Overview: Intel 8086 Microprocessor Organisation

Registers

ax: accumulator, main input to ALU

bx: Base register, can contain a modifiable component of an operand address

cx: Count register, can contain an iteration count for use with the LOOP instruction

dx: Data register, general-purpose

cs; Code segment register. Address of code given by CS:IP

ds: Data segment register. Most data located in DS by default

ss: Stack segment register. TOS indicated by SS:SP.

es: Extra segment register. A second data segment, used in string operations.

SP: Stack Pointer. Offset in SS of TOS.

BP: Base Pointer Offset in SS of TOS.

IP: Instruction Pointer: Offset in CS of next instruction to be executed.

Flags: 16-bit flags register Contains 6 condition flags + 3 control flags

Memory

Byte-organised memory, single contiguous organisation, viewable as segmentable from a programmer's point of view.

4 segments: Code, Data, Stack, Extra, supported in hardware by segment registers that provide the most significant 16 bits of a 20-bit memory address.

The 4 LSBs of a segment address are assumed to be zero. Thus, segments are aligned (i.e. begin at) a 16-byte boundary (an address evenly divisible by 16).

All other addresses produced in the processor are treated as 16-bit offsets within a particular segment. Thus, the maximum size of a memory segment is 64K.

Segments are not mutually exclusive: they can overlap in any number of locations. Indeed, setting all segment registers to zero results in the previous 64K memory model.

Offsets are combined with segment register contents to generate effective addresses.

Usually, the IP contains an offset within the code segment and the SP contains an offset within the stack segment.

This may be specified at assembly level by writing CS:IP and SS:SP.

Operand offsets usually act with respect to the DS or ES registers.

To form a 20-bit effective address, the four LSBs of the offset become bits 0-3 of the effective address. Bits 4 - 19 of the effective address are constructed by adding bits 4-15 of the offset to bits 0-15 of the segment address.

For instance, suppose the CS register contains 302H and the IP contains 2A57H. Then, the code segment address is 3020H and the effective address is 5A77H.

Instructions are variable-length and can occupy 1, 2 or 3, 4, 5 7 or 9 bytes.

Operand addressing modes available are implied, register, immediate, direct, register indirect, memory indirect, based and indexed.

Indeed, although separate based and indexed addressing modes may be discussed as an abstraction, 80x86 addressing is best treated by saying that an operand address may be formed by adding a displacement to the contents of two other registers.

Memory is byte-organised, with the Little-Endian convention employed for the storage of multiple-byte quantities.

The restart address is FFFF0H and Intel literature "reserves" locations FFFF0H to FFFFFH (16 bytes) for a jump to the initial bootstrap loader.

Memory locations 00000H to 003FFH are reserved for the interrupt vector, so 256 entries of 4 bytes each occupy 1K memory bytes.

Note that full address quantities in this processor occupy 4 bytes, 2 each for the segment address and the offset.

Stack is located off-chip, in main memory and is accessed through SS: SP.

The stack expands into memory locations having lower addresses than those already occupied.

The SP is decremented before a byte is saved on the stack.

Only 16-bit quantities are moved to and from the stack.