and minimum of floating point values. **pswapd** reverses the high and low double word of the source operand. **pfrcp** returns an estimates of the reciprocals of floating point values from the source operand, **pfrrsqrt** returns an estimates of the reciprocal square roots of floating point values from the source operand, **pfrcpit1** performs the first step in the Newton–Raphson iteration to refine the reciprocal approximation produced by **pfrcp** instruction, **pfrrsqit1** performs the first step in the Newton–Raphson iteration to refine the reciprocal square root approximation produced by **pfrrsqrt** instruction, **pfrcpit2** performs the second final step in the Newton–Raphson iteration to refine the reciprocal approximation or the reciprocal square root approximation. **pfcmpeq**, **pfcmpge** and **pfcmpgt** compare the packed floating point values and sets all bits or zeroes all bits of the corresponding data element in the destination operand according to the result of comparison, first checks whether values are equal, second checks whether destination value is greater or equal to source value, third checks

whether destination value is greater than source value.

prefetch and **prefetchw** load the line of data from memory that contains byte specified with the operand into the data cache, **prefetchw** instruction should be used when the data in the cache line is expected to be modified, otherwise the **prefetch** instruction should be used. The operand should be an 8-bit memory location.

femms performs a fast clear of MMX state. This instruction has no operands.

2.2 Control directives

This section describes the directives that control the assembly process, they are processed during the assembly and may cause some blocks of instructions to be assembled differently or not assembled at all.

2.2.1 Repeating blocks of instructions

times directive repeats one instruction specified number of times. It should be followed by numerical expression specifying number of repeats and the instruction to repeat (optionally colon can be used to separate number and instruction). When special symbol **%** is used inside the instruction, it is equal to the number of current repeat. For example **times 5 db %** will define five bytes with values 1, 2, 3, 4, 5. Recursive use of **times** directive is also allowed, so **times 3 times % db %** will define six bytes with values 1, 1, 2, 1, 2, 3.

repeat directive repeats the whole block of instructions. It should be followed by numerical expression specifying number of repeats. Instructions to repeat are expected in next lines, ended with the **end repeat** directive, for example:

```
repeat 8
    mov byte [bx],%
    inc bx
end repeat
```

The generated code will store byte values from one to eight in the memory addressed by **bx** register.

Number of repeats can be zero, in that case the instructions are not assembled at all.

2.2.2 Conditional assembly

if directive causes some block of instructions to be assembled only under certain condition. It should be followed by logical expression specifying the condition, instructions in next lines will be assembled only when this condition is met, otherwise they will be skipped. The optional **else if** directive followed with logical expression specifying additional condition begins the next block of instructions that will be assembled if previous conditions were not met, and the additional condition is met. The optional **else** directive begins the block of instructions that will be assembled if all the conditions were not met. The **end if** directive ends the last block of instructions.

You should note that **if** directive is processed at assembly stage and therefore it doesn't affect any preprocessor directives, so if you put some symbolic constants or macroinstructions inside such block, they will get defined even when the condition is not met.

The logical expression consist of logical values and logical operators. The logical operators are **~** for logical negation, **&** for logical and, **|** for logical or. The negation has the highest priority. Logical value can be a numerical expression, it will be false if it is equal to zero, otherwise it will be true. Two numerical expression can be compared using one of the following operators to make the logical value: **=** (equal), **<** (less), **>** (greater), **<=** (less or equal), **>=** (greater or equal), **<>** (not equal).

There are also operators that allow comparison of values being any chains of symbols. The **eq** compares two such values whether they are exactly the same. The **in** operator checks whether given value is a member of the list of values following this operator, the list should be enclosed between **<** and **>** characters, its members should be separated with commas. The **eqtype** operator checks whether the two compared values have the same structure, and whether the structural elements are of the same type. The distinguished types include numerical expressions, individual quoted strings, floating point numbers, address expressions (the expressions enclosed in square brackets or preceded by **ptr** operator), instruction mnemonics, registers, size operators, jump type and code type operators. For example, two values, each one consisting of register name followed by comma and numerical expression, will be regarded as of the same type, no matter what kind of register and how complicated numerical expression is used (with exception for quoted strings and floating point values, which are special kinds of numerical expressions and are treated as different types).

The **used** operator should be followed by a symbol name, it checks whether the given symbol is used somewhere (it returns correct result even if symbol is used only after this check). The **defined** operator can be followed by any

expression, usually just by a single symbol name; it checks whether the given expression contains only symbols that are defined in the source and accessible from the current position.

The following simple example uses the `count` constant that should be defined somewhere in source:

```
if count>0
    mov cx,count
    rep movsb
end if
```

These two assembly instructions will be assembled only if the `count` constant is greater than 0.

The next example is more complex and assumes that the symbolic constant `reg` has been defined:

```
if reg in <cs,ds,es,fs,gs,ss>
    mov dx,reg
    add ax,dx
    shl ax,1
else if reg eq ax
    shl ax,2
else
    add ax,reg
    shl ax,1
end if
```

The first block of instructions will be assembled only if the value of `reg` is segment register, otherwise the second or third block will be assembled whether the value of `reg` is `ax` register or not.

2.2.3 Other directives

`align` directive aligns code or data to the specified boundary. It should be followed by a numerical expression specifying the number of bytes, to the multiply of which the current address has to be aligned. The boundary value has to be the power of two.

`virtual` defines virtual data at specified address. This data won't be included in the output file, but labels defined there can be used in other parts of source. This directive can be followed by `at` operator and the numerical expression specifying the address for virtual data, otherwise it uses current address, the same as `virtual at $`. Instructions defining data are expected

in next lines, ended with `end virtual` directive. This directive can be used to create union of some variables, for example:

```
GDTR dp ?
virtual at GDTR
    GDT_limit dw ?
    GDT_address dd ?
end virtual
```

It defines two labels for parts of the 48-bit variable at `GDTR` address.

It can be also used to define labels for some structures addressed by a register, for example:

```
virtual at bx
    LDT_limit dw ?
    LDT_address dd ?
end virtual
```

With such definition instruction `mov ax,[LDT_limit]` will be assembled to `mov ax,[bx]`.

Declaring defined data values or instructions inside the virtual block would also be useful, because the `load` directive (already described in section 1.2.3) can be used to load the values from the virtually generated code into a constants. This directive should be used after the code it loads but before the virtual block ends, because it can only load the values from the same code space. For example:

```
virtual at 0
    xor eax,eax
    and edx,eax
    load zeroq dword from 0
end virtual
```

The above piece of code will define the `zeroq` constant containing four bytes of the machine code of the instructions defined inside the virtual block. This method can be also used to load some binary value from external file. For example this code:

```
virtual at 0
    file 'a.txt':10h,1
    load char from 0
end virtual
```

loads the single byte from offset 10h in file `a.txt` into the `char` constant.

`display` directive displays the message at the assembly time. It should be followed by the quoted strings or byte values, separated with commas. It can be used to display values of some constants, for example:

```
d1 = '0'+ $ shr 12 and 0Fh
d2 = '0'+ $ shr 8 and 0Fh
d3 = '0'+ $ shr 4 and 0Fh
d4 = '0'+ $ and 0Fh
if d1>'9'
    d1 = d1 + 'A'-'9'-1
end if
if d2>'9'
    d2 = d2 + 'A'-'9'-1
end if
if d3>'9'
    d3 = d3 + 'A'-'9'-1
end if
if d4>'9'
    d4 = d4 + 'A'-'9'-1
end if
display 'Current offset is 0x',d1,d2,d3,d4,13,10
```

Instructions before the `display` directive calculate four digits of 16-bit value and convert them into characters for displaying.

The `store` directive can modify the already generated code by replacing some of the previously generated data with the value defined by given numerical expression, which follow. The expression can be preceded by the optional size operator to specify how large value the expression defines, and therefore how much bytes will be stored, if there is no size operator, the size of one byte is assumed. Then the `at` operator and the numerical expression defining the valid address in currently generated code space, at which the given value have to be stored should follow. This is a directive for advanced appliances and should be used carefully.

2.3 Preprocessor directives

All preprocessor directives are processed before the main assembly process, and therefore are not affected by the control directives. At this time also all comments are stripped out.

2.3.1 Including source files

include directive includes the specified source file at the position where it is used. It should be followed by the quoted name of file that should be included, for example:

```
include 'macros.inc'
```

The whole included file is preprocessed before preprocessing the lines next to the line containing the **include** directive. There are no limits to the number of included files as long as they fit in memory.

The quoted path can contain environment variables enclosed within % characters, they will be replaced with their values inside the path, both the \ and / characters are allowed as a path separators. If no absolute path is given, the file is first searched for in the directory containing file which included it and when it's not found there, in the directory containing the main source file (the one specified in command line). These rules concern also paths given with the **file** directive.

2.3.2 Symbolic constants

The symbolic constants are different from the numerical constants, before the assembly process they are replaced with their values everywhere in source lines after their definitions, and anything can become their values.

The definition of symbolic constant consists of name of the constant followed by the **equ** directive. Everything that follows this directive will become the value of constant. If the value of symbolic constant contains other symbolic constants, they are replaced with their values before assigning this value to the new constant. For example:

```
d equ dword
NULL equ d 0
d equ edx
```

After these three definitions the value of **NULL** constant is **dword 0** and the value of **d** is **edx**. So, for example, **push NULL** will be assembled as **push dword 0** and **push d** will be assembled as **push edx**.

restore directive allows to get back previous value of redefined symbolic constant. It should be followed by one more names of symbolic constants, separated with commas. So **restore d** after the above definitions will give **d** constant back the value **dword**. If there was no constant defined of given name, **restore** won't cause an error, it will be just ignored.

Symbolic constant can be used to adjust the syntax of assembler to personal preferences. For example the following set of definitions provides the handy shortcuts for all the size operators:

```
b equ byte
w equ word
d equ dword
p equ pword
f equ fword
q equ qword
t equ tword
x equ dqword
```

Because symbolic constant may also have an empty value, it can be used to allow the syntax with `offset` word before any address value:

```
offset equ
```

After this definition `mov ax,offset char` will be valid construction for copying the offset of `char` variable into `ax` register, because `offset` is replaced with an empty value, and therefore ignored.

Symbolic constants can also be defined with the `fix` directive, which has the same syntax as `equ`, but defines constants of high priority - they are replaced with their symbolic values even before processing the preprocessor directives and macroinstructions, the only exception is `fix` directive itself, which has the highest possible priority, so it allows redefinition of constants defined this way. But when such high priority constants are found inside the value following the `fix` directive, they are replaced with their values before assigning this value to the new constant.

The `fix` directive can be used for syntax adjustments related to directives of preprocessor, what cannot be done with `equ` directive. For example:

```
incl fix include
```

defines a short name for `include` directive, while the similar definition done with `equ` directive wouldn't give such result, as standard symbolic constants are replaced with their values after searching the line for preprocessor directives.

2.3.3 Macroinstructions

`macro` directive allows you to define your own complex instructions, called macroinstructions, using which can greatly simplify the process of programming. In its simplest form it's similar to symbolic constant definition. For

example the following definition defines a shortcut for the `test al,0xFF` instruction:

```
macro tst {test al,0xFF}
```

After the `macro` directive there is a name of macroinstruction and then its contents enclosed between the `{` and `}` characters. You can use `tst` instruction anywhere after this definition and it will be assembled as `test al,0xFF`. Defining symbolic constant `tst` of that value would give the similar result, but the difference is that the name of macroinstruction is recognized only as an instruction mnemonic. Also, macroinstructions are replaced with corresponding code even before the symbolic constants are replaced with their values. So if you define macroinstruction and symbolic constant of the same name, and use this name as an instruction mnemonic, it will be replaced with the contents of macroinstruction, but it will be replaced with value if symbolic constant is used somewhere inside the operands.

The definition of macroinstruction can consist of many lines, because `{` and `}` characters don't have to be in the same line as `macro` directive. For example:

```
macro stos0
{
    xor al,al
    stosb
}
```

The macroinstruction `stos0` will be replaced with these two assembly instructions anywhere it's used.

Like instructions which needs some number of operands, the macroinstruction can be defined to need some number of arguments separated with commas. The names of needed argument should follow the name of macroinstruction in the line of `macro` directive and should be separated with commas if there is more than one. Anywhere one of these names occurs in the contents of macroinstruction, it will be replaced with corresponding value, provided when the macroinstruction is used. Here is an example of a macroinstruction that will do data alignment for binary output format:

```
macro align value { rb (value-1)-($+value-1) mod value }
```

When the `align 4` instruction is found after this macroinstruction is defined, it will be replaced with contents of this macroinstruction, and the `value` will there become 4, so the result will be `rb (4-1)-($+4-1) mod 4`.

If a macroinstruction is defined that uses an instruction with the same name inside its definition, the previous meaning of this name is used. Useful redefinition of macroinstructions can be done in that way, for example:

```
macro mov op1,op2
{
    if op1 in <ds,es,fs,gs,ss> & op2 in <cs,ds,es,fs,gs,ss>
        push op2
        pop op1
    else
        mov op1,op2
    end if
}
```

This macroinstruction extends the syntax of `mov` instruction, allowing both operands to be segment registers. For example `mov ds,es` will be assembled as `push es` and `pop ds`. In all other cases the standard `mov` instruction will be used. The syntax of this `mov` can be extended further by defining next macroinstruction of that name, which will use the previous macroinstruction:

```
macro mov op1,op2,op3
{
    if op3 eq
        mov op1,op2
    else
        mov op1,op2
        mov op2,op3
    end if
}
```

It allows `mov` instruction to have three operands, but it can still have two operands only, because when macroinstruction is given less arguments than it needs, the rest of arguments will have empty values. When three operands are given, this macroinstruction will become two macroinstructions of the previous definition, so `mov es,ds,dx` will be assembled as `push ds`, `pop es` and `mov ds,dx`.

When it's needed to provide macroinstruction with argument that contains some commas, such argument should be enclosed between `<` and `>` characters. If it contains more than one `<` character, the same number of `>` should be used to tell that the value of argument ends.

purge directive allows removing the last definition of specified macroinstruction. It should be followed by one or more names of macroinstructions,

separated with commas. If such macroinstruction has not been defined, you won't get any error. For example after having the syntax of `mov` extended with the macroinstructions defined above, you can disable syntax with three operands back by using `purge mov` directive. Next `purge mov` will disable also syntax for two operands being segment registers, and all the next such directives will do nothing.

If after the `macro` directive you enclose some group of arguments' names in square brackets, it will allow giving more values for this group of arguments when using that macroinstruction. Any more argument given after the last argument of such group will begin the new group and will become the first argument of it. That's why after closing the square bracket no more argument names can follow. The contents of macroinstruction will be processed for each such group of arguments separately. The simplest example is to enclose one argument name in square brackets:

```
macro stoschar [char]
{
    mov al,char
    stosb
}
```

This macroinstruction accepts unlimited number of arguments, and each one will be processed into these two instructions separately. For example `stoschar 1,2,3` will be assembled as the following instructions:

```
mov al,1
stosb
mov al,2
stosb
mov al,3
stosb
```

There are some special directives available only inside the definitions of macroinstructions. `local` directive defines local names, which will be replaced with unique values each time the macroinstruction is used. It should be followed by names separated with commas. This directive is usually needed for the constants or labels that macroinstruction defines and uses internally. For example:

```
macro movstr
{
    local move
    move:
```

```

        lodsb
        stosb
        test al,al
        jnz move
    }

```

Each time this macroinstruction is used, `move` will become other unique name in its instructions, so you won't get an error you normally get when some label is defined more than once.

`forward`, `reverse` and `common` directives divide macroinstruction into blocks, each one processed after the processing of previous is finished. They differ in behavior only if macroinstruction allows multiple groups of arguments. Block of instructions that follows `forward` directive is processed for each group of arguments, from first to last – exactly like the default block (not preceded by any of these directives). Block that follows `reverse` directive is processed for each group of argument in reverse order – from last to first. Block that follows `common` directive is processed only once, commonly for all groups of arguments. Local name defined in one of the blocks is available in all the following blocks when processing the same group of arguments as when it was defined, and when it is defined in common block it is available in all the following blocks not depending on which group of arguments is processed.

Here is an example of macroinstruction that will create the table of addresses to strings followed by these strings:

```

macro strtbl name,[string]
{
    common
        label name dword
    forward
        local label
        dd label
    forward
        label db string,0
}

```

First argument given to this macroinstruction will become the label for table of addresses, next arguments should be the strings. First block is processed only once and defines the label, second block for each string declares its local name and defines the table entry holding the address to that string. Third block defines the data of each string with the corresponding label.

The directive starting the block in macroinstruction can be followed by the first instruction of this block in the same line, like in the following example:

```
macro stdcall proc,[arg]
{
    reverse push arg
    common call proc
}
```

This macroinstruction can be used for calling the procedures using STD-CALL convention, arguments are pushed on stack in the reverse order. For example `stdcall foo,1,2,3` will be assembled as:

```
push 3
push 2
push 1
call foo
```

If some name inside macroinstruction has multiple values (it is either one of the arguments enclosed in square brackets or local name defined in the block following `forward` or `reverse` directive) and is used in block following the `common` directive, it will be replaced with all of its values, separated with commas. For example the following macroinstruction will pass all of the additional arguments to the previously defined `stdcall` macroinstruction:

```
macro invoke proc,[arg]
{ common stdcall [proc],arg }
```

It can be used to call indirectly (by the pointer stored in memory) the procedure using STDCALL convention.

Inside macroinstruction also special operator `#` can be used. This operator causes two names to be concatenated into one name. It can be useful, because it's done after the arguments and local names are replaced with their values. The following macroinstruction will generate the conditional jump according to the `cond` argument:

```
macro jif op1,cond,op2,label
{
    cmp op1,op2
    j#cond label
}
```

For example `jif ax,ae,10h,exit` will be assembled as `cmp ax,10h` and `jae exit` instructions.

The `#` operator can be also used to concatenate two quoted strings into one. Also conversion of name into a quoted string is possible, with the `'` operator, which likewise can be used inside the macroinstruction. It convert the name that follows it into a quoted string – but note, that when it is followed by a macro argument which is being replaced with value containing more than one symbol, only the first of them will be converted, as the `'` operator converts only one symbol that immediately follows it. Here's an example of utilizing those two features:

```
macro label name
{
    label name
    if ~ used name
        display 'name # " is defined but not used.",13,10
    end if
}
```

When label defined with such macro is not used in the source, macro will warn you with the message, informing to which label it applies.

To make macroinstruction behaving differently when some of the arguments are of some special type, for example a quoted strings, you can use `eqtype` comparison operator. Here's an example of utilizing it to distinguish a quoted string from an other argument.

```
macro message arg
{
    if arg eqtype ""
        local str
        jmp    @f
        str    db arg,0Dh,0Ah,24h
        @@:
        mov    dx,str
    else
        mov    dx,arg
    end if
    mov    ah,9
    int     21h
}
```

The above macro is designed for displaying messages in DOS programs. When the argument of this macro is some number, label, or variable, the

string from that address is displayed, but when the argument is a quoted string, the created code will display that string followed by the carriage return and line feed.

It is also possible to put a declaration of macroinstruction inside another macroinstruction, so one macro can define another, but there is a problem with such definitions caused by the fact, that `}` character cannot occur inside the macroinstruction, as it always means the end of definition. But it's easy to overcome this problem with use of the `fix` directive. Here is an example:

```
macro ext instr
{
  macro instr op1,op2,op3
    _%
    if op3 eq
      instr op1,op2
    else
      instr op1,op2
      instr op2,op3
    end if
    %_
  }

  _% fix {
  %_ fix }

  ext add
  ext sub
```

The macro `ext` is defined correctly, but when it is used, the `_%` and `%_` are defined as high priority symbolic constants and are replaced with their values before doing anything else. So when the `ext add` is processed, the contents of macro becomes valid definition of a macroinstruction and this way the `add` macro becomes defined. In the same way `ext sub` defines the `sub` macro. The use of `_%` constant instead of `{` wasn't really necessary here, but is done this way to make the definition more clear. Note that the right order of definitions is crucial here.

If some directives specific to macroinstructions, like `local` or `common` are needed inside some macro embedded this way, they also have to be replaced with some symbols defined later as a high priority constants with the values being the intended directives. And when the `#` operator is needed in the embedded macro, the help comes from the fact that any sequence of `#` characters inside the macroinstruction is reduced to chain shorter by one character, and

only if it was the single `#` character, the concatenation is performed. So to put the concatenation operator inside the embedded macro it is enough to write `##` there.

The above techniques can be even applied to create the new ways of defining macroinstructions. For example:

```
macro tmacro params
{
    macro params {
    }

    MACRO fix tmacro
    ENDM fix }
```

defines an alternative syntax for defining macroinstructions, which looks like:

```
MACRO stoschar char
    mov al,char
    stosb
ENDM
```

Note that symbol that has such customized definition must be defined with `fix` directive, because only the prioritized symbolic constants are processed before the preprocessor looks for the `}` character while defining the macro. This might be a problem if one needed to perform some additional tasks one the end of such definition, but there is one more feature which helps in such cases. Namely it is possible to put any directive, instruction or macroinstruction just after the `}` character that ends the macroinstruction and it will be processed in the same way as if it was put in the next line.

2.3.4 Structures

`struc` directive is a special variant of `macro` directive that is used to define data structures. Macroinstruction defined using the `struc` directive must be preceded by a label (like the data definition directive) when it's used. This label will be also attached at the beginning of every name starting with dot in the contents of macroinstruction. The macroinstruction defined using the `struc` directive can have the same name as some other macroinstruction defined using the `macro` directive, structure macroinstruction won't prevent the standard macroinstruction being processed when there is no label before it and vice versa. All the rules and features concerning standard macroinstructions apply to structure macroinstructions.

Here is the sample of structure macroinstruction:

```

struc point x,y
{
    .x dw x
    .y dw y
}

```

For example `my point 7,11` will define structure labeled `my`, consisting of two variables: `my.x` with value 7 and `my.y` with value 11.

Next example shows how to extend the data definition directive `db` with ability to calculate the size of defined data by using the structure macroinstruction:

```

struc db [data]
{
    common
    label .data byte
    db data
    .size = $-.data
}

```

With such definition for example `msg db 'Hello!',13,10` will define also `msg.size` constant, equal to the size of defined data in bytes and also additional label `msg.data`, which will be recognized as a label for data of byte size.

Defining data structures addressed by registers or absolute values should be done using the `virtual` directive with structure macroinstruction (see 2.2.3).

2.4 Formatter directives

`format` directive followed by the format identifier allows to select the output format. This directive should be put at the beginning of the source. Default output format is a flat binary file, it can also be selected by using `format binary` directive.

`use16` and `use32` directives force the assembler to generate 16-bit or 32-bit code, omitting the default setting for selected output format.

`org` directive sets address at which the following code is expected to appear in memory. It should be followed by numerical expression specifying the address. You can also use this directive in the `$=` form followed by numerical expression.

Below are described different output formats with the directives specific to these formats.

2.4.1 MZ executable

To select the MZ output format, use **format MZ** directive. The default code setting for this format is 16-bit.

segment directive defines a new segment, it should be followed by label, which value will be the number of defined segment, optionally **use16** or **use32** word can follow to specify whether code in this segment should be 16-bit or 32-bit. The origin of segment is aligned to paragraph (16 bytes). All the labels defined then will have values relative to the beginning of this segment.

entry directive sets the entry point for MZ executable, it should be followed by the far address (name of segment, colon and the offset inside segment) of desired entry point.

stack directive sets up the stack for MZ executable. It can be followed by numerical expression specifying the size of stack to be created automatically or by the far address of initial stack frame when you want to set up the stack manually. When no stack is defined, the stack of default size 4096 bytes will be created.

heap directive should be followed by a 16-bit value defining maximum size of additional heap in paragraphs (this is heap in addition to stack and undefined data). Use **heap 0** to always allocate only memory program really needs. Default size of heap is 65535.

2.4.2 Portable Executable

To select the Portable Executable output format, use **format PE** directive, it can be followed by additional format settings: use **console**, **GUI** or **native** operator selects the target subsystem (floating point value specifying subsystem version can follow), **DLL** marks the output file as a dynamic link library. Then can follow the **at** operator and the numerical expression specifying the base of PE image and then optionally **on** operator followed by the quoted string containing file name selects custom MZ stub for PE program (when specified file is not a MZ executable, it is treated as a flat binary executable file and converted into MZ format). The default code setting for this format is 32-bit. The example of fully featured PE format declaration:

```
format PE GUI 4.0 DLL at 7000000h on 'stub.exe'
```

section directive defines a new section, it should be followed by quoted string defining the name of section, then one or more section flags can follow. Available flags are: **code**, **data**, **readable**, **writable**, **executable**, **shareable**, **discardable**, **notpageable**. The origin of section is aligned to page (4096 bytes). Example declaration of PE section:

```
section '.text' code readable executable
```

Among with flags also on of special PE data identifiers can be specified to mark the whole section as a special data, possible identifiers are **export**, **import**, **resource** and **fixups**. If the section is marked to contain fixups, they are generated automatically and no more data needs to be defined in this section. Also resource data can be generated automatically from the resource file, it can be achieved by writing the **from** operator and quoted file name after the **resource** identifier. Below are the examples of sections containing some special PE data:

```
section '.reloc' data discardable fixups
section '.rsrc' data readable resource from 'my.res'
```

entry directive sets the entry point for Portable Executable, the value of entry point should follow.

stack directive sets up the size of stack for Portable Executable, value of stack reserve size should follow, optionally value of stack commit separated with comma can follow. When stack is not defined, it's set by default to size of 4096 bytes.

heap directive chooses the size of heap for Portable Executable, value of heap reserve size should follow, optionally value of heap commit separated with comma can follow. When no heap is defined, it is set by default to size of 65536 bytes, when size of heap commit is unspecified, it is by default set to zero.

data directive begins the definition of special PE data, it should be followed by one of the data identifiers (**export**, **import**, **resource** or **fixups**) or by the number of data entry in PE header. The data should be defined in next lines, ended with **end data** directive. When fixups data definition is chosen, they are generated automatically and no more data needs to be defined there. The same applies to the resource data when the **resource** identifier is followed by **from** operator and quoted file name – in such case data is taken from the given resource file.

2.4.3 Common Object File Format

To select Common Object File Format, use **format COFF** or **format MS COFF** directive whether you want to create simple or Microsoft COFF file. The default code setting for this format is 32-bit.

section directive defines a new section, it should be followed by quoted string defining the name of section, then one or more section flags can follow. Available flags are: **code** and **data** for both COFF variants, **readable**,

`writable`, `executable`, `shareable`, `discardable` and `notpageable` only for Microsoft COFF variant. By default section is aligned to double word (four bytes), in case of Microsoft COFF variant other alignment can be specified by providing the `align` operator followed by alignment value (any power of two up to 8192) among the section flags.

`extrn` directive defines the external symbol, it should be followed by the name of symbol and optionally the size operator specifying the size of data labeled by this symbol. The name of symbol can be also preceded by quoted string containing name of the external symbol and the `as` operator. Some example declarations of external symbols:

```
extrn exit
extrn '__imp__MessageBoxA@16' as MessageBox:dword
```

`public` directive declares the existing symbol as public, it should be followed by the name of symbol, optionally it can be followed by the `as` operator and the quoted string containing name under which symbol should be available as public. Some examples of public symbols declarations:

```
public main
public start as '_start'
```

2.4.4 Executable and Linkable Format

To select ELF output format, use `format ELF` directive. The default code setting for this format is 32-bit.

`section` directive defines a new section, it should be followed by quoted string defining the name of section, then can follow one or both of the `executable` and `writable` flags, optionally also `align` operator followed by the number specifying the alignment of section (it has to be the power of two), if no alignment is specified, the default value 4 is used.

`extrn` and `public` directives have the same meaning and syntax as when the COFF output format is selected (described in previous section).

To create executable file, use `format ELF executable` directive. It allows to use `entry` directive followed by the value to set as entry point of program. On the other hand it makes `extrn` and `public` directives unavailable. `section` directive in this case can be followed only by one or more section flags and its origin is aligned to page (4096 bytes). Available flags for section are: `readable`, `writable` and `executable`.