1

INTRODUCTION TO MICROPROCESSORS AND MICROCOMPUTER-BASED APPLICATIONS

This chapter provides a brief summary of the features of microprocessors and microcomputer-based applications.

The basic elements of a computer are the Central Processing Unit (CPU), the Memory, and Input/Output (I/O) units. The CPU translates instructions, performs arithmetic or logic operations, and temporarily stores instructions and data in its internal high-speed registers. The memory stores programs and data. The I/O unit interfaces the computer with external devices such as keyboard and display.

With the advent of semiconductor technology, it is possible to integrate the CPU in a single chip. The result is the microprocessor. Metal Oxide Semiconductor (MOS) technology is typically used to fabricate the standard off-the-shelf microprocessors such as those manufactured by Intel and Motorola. Appropriate memory and I/O chips are interfaced to the microprocessor to design a microcomputer. Single-chip microcomputers are also available in which the microprocessor, memory, and I/O are all fabricated in the same chip. These single-chip microcomputers offer limited capabilities. However, they are ideal for certain applications such as peripheral controllers.

Single chip microcomputers are also referred to as "microcontrollers". The microcontrollers are typically used for dedicated applications such as automotive systems, home appliances and home entertainment systems. Typical microcontrollers, therefore, include on-chip timers, A/D (Analog to Digital) and D/A (Digital to Analog) converters. Two popular microcontrollers are Intel 8751 (8-bit)/8096 (16-bit) and Motorola HC11 (8-bit)/HC16 (16-bit). The 16-bit microcontrollers include more on-chip ROM, RAM, and I/O compared to the 8-bit microcontrollers.

The efficient development of microprocessor-based systems necessitates the use of a microcomputer development system. The microcomputer development system is used for the design, debugging, and sometimes the documentation of a microprocessor-based system.

This chapter first covers the evolution of 8-, 16-, and 32-bit microprocessors along with an overview of programming languages, microcomputer hardware, and software. The attributes of typical microcomputer development systems features as well as some specific microprocessor applications are also included.

1.1 Evolution of the Microprocessor

Intel Corporation introduced the first microprocessor, the 4004 (4-bit), in 1971. The 4004 evolved from a development effort while designing a calculator chip set.

Soon after the 4004 appeared in the commercial market, three other microprocessors were introduced. These were the Rockwell International 4-bit PPS-4, the Intel 8-bit 8008, and the National Semiconductor 16-bit IMP-16.

The microprocessors introduced between 1971 and 1973 were the first-generation systems. They were designed using the PMOS (P-type MOS) technology. This technology provided low cost, slow speed, and low output currents and was not compatible with TTL (Transistor Transistor Logic).

After 1973, second-generation microprocessors such as MOS Technology 6502, Motorola 6800 and 6809, Intel 8085, and Zilog Z80 evolved. These 8-bit processors were fabricated using the NMOS (N-type MOS) technology. The NMOS process offers faster speed and higher density than PMOS and is TTL-compatible.

After 1978, the third-generation microprocessors were introduced. These processors are 16 bits wide (16-bit ALU) and include typical processors such as Intel 8086/80186/80286 and Motorola 68000/68010. These microprocessors were designed using the HMOS (high-density MOS) technology. HMOS provides the following advantages over NMOS:

- Speed-Power-Product (SPP) of HMOS is four times better than NMOS:
  NMOS = 4 Picojoules (PJ)
  HMOS = 1 Picojoule (PJ)

  Note that Speed-Power-Product
  = speed * power
  = nanosecond * milliwatt
  = picojoules

- Circuit densities provided by HMOS are approximately twice those of NMOS:
  NMOS = 4128 gates/μm^2
  HMOS = 1852.5 gates/μm^2

  where 1 μm (micrometer) = 10^{-6} meter.

Later, Intel utilized the HMOS technology to fabricate the 8085A. Thus, Intel offers a high-speed version of the 8085A called 8085AH. The price of the 8085AH is higher than the 8085A.

In 1980, fourth-generation microprocessors evolved. Intel introduced the first commercial 32-bit microprocessor, the problematic Intel 432. This processor was eventually discontinued by Intel. Since 1985, more 32-bit microprocessors have been introduced. These include Motorola's MC 68020/68030/68040/PowerPC, Intel's 80386/80486 and the Intel Pentium processor. These processors are fabricated using the low-power version of the HMOS technology called the HCMOS.

The performance offered by the 32-bit microprocessor is more comparable to that of superminicomputers such as Digital Equipment Corporation's VAX11/750 and VAX11/780. Both Intel and Motorola introduced 32-bit RISC (Reduced Instruction Set Computer) microprocessors, namely the Intel 80960 and Motorola MC88100/PowerPC with simplified instruction sets. Note that the purpose of RISC microprocessors is to maximize speed by reducing clock cycles per instruction. Almost all computations can be obtained from a simple instruction set.

RISC microprocessors have hardwired instruction sets like the non-RISC microprocessors such as the Intel 8085 and MOS 6502. This means that for every instruction, there exists actual, physical connections that provide the desired instruction decoding. The microprocessor,

therefore, does not use valuable clock cycles (machine cycles) in instruction decoding. Most RISC instructions require a maximum of only two clock cycles to complete. These instructions are restricted to register-to-register operations with load and store for memory access. Since RISC-type microprocessors are hardwired, instructions may be executed simultaneously (as long as the instructions do not share the same register.) This technique is known as pipelining.

Pipelining is the mechanism that actually enables simultaneous processing to occur. Instructions are fetched in sequential order. The processor continues to fetch instructions even though it has not executed the present instruction. This provides an input pipeline to the instruction cache. At this stage, the instructions are intelligently fed into the instruction execution unit for processing. When an instruction is executed, the registers used by the instruction clear respective bits in a scoreboard register. As the microprocessor is executing a given instruction, it checks the scoreboard register to see if the next instructions (residing in a fast read/write memory internal to the microprocessor called the instruction cache) use registers that are currently in use. Instructions not using the same registers can therefore be executed at the same time. The RISC microprocessors can execute as many as five instructions simultaneously.

The trend in microprocessors is implementation of more on-chip functions and for improvement of the speeds of memory and I/O devices. Some manufacturers are speeding up the processors for data crunching type applications. Digital Equipment Corporation's Alpha 21164 with 300 MHz clock, four instruction-per-cycle rate and RISC-based architecture is the fastest microprocessor available today.

1.2 Microprocessor Data Types

This section discusses the data types used by microprocessor. Typical data types include signed and unsigned binary integers, binary coded decimal (BCD), American Standard Code for Information Interchange (ASCII) and floating-point numbers. Note that binary integer numbers do not support fractions in microprocessors. Fractions and mixed numbers (numbers comprised of integers and fractions) use either binary or BCD floating-point formats. Floating-point numbers are often referred to as real numbers.

1.2.1 Unsigned and Signed Binary Integers

An unsigned binary integer has no arithmetic sign. Unsigned binary numbers are therefore always positive. An example is memory address which is always a positive number. An 8-bit unsigned binary integer represents all numbers from 00_{16} through FF_{16} (0_{10} through 255_{10}).

A signed binary integer, on the other hand, includes both positive and negative numbers. It is represented in true form for a positive number and in two's complement form for a negative number. For example, the decimal number +10 can be represented as a true form 8-bit number in a microprocessor as 0000 1010 (binary) or 0A (hexadecimal). The decimal number -10 can be represented in two's complement form as 1111 0110 (binary) or F6 (hexadecimal). The most significant bit of a signed binary number represents the sign of the number. For example, bit 7 of an 8-bit signed number, bit 15 of a 16-bit signed number and bit 31 of a 32-bit signed number represent the signs of the respective numbers. A 32-bit signed binary integer includes all numbers from -2,147,483,648 to 2,147,483,647 with 0 being a positive number.

Finally, note that the hexadecimal number FF_{16} is 255_{10} when represented as an unsigned number. On the other hand, FF_{16} is -1_{10} when represented as a signed number.

1.2.2 BCD (Binary Coded Decimal) Numbers

A BCD digit consists of four bits with a value ranging from 0000_2 to 1001_2 (0 through 9 decimal). A BCD digit greater than 9_{10} can be represented as two or more BCD digits.

Microprocessors store BCD numbers in two forms, packed and unpacked. The unpacked BCD number represents each BCD digit as a byte while the packed BCD number represents two BCD digits in a byte. For example, 23_{10} is represented as $0000\ 0010\ 0000\ 0011_2$ as two unpacked BCD numbers, while it is represented as $0010\ 0011_2$ as a packed BCD number.

Microprocessors normally input data from a keypad in unpacked BCD form. This data can then be converted by writing a program in the microprocessor to packed BCD form for arithmetic operations or for storing in memory. After processing, the packed BCD result is then converted to unpacked BCD form by another microprocessor program written by the user for displays. Typical displays use either unpacked BCD data or unpacked BCD data converted to seven segment code by the microprocessor's program.

1.2.3 ASCII

ASCII (American Standard Code for Information Interchange) is a code that represents alphanumeric (alpha characters and numbers) in a microcomputer's memory. ASCII also represents special symbols such as # and %. It is a 7-bit code. The most significant bit (bit 7) is sometimes used as a parity bit. The parity bit represents the number of ones in the byte. If the number of ones is odd, the parity is odd; otherwise, the parity is even.

Also, note that the hexadecimal numbers 30_{16} through 39_{16} are ASCII codes representing the decimal numbers 0 through 9. A listing of the ASCII codes is included in Appendix F.

1.2.4 Floating-Point Numbers

Floating-Point numbers contain three components. These are a sign, exponent and mantissa. For example, consider the decimal value -2.5×10^{-2}. The sign is negative, the exponent is -2 and the significand or mantissa is 2.5. In floating point numbers, it is possible to store the same number in several different ways. For example, 1 can be represented as 10.0×10^{-1}, 1.0×10^0, 0.1×10^1 and 0.01×10^2. To make computations yield the maximum accuracy, the numbers are normalized. This means that the exponent is adjusted so that the mantissa always follows a specific format. A binary floating point number is represented as a normalized binary fraction raised to a power of 2. To convert a binary number to its standard floating-point form, the binary number is converted to a normalized floating-point number. Note that a normalized binary floating-point number is represented as 1.XXXXXX raised to some power of 2 where X can be 0 or 1. The 1. part is implied and is not stored by the microprocessor. First, the binary number is converted to a common 32-bit floating-point format. The most significant bit (bit 31) is the sign bit (S). If S is zero, the number is positive, while if S is one, the number is negative. The next 8 bits contain the biased exponent. This means that an 8-bit number $7F_{16}$ or 127_{10} is added to the exponent. Biased exponent makes numeric comparison (such as less than or greater than) easy. The bias is usually chosen such that the most negative number allowed in the exponent becomes zero and the most positive number becomes the largest value of the representation. For example, with an 8-bit exponent, and $+127_{10}$ bias, the smallest and the largest values of the bias exponent are 0_{10} and 255_{10} respectively. Note that the zero biased exponent is represented as 127_{10}, while the unbiased minimum and maximum values of the 8-bit exponent are -127_{10} and $+128_{10}$ respectively. The remaining 23 bits represent the fractional part of the number. Note that zero is stored as 32 zeros while infinity is stored as 32 ones. As an example, consider converting the decimal number 10 to the standard floating-point format as follows:

1. $10_{10} = 1010_2$
2. Normalize the binary number as $1.XXXXXX \times 2^n = 1.010 \times 2^3$
3. Sign, $S = 0$ for positive
4. Biased exponent $= 7F_{16} + 3 = 82_{16}$

5. 23-bit Fraction = 0100 0000 0000 0000 0000 000
6. The floating-point equivalent of 10_{10} is

S	Exponent	Fraction
0	1000 0010	0100 0000 0000 0000 0000 0000

A special case of the floating-point format is called NaN (Not a Number). NaNs are results generated by floating-point operations that have no mathematical interpretations. These results may be generated by operations such as multiplication of infinity by infinity.

The BCD floating-point form represents a number in BCD scientific notation. The number is represented as normalized significand raised to some power of 10. Each BCD floating-point number is represented in typical microprocessors as 80 bits. The BCD fraction is 16 digits wide (64-bit) and is stored as packed BCD digits. The whole number portion of the significand is stored as one digit BCD from 0 to 9. The BCD exponent along with the sign is expressed as 12 bits.

Typical floating-point coprocessors such as 80387 (for the Intel 80386 microprocessor) and 68881/68882 (for the Motorola 68020/68030 microprocessor) support several data types. For example, the 80387 coprocessor supports seven data types. These are word integer (16-bits), short integer (32-bit), long integer (64-bit), packed BCD (80-bit), short real (32-bit), long real (64-bit), and temporary real (80-bit).

The 80387 integer data types are represented by the two's complement same as those used by the 80386. The only difference is that the 80386 supports an 8-bit integer while the 80387 supports a 64-bit integer.

The 80387 supports 80-bit packed BCD with 18 decimal digits (bits 0-71), bit 79 as the sign-bit and seven (bits 72-78) unused bits. With 18-digits representation, the COBOL standard (the High level language utilizing BCD) is followed.

The 80387 supports three real data formats. These are short real (32-bit with one sign bit, 8-bit exponent and 23-bit significand), long real (64-bit with one sign-bit, 11-bit for exponent and 62-bit for the significand), temporary real (80-bit with one sign-bit, 15-bit exponent, and 64-bit significand).

The 80387 uses the temporary real format internally. All data types are converted by the 80387 immediately into temporary real. This is done to maximum precision and range of computations.

The 80387 supports four special cases. These are zeros, infinities (both positive and negative), denormals, and NaNs (signalling and quiet).

Denormals represent very small numbers that are not normalized. Normally, numbers are required to be normalized by shifting to left until the most significant is one. Denormals do not have one as the most significant bit of the significand. Denormals permit a gradual underflow. That is the precision is lost gradually rather than abruptly. When the least normal-izable number is reached, the next small representation is zero. Denormals provide gradual underflow of numbers that are not normalized. That is, denormals extend the range of very small numbers significantly, but with some loss in precision.

A signaling NaN causes an invalid operation exception when used in an operation. A quiet NaN, on the other hand, does not cause an invalid operation exception.

The floating-point data types supported by the Motorola 68881/68882 floating-point coprocessors are summarized next. Note that the 68881 and 68882 differ in execution speed. They are basically identical. The 68882 is an enhanced version of the 68881 in that it executes several floating-point instructions concurrently with the 68020/68030. The 68881/68882 supports integers, binary floating-point numbers, and packed floating-point BCD. Data are represented externally by using these formats. The 68881/68882 utilizes an 80-bit binary floating point form to represent all data internally.

The 68881/68882 supports three signed integer formats. These are 8-bit byte, 16-bit word, and 32-bit long word.

Binary floating-point format is also called binary real form. The 68881/68882 supports binary floating-point form which contains three fields. These are a sign, biased exponent and a significand. The 68881/68882 operates on these sizes, namely 32-bit single-precision, 64-bit double precision, and 96-bit extended precision.

For single precision (bit 31: sign bit, bits 23-30: 8-bit exponent, bits 0-22: 23-bit significand), 64-bit double precision (bit 63: sign bit, bits 62-52: 11-bit exponent, bits 0-51: 52-bit significand), and 96-bit extended-precision (bit 95: sign bit, bits 80-94: 15-bit exponent, bits 64-79: zero; sixteen unused bits, bits 0-63: 64-bit significand).

The biased exponent is used. The Single Precision adds a bias of 127_{10} ($7F_{16}$), double Precision uses a bias of 1023_{10} ($3FF_{16}$), and the extended-Precision uses a bias of 16383_{10} ($3FFFF_{16}$). The bias is added to the exponent before it is stored in this format and subtracted to convert to a true exponent when the number is interpreted.

A few special cases that do not conform to the floating-point form are also handled by the 68881/68882. For example, a zero is represented with all bits of the exponent and significand as zeros. The sign bit may be a one or a zero representing +0 or -0. The infinity, on the other hand is represented by all bits in the exponent and significand set to ones. The sign bit may be zero or one representing positive or negative infinity. The 68881/68882 supports BCD floating point form which represent each number as normalized significand raised to a power of 10. This format stores a number as 96 bits. The least significant 64 bits (8 bytes) contain the 16-digit BCD fraction. The next byte contains the whole number portion of the significand (0-9). The most significant bit (bit 95) contains the sign of significand while bit 94 includes the sign of the exponent. The exponent is represented by three digit BCD packed exponent (000-999) in bits 80–91. Similar to the 80387, the 68881/68882 also represents NaN's and also provides exceptions for signaling NaN's.

IEEE has established the standard for floating-point arithmetic specified by ANSI-IEEE754-1985. Typical 32-bit microprocessors use this standard.

1.3 Microcomputer Hardware

In this section, some unique features associated with various microcomputer components will be described.

The microcomputer contains a microprocessor, a memory unit, and an input/output unit. These elements are explained in the following in detail. Figure 1.1 shows a simplified block diagram of a microcomputer.

1.3.1 The System Bus

The system bus contains three buses. These are the address bus, the data bus, and the control bus. These buses connect the microprocessor to each of the memory and I/O elements so that information transfer between the microprocessor and any of the other elements can take place.
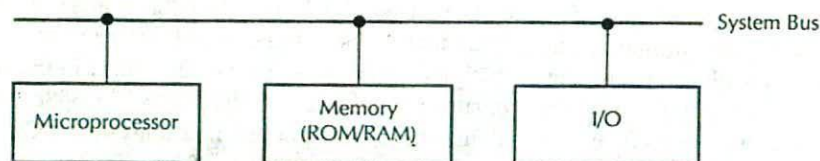

FIGURE 1.1   Simplified block diagram of a microprocessor.

On the address bus, information transfer normally takes place only in one direction, from the microprocessor to the memory or I/O elements. Therefore, this is called a unidirectional bus. This bus is usually 16 to 32 bits wide. The number of unique addresses that the microprocessor can generate on this bus depends on the width of this bus. For example, for a 16-bit address bus, the microprocessor can generate $2^{16} = 65,536$ different possible addresses. A different memory location or an I/O element can be represented by each one of these addresses.

The data bus is a bidirectional bus, that is, information can flow in both directions, to or from the microprocessor. This bus is normally 8, 16, or 32 bits wide.

The control bus is used to transmit signals that are used to synchronize the operation of the individual microcomputer elements. Typical control signals include READ, WRITE, and RESET. Some signals on the control bus such as interrupt signals are unidirectional, while some others such as RESET may be bidirectional.

1.3.2 The Microprocessor

The commercial microprocessor, fabricated using the MOS technology, is normally contained in a single chip. The microprocessor is comprised of a register section, one or more ALUs (Arithmetic Logic Units), and a control unit. Depending on the register section, the microprocessor can be classified either as an accumulator-based or a general-purpose register-based machine.

In an accumulator-based microprocessor such as the Intel 8085 and Motorola 6809, one of the operands is assumed to be held in a special register called the "accumulator". All arithmetic and logic operations are performed using this register as one of the data sources. The result after the operation is stored in the accumulator. One-address instructions are very predominant in this organization. Eight-bit microprocessors are usually accumulator-based.

The general-purpose register-based microprocessor is usually popular with 16- and 32-bit microprocessors, such as Intel 8086/80386/80486 and Motorola 68000/68020/68030/68040, and is called general-purpose, since its registers can be used to hold data, memory addresses, or the results of arithmetic or logic operations. The number, size, and types of registers vary from one microprocessor to another. Most registers are general-purpose registers, while some are provided with dedicated functions.

Typical dedicated registers include the Program Counter (PC), the Instruction Register (IR), Status Register (SR), the Stack Pointer (SP) and the Index Register. The 32-bit microprocessors include special on-chip combinational network called the Barrel Shifter.

The PC normally contains the address of the next instruction to be executed. Upon activating the microprocessor chip's RESET input, the PC is normally initialized with the address of the first instruction. For example, the 80486, upon hardware reset, reads the first instruction from the 32-bit address $FFFFFFF0_{16}$. In order to execute the instruction, the microprocessor normally places the PC contents on the address bus and reads (fetches) the first instruction from external memory. The program counter contents are then automatically incremented by the ALU. The microcomputer thus executes a program sequentially unless it encounters a jump or branch instruction. The size of the PC varies from one microprocessor to another depending on the address size. For example, the 68000 has a 24-bit PC, while the 68040 contains a 32-bit PC.

The instruction register (IR) contains the instruction to be executed. After fetching an instruction from memory, the microprocessor places it in the IR for translation.

The status register contains individual bits with each bit having a special meaning. The bits in the status register are called flags. Each flag is usually set or reset by an ALU operation. The flags are used by the Conditional Branch instructions. Typical flags include carry, sign, zero, and overflow.

The carry (C) flag is used to reflect whether or not an arithmetic operation such as ADD generates a carry. The carry is generated out of the 8th bit (bit 7) for byte operations, 16th bit (bit 15) for 16-bit, or 32nd bit (bit 31) for 32-bit operations. The carry is used as the borrow flag for subtraction. In multiple word arithmetic operations, any carry from a low-order word must be reflected in the high-order word for correct results.

The zero (Z) flag is used to indicate whether the result of an arithmetic or logic operation is zero. $Z = 1$ for a zero result and $Z = 0$ for a non-zero result. The sign flag (sometimes also called the negative flag) indicates whether a number is positive or negative. $S = 1$ indicates a negative number if the most significant bit of the number is one; $S = 0$ indicates a positive number if the most significant bit of the number is zero.

The overflow (V) flag is set to one if the result of an arithmetic operation on signed (two's complement) numbers is too large for the microprocessor's maximum word size; the C flag is overflow for unsigned numbers. The overflow flag for signed 8-bit numbers can be shown as $V = C7 \oplus C6$, where C7 is the final carry and C6 is the previous bit's carry. The $\oplus$ symbol indicates exclusive-OR operation. This can be illustrated by the numerical examples shown below:

$$
\begin{array}{ll}
\ 0000\ 0100 & 04_{10} \\
+\ 0000\ 0010 & +\ 02_{10} \\
\hline
C7 = 0 \leftarrow 0000\ 0110 & \overline{06_{10}} \\
\\
C6 = 0 \leftarrow \ 0 \quad 6_{16} &
\end{array}
$$

From the above, the result is correct when C6 and C7 have the same values (0 in this case). When C6 and C7 are different, an overflow occurs. For example, consider the following:

$$
\begin{array}{ll}
\ 1011\ 1110 & -\ 66_{10} \\
\ 1011\ 1111 & -\ 65_{10} \\
\hline
\ 0111\ 1101 & +\ 125_{10(?)} \\
& \text{result is} \\
C7 = 1 \leftarrow & \text{incorrect} \\
C6 = 0 \leftarrow &
\end{array}
$$

The result is incorrect. Since $V = C6 \oplus C7 = 0 \oplus 1 = 1$, the overflow flag is set. Note that this applies to signed two's complemented numbers only.

The stack pointer (SP) register addresses the stack. A stack is Last-In First-Out (LIFO) read/write memory in the sense that items that go in last will come out first. This is because stacks perform all read (POP) and write (PUSH) operations from one end.

The stack is addressed by a register called the stack pointer (SP). The size of the SP is dependent on the microprocessor's address size. The stack is normally used by subroutines or interrupts for saving certain registers such as the program counter.

Two instructions, PUSH (stack write) and POP (stack read), can usually be performed by the programmer to manipulate the stack. If the stack is accessed from the top, the stack pointer is decremented before a PUSH and incremented after a POP. On the other hand, if the stack is accessed from the bottom, the SP is incremented before a PUSH and decremented after a POP. Typical microprocessors access the stack from the top. Depending on the microproces-
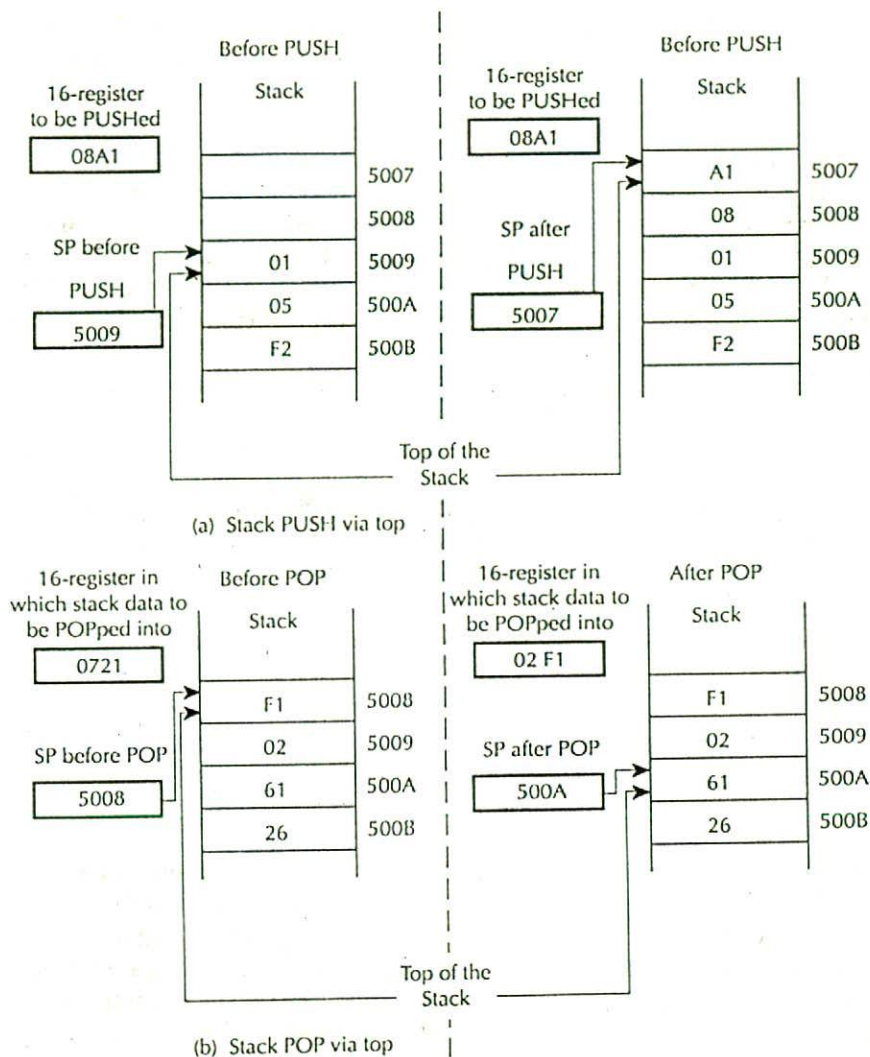
FIGURE 1.2   PUSH and POP operations via top.

sor, an 8-, 16-, or 32-bit register can be pushed onto or popped from the stack. The value by which the SP is incremented or decremented after POP or PUSH operations depends on the register size. For example, values of one for an 8-bit register, two for 16-bit registers, and four for 32-bit registers are used. Figure 1.2 shows the stack data when accessed from the top before and after PUSHing a 16-bit register onto the stack or POPping 16 bits from the stack into the 16-bit register. Note that stack items PUSHed must be POPped in reverse order. The item pushed last must be popped first.

Consider the PUSH operation in Figure 1.2a when the stack is accessed from the top. The SP is decremented by 2 after the PUSH. The SP is decremented since it is accessed from the top. A decrement value of 2 is used since the register to be pushed is 16 bits wide.

The POP operation shown in Figure 1.2b is the reverse of the PUSH. The SP is incremented after POP. The contents of locations 5008_{16} and 5009_{16} are assumed to be empty.

An index register is typically used as a counter for an instruction or for general storage functions. The index register is useful with instructions where tables or arrays of data are accessed. The general-purpose register-based microprocessor can use any general-purpose register as the index register.

Typical 32-bit microprocessors such as the Intel 80386/80486 and Motorola 68020/68030/68040 include a special type of shifter called barrel shifter for performing fast shift operations.

The barrel shifter is an on-chip combinational network for 32-bit microprocessors and provides fast shift operations. For example, the 80386 barrel shifter can shift a number from 0 through 64 positions in one clock period (clock rate is 16.67 MHz).

The ALU in the microprocessor performs all arithmetic and logic operations on data. The size of the ALU defines the size of the microprocessor. For example, Intel 8086 (or Motorola 68000) is a 16-bit microprocessor since its ALU is 16 bits wide. The Intel 8088 (or Motorola 68008) is also a 16-bit microprocessor since its ALU is 16 bits wide, even though its data bus is 8 bits wide. Motorola 68040 (or Intel 80486) is a 32-bit microprocessor since its ALU is 32 bits wide. The ALU usually performs operations such as binary addition and subtraction. The 32-bit microprocessors include multiple ALUs for parallel operations and thus achieve fast speed.

The control unit of the microprocessor performs instruction interpreting and sequencing. In the fetch phase, the control unit reads instructions from memory using the PC as a pointer. It then recognizes the instruction type, gets the necessary operands, and routes them to the appropriate functional units of the execution unit. Necessary signals are issued to the execution unit to perform the desired operations, and the results are routed to the specified destination.

In the sequencing phase, the control unit determines the address of the next instruction to be executed and loads it into the PC. The control unit is typically designed using one of three techniques:

- Hardwired control
- Microprogramming
- Nanoprogramming

The hardwired control unit is designed by physically connecting typical components such as gates and flip-flops. Typical 32-bit RISC microprocessors such as the Intel 80960 and Motorola 88100 are designed using hardwired control. The microprogrammed control unit includes a control ROM for translating the instructions. Intel 8086 is a microprogrammed microprocessor. Nanoprogramming includes two ROMs inside the control unit. The first ROM (microROM) stores all the addresses of the second ROM (nanoROM). If the microinstructions (which is the case with the 68000/68020/68030/68040) repeat many times in a microprogram, use of two-level ROMs provides tremendous memory savings. This is the reason that the control units of the 68000, 68020, 68030, and 68040 are nanoprogrammed.

1.3.3 Memory Organization

1.3.3.a Introduction

A memory unit is an integral part of any microcomputer system and its primary purpose is to hold programs and data.

In a broad sense, a microcomputer memory system can be logically divided into three groups:

- Processor memory
- Primary or main memory
- Secondary memory

Processor memory refers to the microprocessor registers. These registers are used to hold temporary results when a computation is in progress. Also, there is no speed disparity between these registers and the microprocessor because they are fabricated using the same technology.

However, the cost involved in this approach forces a microcomputer architect to include only a few registers (usually 8 or 16) in the microprocessor.

Primary or main memory is the storage area in which all programs are executed. The microprocessor can directly access only those items that are stored in primary memory. Therefore, all programs and data must be within the primary memory prior to execution.

Secondary memory refers to the storage medium comprising slow devices such as magnetic tapes and disks. These devices are used to hold large data files and huge programs such as compilers and data base management systems which are not needed by the processor frequently. Sometimes secondary memories are also referred to as auxiliary or backup store or virtual memory.

Secondary memory stores programs and data in excess of the main memory. The microcomputer cannot directly execute programs stored in the secondary memory. In order to execute these programs, the microcomputer must transfer them to its main memory by a system program called the operating system. This topic is covered later in the chapter.

Data in disk memories are stored in tracks. A track is a concentric ring of data stored on a surface of a disk. Each track is further subdivided into several sectors. Each sector typically stores 512 or 1024 bytes of data. All disk memories use magnetic media except the optical disk memory which stores data on a plastic disk. Data is read or sometimes written on the optical disk with a laser beam. There are two types of optical disks. These are the CD-ROM (Compact Disk Read Only Memory) and the WORM (Write Once Read Many). The CD-ROM is inexpensive compared to the WORM drive. However it suffers from lack of speed and has limited software applications at the present time. The WORM drive is typically used in huge data storing applications such as insurance and banking since data can be written only once. The optical disk memory is currently becoming popular with microcomputer systems. One of the commonly used disk memories with microcomputer systems is the floppy or flexible disk. The floppy disk is a flat, round piece of plastic coated with magnetically sensitive oxide material. The disk is provided with a protective jacket to prevent fingerprints or foreign matter from contaminating the disk's surface. The floppy disk is available in three sizes. These are the 8 inch, 5.25 inch, and 3.5 inch. The 8 inch floppy disk is not used in present systems. These days, the 5.25 inch and 3.5 inch are very popular. Also, the 3.5 inch floppy is replacing the 5.25 inch floppy in newer systems since it is smaller in size and does not bend easily. All floppy disks are provided with an off-center index hole that allows the electronic system reading the disk to find the start of a track and the first sector.

Hard disk memory is also frequently used with microcomputer systems. The hard disk, also known as the fixed disk, is not removable like the floppy disk.

A comparison of the some of the features associated with the hard disk and floppy disk is provided below:

Characteristic	Hard Disk	Floppy Disk
Size	5 Mbytes to several Gbytes	1.2 Mbytes typical for 5.25 inch floppy. 1.44 Mbytes typical for 3.5 inch floppy.
Rotational Speed	3600 rpm	300 rpm
Number of heads	May have up to 8 disk surfaces with up to two heads per surface	Two heads; One head for the upper surface and the other head for the lower surface

Primary memory normally includes ROM (Read-only Memory) and RAM (Random Access Memory). As the name implies, a ROM permits only a read access. Some ROMs are custom made, that is, their contents are programmed by the manufacturer. Such ROMs are called mask programmable ROMs. Sometimes a user may have to program a ROM in the field. For instance, in a fusible-link ROM, programmable read-only memory (PROM) is available. The main disadvantage of a PROM is that it cannot be reprogrammed.

Some ROMs can be reprogrammed, these are called Erasable Programmable Read-Only Memories (EPROMs).

In an EPROM, programs are entered using electrical impulses and the stored information is erased by using ultraviolet rays. Usually an EPROM is programmed by inserting the EPROM chip into the socket of a PROM programmer and providing program addresses and voltage pulses at the appropriate pins of the chip. Typical erase times vary between 10 and 30 minutes.

With advances in IC technology, it is possible to achieve an electrical means of erasure. These new ROMs are called Electrically Alterable ROMs (EAROMs) or Electrically Erasable PROMs (EEPROMs or E²PROMs) and these ROM chips can be programmed even when they are in the circuit board. These memories are also called Read Mostly Memories (RMMs), since they have much slower write times than read times. Random Access Memories (RAMs) are read/write memories.

Information stored in random access memories will be lost if the power is turned off. This property is known as volatility and, hence, RAMs are usually called volatile memories. RAMs can be backed up by batteries for a certain period of time and are sometimes called nonvolatile RAMs. Stored information in a magnetic tape or magnetic disk is not lost when the power is turned off. Therefore, these storage devices are called nonvolatile memories. Note that a ROM is a nonvolatile memory.

Some RAMs are constructed using bipolar transistors, and the information is stored in the form of voltage levels in flip-flops. These voltage levels do not usually drift away, or decay. Such memories are called static RAMs because the stored information remains constant for some period of time.

On the other hand, in RAMs that are designed using MOS transistors, the information is held in the form of electrical charges in capacitors. Here, the stored charge has the tendency to decay. Therefore, a stored 1 would become a 0 if no precautions were taken. These memories are referred to as dynamic RAMs. In order to prevent any information loss, dynamic RAMs have to be refreshed at regular intervals. Refreshing means boosting the signal level and writing it back. This activity is performed by a hardware unit called "refresh logic" which can either be a separate chip or is contained in the microprocessor chip.

Since the static RAM maintains information in active circuits, power is required even when the chip is inactive or in standby mode. Therefore, static RAMs require large power supplies. Also, each static RAM cell is about four times larger in area than an equivalent dynamic cell; a dynamic RAM chip contains about four times as many bits as a static RAM chip using the same or comparable semiconductor technology. Figure 1.3 shows the subcategories of ROMs, RAMs, and their associated technologies.
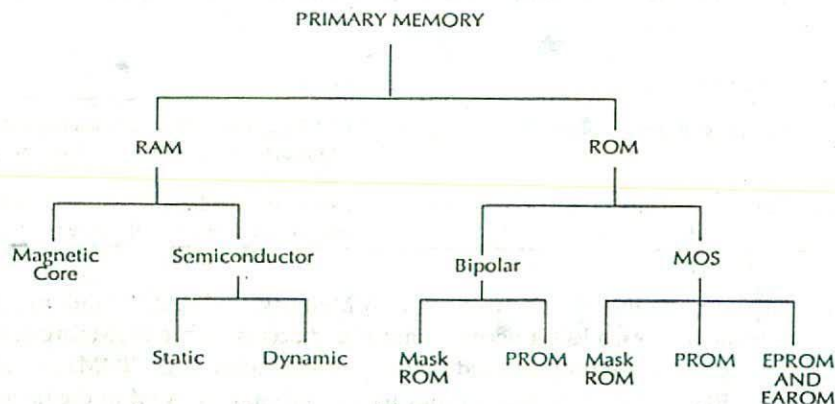

FIGURE 1.3   Subcategories of RAMs and ROMs.

Today, one megabit of data can be stored in an ordinary dynamic RAM chip. The data can be accessed in 80 nanoseconds or less. The RAM chip costs $5. In contrast, it takes 150 nanoseconds to access a one-megabit EEPROM which costs $150. Sixteen megabit DRAMs are very popular these days at a price of approximately 0.3 millicent per bit. Recently, IBM, Hitachi, Toshiba and others have introduced 64 mega-bit DRAMs. It is expected that giga-bit DRAMs will not be introduced until the next century.

In the mid 1980s, Toshiba Semiconductor invented flash memory. About the same time, Intel and Seeq Semiconductor were also working on flash memories. While each manufacturer implemented its flash memory differently, they operate in a similar way.

Like EPROMs and EEPROMs or EAROMs, flash memory is nonvolatile and reprogrammable. Flash memory is fabricated by using ETOX II (EPROM Tunnel Oxide) technology which is a combination of EPROM and EEPROM technologies. Flash memory is relatively inexpensive compared to EEPROM. A one megabit flash memory costs about $15. Flash memory can be reprogrammed electrically while embedded in the board. However, one can only change a sector or a block (consisting of multiple bytes) at a time.

Flash memory cells contain a single transistor like the EPROM cell. In contrast, a DRAM cell typically contains a transistor and a capacitor, an EEPROM cell contains two transistors while a static RAM cell requires four or six transistors.

The non-volatility and DRAM-like speed of flash memory are ideal for solid-state "disk" drives. Flash based disks do not have any disks or moving parts. Flash disks are very fast compared to most available disk drives.

Data can be accessed in 120 nanoseconds in flash memories while it takes 15 to 30 milliseconds to access data stored in today's typical hard disk. However, flash disks are limited to up to 40 megabytes in capacity whereas hard disk drives can store from 5 megabytes to several gigabytes.

A flash disk can be built from one or more flash-memory IC chips and some controlling logic devices. For example, to build a 512Kbyte flash disk, four one-megabit flash memory chips can be connected on a small card. An example of such a flash memory system is the Intel iMC004FLKA 4 Megabyte flash memory card. In addition to flash-disk hardware, software to manage files on a flash disk is required. The file system software handles creating and deleting files, changing the file sizes and formatting the flash disk. Microsoft offers flash file system software for the MS-DOS operating system.

The most severe limitation of flash disks has been its cost.However, the cost of flash ROM is significantly decreasing. In the future, high density flash memory is expected to be available at an inexpensive cost.

Flash memory can be programmed using either 5V or 12V. The 5V feature becomes more desirable for portable equipment wℓ :e no 12V power is available. The speed, rugged construction, and lower power consumption of flash disks is ideal for laptop and notebook computers.

In summary, due to the high cost of flash disks, desktop computers will continue to use hard disk drives. Since flash memory combines the advantages of an EPROM's low cost with an EEPROM's ease of reprogramming, flash memories are being extensively used these days as a microcomputer's main non-volatile memory. An example of flash memory is the Intel 28F020 256K x 8 flash memory. By 1997, the cost of a megabyte of flash memory is expected to move from its current level of $120 to about $5. At that time, flash disks will be able to replace hard disks in many applications.

1.3.3.b Main Memory Array Design

In many applications, a memory of large capacity is often realized by interconnecting several small-size memory blocks. In this section, design of a large main memory using small-size memories as building blocks is presented. The memory map defining all memory addresses is determined. Note that the microprocessor's reset vector must be included in the memory map.

There are three types of techniques used for designing the main memory. These are linear decoding, full decoding/partial decoding and memory decoding using PALs. We will illustrate the concepts associated with these techniques in the following.

First, consider the block diagram of a typical static RAM chip shown in Figure 1.4.
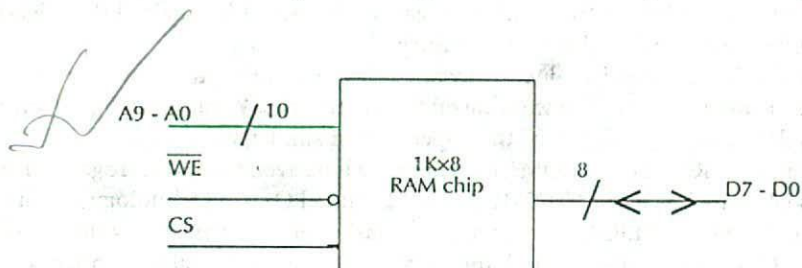

FIGURE 1.4   Typical Static RAM Chip.

The capacity of this chip is 8192 bits and these bits are organized as 1024 words with 8 bits/ word. Each word has a unique address and this is specified on 10-bit address lines A9—A0 (note that $2^{10} = 1024$). The inputs and outputs are routed through the 8-bit bidirectional data lines D7 through D0. The operation of this chip is governed by the two control inputs: WE (write Enable) and CS (chip select). The truth table that describes the operation of this chip is shown in Table 1.1.

TABLE 1.1   Truth Table for 1K × 8 Static RAM

CS	WE	MODE	Status of D7—D0	Power
L	X	Not selected	High impedance	Standby
H	L	Write	Acts as an input bus	Active
H	H	Read	Acts as an output bus	Active

Note: H — high, L — low, X — don't care.

From this table, it is easy to see that when CS input is low, the chip is not selected and thus the lines D7 through D0 are driven to the high impedance state. When CS = 1 and WE is LOW, data on lines D7—D0 are written into the word addressed by A0 through A9. Similarly, when CS = 1 and WE is high, the contents of the memory word (whose address is specified on address lines A9 through A0) will appear on lines D7 through D0. Note that when the chip select input CS goes to low, the device is disabled and the chip automatically reduces its power requirements and remains in this low-power standby mode as long as CS remains low. This feature results in system power savings as high as 85% in larger systems, where the majority of devices are disabled.

1.3.3.b.i Linear Decoding.   This technique uses the unused address lines of the microprocessor as chip selects for the memory chips. This method is used for small systems.

A simple way to connect an 8-bit microprocessor to a 6K RAM system using linear decoding is shown in Figure 1.5. In this approach, the address lines A9 through A0 of the microprocessor are used as a common input to each 1K × 8 RAM chip. The remaining 6 high-order lines are used to select one of the 6 RAM chips. For example, if A15A14A13A12A11A10 = 000010, then the RAM chip 1 is selected. The address map realized by this arrangement is summarized in Figure 1.6. This method is known as the linear select decoding technique. The principal advantage of this method is that it does not require any decoding hardware. However, this approach has some disadvantages:

- Although with a 16-bit address bus we have 64K bytes of RAM space, we are able to interface only 6K bytes of RAM. This means that this idea wastes address space.
- The address map is not contiguous; rather, it is sparsely distributed.
- If both A11 and A10 are high at the same time, both RAM chips 0 and 1 are selected and thus a bus conflict occurs. This can be avoided by proper programming to select the desired memory chip and deselect the others.
- Also, if all unused address lines are not utilized as chip selects for memory, then these unused pins become don't cares (can be 0 or 1). This results in foldback, meaning that a memory location will have its image in the memory map. For example, if A15 is don't care in design and if A14 to A0 address lines are used, then address 0000_{16} and address 8000_{16} are the same locations. This is called foldback and it wastes memory space.

1.3.3.b.ii Full/Partial Decoding.   Difficulties such as the bus conflict and sparse address distribution are eliminated by the use of the full/partial decoded addressing technique. To see this, consider the organization shown in Figure 1.7. In this setup, we use a 2-to-4 decoder and interface the 8-bit microprocessor with 4K bytes of RAM. In particular, the four combinations of the lines A11 and A10 select the RAM chips as follows:

A11	A10	Device Selected
0	0	RAM chip 0
0	1	RAM chip 1
1	0	RAM chip 2
1	1	RAM chip 3

Also observe that this hardware makes sure that the memory system is enabled only when the lines A15 through A12 are zero. The complete address map corresponding to this organization is summarized in Figure 1.8.

1.3.3.b.iii   *Memory Decoding by using Programmable Array Logic (PAL).*   A Programmable Array Logic (PAL) is similar to a ROM in concept except that it does not provide full decoding of the input lines. Instead, a PAL provides a partial sum of products which can be obtained via programming and saves a lot of space on the board. The PAL chip contains a fused programmable AND array and a fixed OR array. Note that in PLA (Programmable Logic Array) both AND and OR arrays are programmable. The AND and OR gates are fabricated inside the PAL without interconnections. The specific functions desired are implemented during programming via software. Programming of the PAL provides connections of the inputs of the AND gates and the outputs of the AND gates to the inputs of the OR gates. Therefore, the PAL implements sum of products of the inputs. PALs are used extensively these days with 32-bit microprocessors such as the Intel 80386/80486 and Motorola 68030/68040 for performing the memory decode function. PALs connect these microprocessors to memory, I/O and other chips without the use of any additional logic gates or circuits.

Each input has both true and complemented forms. A look at the NOR gate output $\varnothing_1$ indicates that there are two X connections at the inputs of this NOR gate. The three X inputs are wire-ANDed together with programming the PAL so that $\varnothing_1 = \left(I_0 \cdot \overline{I_1}\right) + I_2$.

The PAL chips are usually identified by a two-digit number followed by a letter and then a digit. The two-digit number specifies the number of input lines while the last digit defines the number of output lines. The fixed number of AND gates are connected to either an OR or a NOR gate. The letter 'H' indicates that the output gates are OR gates. The letter 'L' is used when the outputs are NOR gates.

As an example, the 10H8 provides eight OR gate outputs driven by ten AND gate inputs. The 10L8, on the other hand, is the same as the 10H8 except that the eight output gates are NOR gates.
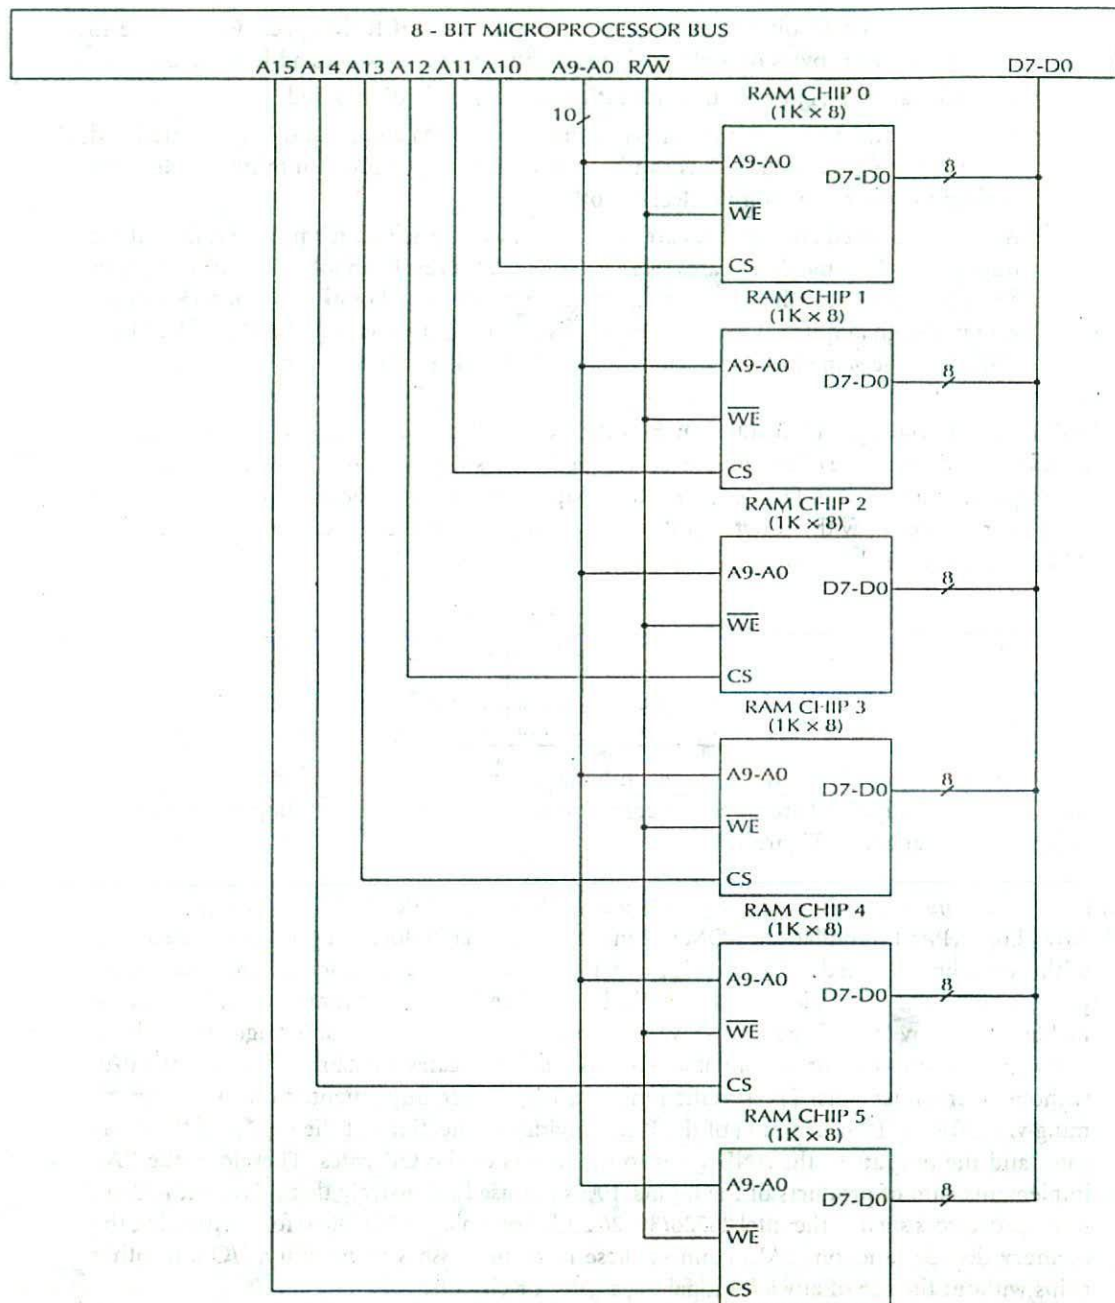
FIGURE 1.5   An 8-bit microprocessor interfaced to a 6K RAM system using the linear select decoding technique.

Some PALs provide additional features. The 16L8 includes tristate outputs. Each of the eight NOR gate outputs is driven internally by six AND gates. A seventh AND gate provides the enable signal for the tristate buffers.

The 16L8 is a popular PAL used with 32-bit microprocessors. The 16L8 is a 20-pin PAL manufactured using bipolar transistors. It has ten input pins (labeled I), two outputs (labeled O) and six programmable Input/Output (labeled I/O) lines. Using the programmable I/O lines, the number of input lines can be increased to a maximum of 16 and the number of output lines can be increased to 8.

Binary Address Pattern															Device Selected	Address Assignment in Hex
A15	A14	A13	A12	A11	A10	A9	A7	A6	A5	A4	A3	A2	A1	A0		
0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	RAM CHIP 0	0400 to 07FF
.	.	.	.	.	.	.	.	.	.	.	.	.	.	.		
0	0	0	0	0	1	1	1	1	1	1	1	1	1	1		
0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	RAM CHIP 1	0800 to 0BFF
.	.	.	.	.	.	.	.	.	.	.	.	.	.	.		
0	0	0	0	1	0	1	1	1	1	1	1	1	1	1		
0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	RAM CHIP 2	1000 to 13FF
.	.	.	.	.	.	.	.	.	.	.	.	.	.	.		
0	0	0	1	0	0	1	1	1	1	1	1	1	1	1		
0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	RAM CHIP 3	2000 to 23FF
.	.	.	.	.	.	.	.	.	.	.	.	.	.	.		
0	0	1	0	0	0	1	1	1	1	1	1	1	1	1		
0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	RAM CHIP 4	4000 to 43FF
.	.	.	.	.	.	.	.	.	.	.	.	.	.	.		
0	1	0	0	0	0	1	1	1	1	1	1	1	1	1		
1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	RAM CHIP 5	8000 to 83FF
.	.	.	.	.	.	.	.	.	.	.	.	.	.	.		
1	0	0	0	0	0	1	1	1	1	1	1	1	1	1		

FIGURE 1.6　Address map realized by the system shown in Figure 1.5.

Programming PALs can be accomplished by first creating a file by using a text editor on a personal computer. The file should include information such as the pin assignments of the PAL and the boolean equation for the outputs. By inserting the PAL into the programming module included with the personal computer, the PAL can then be programmed with the PAL programming software provided with the personal computer. Note that PAL programming hardware and software are sold separately and not usually included with a personal computer.

1.3.3.c Memory Management Concepts

Due to the massive amount of information that must be saved in most systems, the mass storage is often a disk. If each access is to a disk (even a hard disk), then system throughput will be reduced to unacceptable levels.

An obvious solution is to use a large and fast locally accessed semiconductor memory. Unfortunately the storage cost per bit for this solution is very high. A combination of both off-board disk (secondary memory) and on-board semiconductor main memory must be designed into a system. This requires a mechanism to manage the two-way flow of information between the primary (semiconductor) and secondary (disk) media. This mechanism must be able to transfer blocks of data efficiently, keep track of block usage, and replace them in a nonarbitrary way. The primary memory system must therefore be able to dynamically allocate memory space.

An operating system must have resource protection from corruption or abuse by users. Users must be able to protect areas of code from each other, while maintaining the ability to communicate and share other areas of code. All these requirements indicate the need for a device, located between the microprocessor and memory, to control accesses, perform address mappings, and act as an interface between the logical (programmer's memory) and microprocessor physical (memory) address spaces. Since this device must manage memory use, it is appropriately called the memory management unit (MMU). Typical 32-bit microprocessors such as Motorola 68030 and Intel 80386 include on-chip MMU.

The MMU reduces the burden of the memory management function on the operating system.

The basic functions provided by the MMU are address translation and protection. The MMU translates logical program addresses to physical memory addresses. The addresses in a
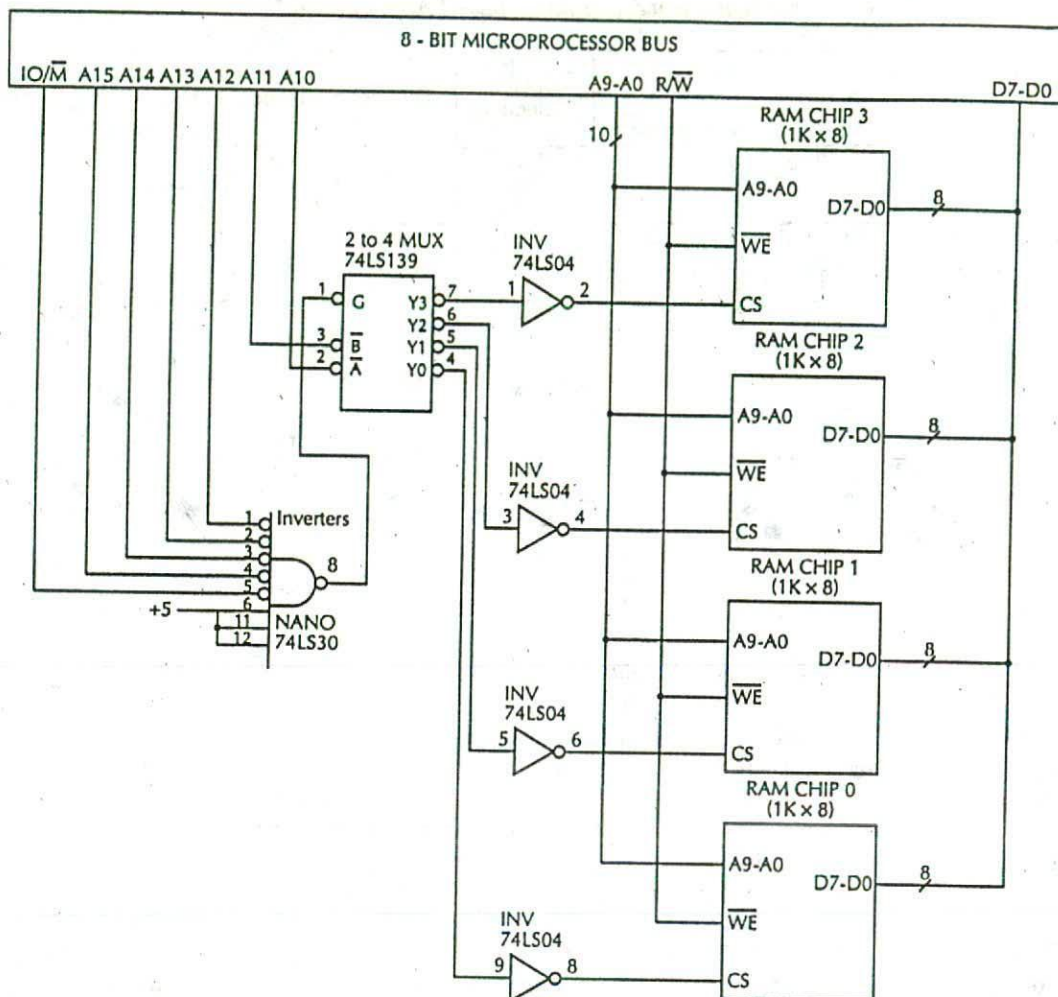
FIGURE 1.7   An 8-bit microprocessor interfaced to a 4K RAM system using a full/partial decoded addressing technique

Binary Address Pattern																Device Selected	Address Assignment in Hex
A15	A14	A13	A12	A11	A10	A9	A7	A6	A5	A4	A3	A2	A1	A0			
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0		RAM CHIP 0	0000 to 03FF
0	0	0	0	0	0	1	1	1	1	1	1	1	1	1			
0	0	0	0	0	1	0	0	0	0	0	0	0	0	0		RAM CHIP 1	0400 to 07FF
0	0	0	0	0	1	1	1	1	1	1	1	1	1	1			
0	0	0	0	1	0	0	0	0	0	0	0	0	0	0		RAM CHIP 2	0800 to 0BFF
0	0	0	0	1	0	1	1	1	1	1	1	1	1	1			
0	0	0	0	1	1	0	0	0	0	0	0	0	0	0		RAM CHIP 3	0C00 to 0FFF
0	0	0	0	1	1	1	1	1	1	1	1	1	1	1			

FIGURE 1.8   Address map corresponding to the organization shown in Figure 1.7.

FIGURE 1.9   A typical PAL.

$$O_1 = (I_0 \cdot \bar{I_1}) + I_2$$
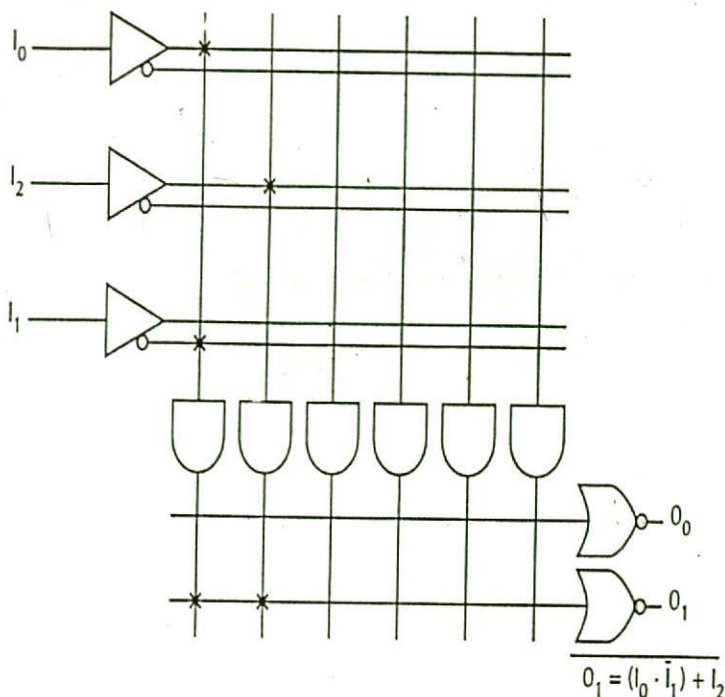
program are called logical addresses since they indicate the logical positions of instructions and data. The MMU translates these logical addresses to physical addresses provided by the memory chips. The MMU can perform address translation in one of two ways:

1. By using the substitution technique as shown in Figure 1.10a
2. By adding an offset to each logical address to obtain the corresponding physical address as shown in Figure 1.10b

Address translation using substitution is faster than the offset method. However, the offset method has the advantage of mapping a logical address to any physical address as determined by the offset value.

Memory is usually divided into small manageable units. The terms "page" and "segment" are frequently used to describe these units. Paging divides the memory into equal-sized pages, while segmentation divides the memory into variable-sized segments.

It is relatively easier to implement the address translation table if the logical and main memory spaces are divided into pages. The term "page" is associated with logical address space, while the term "block" usually refers to a page in main memory space.

There are three ways to map logical addresses to physical addresses. These are paging, segmentation, and combined paging/segmentation.

In a paged system, a user has access to a larger address space than physical memory provides. The virtual memory system is managed by both hardware and software. Note that memory in excess of the main memory such as floppy disk storage is called virtual memory. The hardware included in the memory management unit handles address translation. The memory management software in the operating system performs all functions including page replacement policies in order to provide efficient memory utilization. The memory management software performs functions such as removal of the desired page from main memory to accommodate a new page, transferring a new page from secondary to main memory at the right instant of time, and placing the page at the right location in memory.
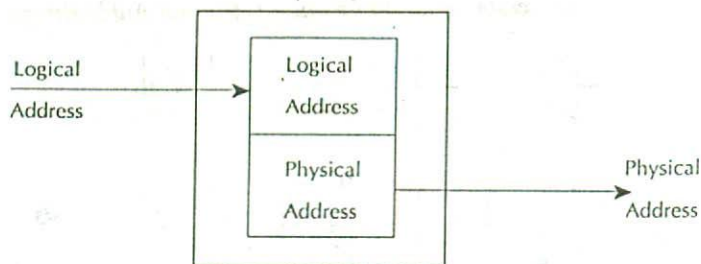
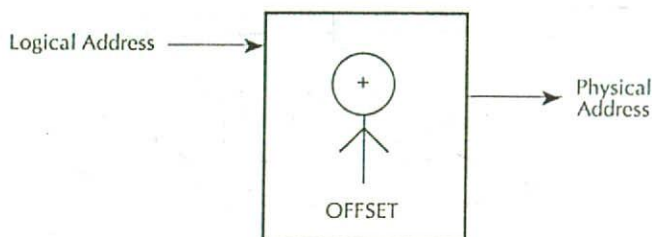FIGURE 1.10a   Address translation using substitution technique.

FIGURE 1.10b   Address translation by offset technique.

If the main memory is full during transfer from secondary to main memory, it is necessary to remove a page from main memory to accommodate the new page. Two popular page replacement policies are first-in first-out (FIFO) and least recently used (LRU). The FIFO policy removes the page from main memory that has been resident in memory for the longest amount of time. The FIFO replacement policy is easy to implement. One of the main disadvantages of the FIFO policy is that it is likely to replace heavily used pages. Note that heavily used pages are resident in main memory for the longest amount of time. Sometimes this replacement policy might be a poor choice. For example, in a time-shared system, several users normally share a copy of the text editor in order to type and correct programs. The FIFO policy on such a system might replace a heavily used editor program page to make room for a new page. This program page might be recalled to main memory immediately. The FIFO, in this case, would be a poor choice.

The LRU policy, on the other hand, replaces that page which has not been used for the longest amount of time.

In the segmentation method, the MMU utilizes the segment selector to obtain a descriptor from a table in memory containing several descriptors. A descriptor contains the physical base address for a segment, the segment's privilege level, and some control bits. When the MMU obtains a logical address from the microprocessor, it first determines whether the segment is already in the physical memory. If it is, the MMU adds an offset component to the segment base component of the address obtained from the segment descriptor table to provide the physical address. The MMU then generates the physical address on the address bus for accessing the memory. On the other hand, if the MMU does not find the logical address in physical memory, it interrupts the microprocessor. The microprocessor executes a service routine to bring the desired program from a secondary memory such as disk to the physical memory. The MMU determines the physical address using the segment offset and descriptor as above and then generates the physical address on the address bus for memory. A segment will usually consist of an integral number of pages, say, each 256 bytes long. With different-sized segments being swapped in and out, areas of valuable primary memory can become unusable. Memory is unusable for segmentation when it is sandwiched between already allocated segments and if it is not large enough to hold the latest segment that needs to be

loaded. This is called external fragmentation and is handled by MMUs using special techniques. An example of external fragmentation is given in Figure 1.11. The advantages of segmented memory management are that few descriptors are required for large programs or data spaces, and internal fragmentation (to be discussed later) is minimized. The disadvantages include external fragmentation, involved algorithms for placing data are required, possible restrictions on starting address, and longer data swap times are required to support virtual memory.
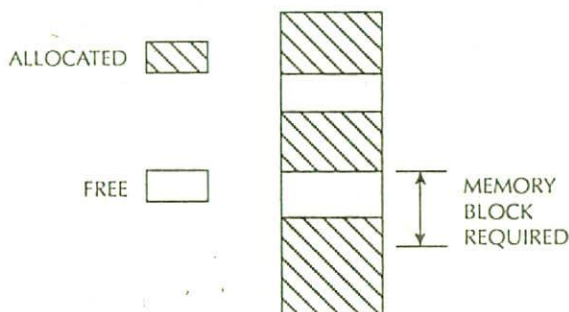

FIGURE 1.11   Memory fragmentation (external).

Address translation using descriptor tables offers a protection feature. A segment or a page can be protected from access by a program section of a lower privilege level. For example, the selector component of each logical address includes one or two bits indicating the privilege level of the program requesting access to a segment. Each segment descriptor also includes one or two bits providing the privilege level of that segment. When an executing program tries to access a segment, the MMU can compare the selector privilege level with the descriptor privilege level. If the segment selector has the same or higher privilege level, then the MMU permits the access. If the privilege level of the selector is lower than the descriptor, the MMU can interrupt the microprocessor informing of a privilege level violation. Therefore, the indirect technique of generating physical address provides a mechanism of protecting critical program sections in the operating system.

Paging divides the memory into equal-sized pages, it avoids the major problem of segmentation-external fragmentation. Since the pages are of the same size, when a new page is requested and an old one swapped out, the new one will always fit into the vacated space. However, a problem common to both techniques remains — internal fragmentation. Internal fragmentation is a condition where memory is unused but allocated due to memory block size implementation restrictions. This occurs when a module needs, say, 300 bytes and page is 1K bytes, as shown in Figure 1.12.

In the paged-segmentation method, each segment contains a number of pages. The logical address is divided into three components: segment, page, and word. The segment component defines a segment number, the page component defines the page within the segment, and the word component provides the particular word within the page. A page component of n bits can provide up to 2^n pages. A segment can be assigned with one or more pages up to a maximum of 2^n pages; therefore, a segment size depends on the number of pages assigned to it.

Protection mechanisms can operate either on physical address or logical address. Physical memory protection can be accomplished by using one or more protection bits with each block to define the access type permitted on the block. This means that each time a page is transferred from one block to another, the block protection bits must be updated. A more efficient approach is to provide a protection feature in logical address space by including protection bits in the descriptors of the segment table in the MMU.
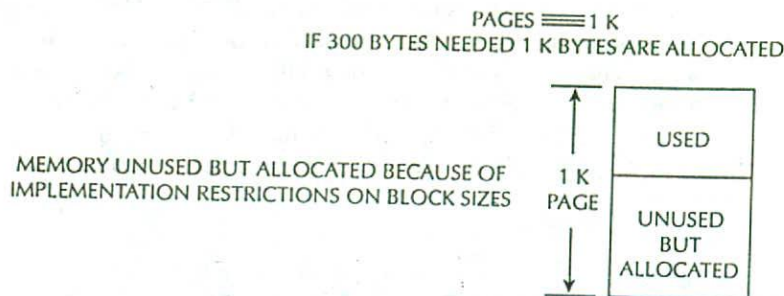
PAGES ≡≡≡ 1 K
IF 300 BYTES NEEDED 1 K BYTES ARE ALLOCATED

MEMORY UNUSED BUT ALLOCATED BECAUSE OF
IMPLEMENTATION RESTRICTIONS ON BLOCK SIZES

| | USED |
| 1 K PAGE | UNUSED BUT ALLOCATED |

FIGURE 1.12   Memory fragmentation (internal).

1.3.3.d Cache Memory Organization

The performance of a microcomputer system can be significantly improved by introducing a small, expensive, but fast memory between the microprocessor and main memory. This memory is called cache memory and this idea was first introduced in the IBM 360/85 computer. Later on, this concept was also implemented in minicomputers such as PDP-11/70. With the advent of VLSI technology, the cache memory technique is gaining acceptance in the microprocessor world. For example, an on-chip cache memory is implemented in Intel's 32-bit microprocessor, the 80486, and Motorola's 32-bit microprocessors, the MC 68020/68030/68040. The 80386 does not have on-chip cache but external cache memory can be interfaced to it. Studies have shown that typical programs spend most of their execution times in loops. This means that the addresses generated by a microprocessor have a tendency to cluster around a small region in the main memory. This phenomenon is known as locality of reference. The 32-bit microprocessor can execute the same instructions in a loop from the on-chip cache rather than reading them repeatedly from the external main memory. Thus the performance offered by 32-bit microprocessors is greatly improved.

The block diagram representation of a microprocessor system that employs an on-chip cache memory is shown in Figure 1.13. Usually, a cache memory is very small in size and its access time is less than that of the main memory by a factor of 5.

Cache hit means that the reference is found in the cache and the data pertaining to the microprocessor reference is transferred to the microprocessor from the cache. However, if the reference is not found in the cache, we call it a cache miss. When there is a cache miss, the main memory is accessed by the microprocessor and the data are then transferred to the microprocessor from the main memory. At the same time, a block of data containing the desired data needed by the microprocessor is transferred from the main memory to the cache. The block normally contains 4 to 16 bytes, and this block is placed in the cache using the standard replacement policies such as FIFO (First-In First-Out) or LRU (Least Recently Used). This block transfer is done with a hope that all future references made by the microprocessor will be confined to the fast cache.

The relationship between the cache and main memory blocks is established using mapping techniques. Three widely used mapping techniques are

- Direct mapping
- Fully-associative mapping
- Set-associative mapping

In direct mapping, the main memory address is divided into two fields: an index field and a tag field. The number of bits in the index field is equal to the number of address bits required to access the cache memory.

Assume that the main memory address is m bits wide and the cache memory address is n bits wide. Then the index field requires n bits and the tag field is (m — n) bits wide. The n-
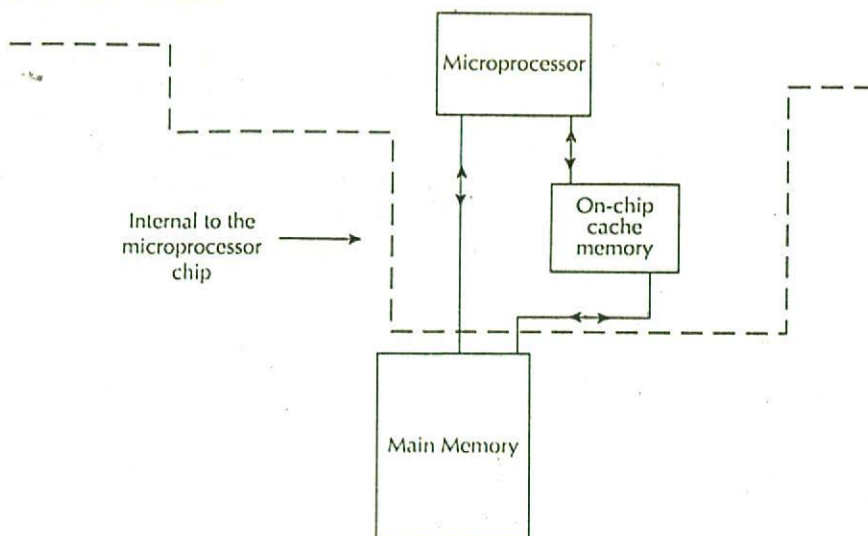
FIGURE 1.13   Memory organization of a computer system that employs a cache memory.

bit address accesses the code. Each word in the cache includes the data word and its associated tag. When the microprocessor generates an address for main memory, the index field is used as the address to access the cache. The tag field of the main memory is compared with the tag field in the word read from cache. A hit occurs if the tags match. This means that the desired data word is in cache. A miss occurs if there is no match, and then the required word is read from main memory. It is written in the cache along with the new tag. A random access memory is used as the cache memory.

One of the main drawbacks of direct mapping is that numerous misses may occur if two or more words with addresses having the same index but different tags are accessed several times. This can be minimized by incorporating a larger cache.

The fastest cache memory utilizes an associative memory. The method is known as fully associative mapping. Each associative memory content contains main memory address and its content (data). When the microprocessor generates a main memory address, it is compared associatively (simultaneously) with all addresses in the associative memory. If there is a match, the corresponding data word is read from the associative cache memory and sent to the microprocessor. If a miss occurs, the main memory is accessed, and the address and its corresponding data are written to the associative cache memory. If the cache is full, certain policies such as FIFO (first-in first-out) are used as replacement algorithm for the cache. The associative cache is expensive but provides fast operation.

The set-associative mapping is a combination of direct and associative mapping. Each cache word stores two or more main memory words using the same index address. Each main memory word consists of a tag and its data word. An index with two or more tags and data words forms a set. When the microprocessor generates a memory request, the index of the main memory address is used as the cache address. The tag field of the main memory address is then compared associatively (simultaneously) with all tags stored under the index. If a match occurs, the desired data word is read. If a miss occurs, the data word, along with its tag, is read from main memory and also written into the cache. The hit ratio improves as the set size increases. This is because more words with the same index but different tags can be stored in cache.

There are two ways of writing into cache: the write-back and write-through methods. In the write-back method, whenever the microprocessor writes something into a cache word, a dirty bit is assigned to the cache word. When a dirty word is to be replaced with a new word, the

dirty word is first copied into the main memory before it is overwritten by the incoming new word. The advantage of this method is that it avoids unnecessary writing into main memory.

In the write-through method, whenever the microprocessor alters cache data, the same alteration is made in the main memory copy of the altered cache data. This policy can be easily implemented and also it insures that the contents of the main memory are always valid. This feature is desirable in a multiprocessor system where the main memory is shared by several processors. However, this approach may lead to several unnecessary writes to main memory.

One of the important aspects of cache memory organization is to devise a method that insures proper utilization of the cache. Usually, the tag directory contains an extra bit for each entry. This additional bit is called a valid bit. When the power is turned on, the valid bit corresponding to each cache block entry of the tag directory is reset to zero. This is done in order to indicate that the cache block holds invalid data. When a block of data is first transferred from the main memory to a cache block, the valid bit corresponding to this cache block is set to 1. In this arrangement, whenever the valid is a zero, it implies that a new incoming block can overwrite the existing cache block. Thus, there is no need to copy the contents of the cache block being replaced into the main memory.

1.3.4 Input/Output (I/O)

This section describes the basic input and output techniques used by microcomputers to transfer data between the microcomputer and external devices. The general characteristics of I/O are described. One communicates with a microcomputer system via the I/O devices interfaced to it. The user can enter programs and data using the keyboard on a terminal and execute the programs to obtain results. Therefore, the I/O devices connected to a microcomputer system provide an efficient means of communication between the computer and the outside world. These I/O devices are commonly called peripherals and include keyboards, CRT displays, printers, and disks.

The characteristics of the I/O devices are normally different from those of the microcomputer. For example, the speed of operation of the peripherals is usually slower compared to the microcomputer, and the word length of the microcomputer may be different from the data format of the peripheral device. To make the characteristics of the I/O devices compatible with those of the microcomputer, interface hardware circuitry between the microcomputer and I/O devices is necessary.

In a typical microcomputer system, the user gets involved with two types of I/O devices: physical I/O and logical I/O. When the microcomputer has no operating system, the user must work directly with physical I/O devices and perform detailed I/O design.

There are three ways of transferring data between the microcomputer and a physical I/O device:

- Programmed I/O
- Interrupt driven I/O
- Direct memory access (DMA)

The microcomputer executes a program to communicate with an external device via a register called the I/O port for programmed I/O.

An external device requests the microcomputer to transfer data by activating a signal on the microcomputer's interrupt line during interrupt I/O. In response, the microcomputer executes a program called the interrupt-service routine to carry out the function desired by the external device, again by way of one or more I/O ports.

Data transfer between the microcomputer's memory and an external device occurs without microprocessor involvement with direct memory access.

For a microcomputer with an operating system, the user works with virtual I/O devices. The user does not have to be familiar with the characteristics of the physical I/O devices. Instead, the user performs data transfers between the microcomputer and the physical I/O devices indirectly by calling the I/O routines provided by the operating system using virtual I/O instructions. This is called logical I/O.

1.3.4.a Programmed I/O

As described earlier, the microcomputer communicates with an external device via one or more registers called I/O ports using programmed I/O. These I/O ports are occasionally fabricated by the manufacturer in the same chip as the memory chip to achieve minimum chip count for small system applications. For example, the Intel 8355/8755 contains 2K bytes of ROM/EPROM with two I/O ports. The Motorola 6846 has 2K bytes of ROM and an 8-bit I/O port.

I/O ports are usually of two types. For one type, each bit in the port can be individually configured as either input or output. For the other type, all bits in the port can be set up as either all parallel input or output bits. Each port can be configured as an input or output port by another register called the command, or data-direction register. The port data register contains the actual input or output data. The data-direction register is an output register and can be used to configure the bits in the port as inputs or outputs.

Each bit in the port can usually be set up as an input or output by respectively writing a 0 or a 1 in the corresponding bit of the data-direction register (DDR). A bidirectional buffer (one input buffer and one output buffer) is connected at each bit of the port. A '1' written to a particular bit in DDR enables the output buffer while a '0' enables the input buffer connected at the corresponding bit of the port. As an example, if an 8-bit data-direction register contains 34_{16}, then the corresponding port is defined as follows:

7	6	5	4	3	2	1	0	Bit position
0	0	1	1	0	1	0	0	Data-direction register
								I/O port

↑ ↑ ↓ ↓ ↑ ↓ ↑ ↑
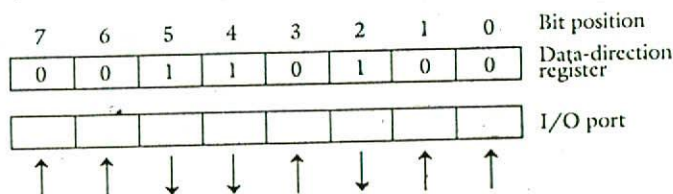
In the preceding example, since 34_{16} ($0011\ 0100_2$) is sent as an output into the data-direction register, bits 0, 1, 3, 6, and 7 of the port are set up as inputs, and bits 2, 4, and 5 of the port are defined as outputs. The microcomputer can then send outputs to external devices, such as LEDs, connected to bits 2, 4, and 5 through a proper interface. Similarly, the microcomputer can input the status of external devices, such as switches, through bits 0, 1, 3, 6, and 7. To input data from the input switches, the 8-bit microcomputer assumed here inputs the complete byte, including the bits to which LEDs are connected. While receiving input data from an I/O port, however, the microcomputer places a value, probably 0, at the bits configured as outputs and the program must interpret them as "don't cares". At the same time, the microcomputer's outputs to bits configured as inputs are disregarded.

For parallel I/O, there is only one register, known as the command register, for all ports. A particular bit in the command register configures all bits in a port as either inputs or outputs.

Consider two I/O ports in an I/O chip along with one command register. Assume that a 0 or a 1 in a particular bit position defines all bits of ports A or B as inputs or outputs.

For example,
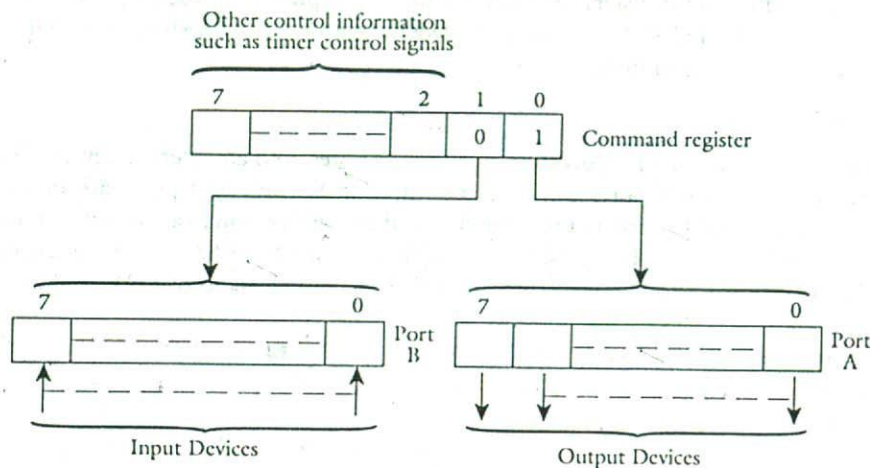
Some I/O ports are called handshake ports. Data transfer occurs via these ports through exchanging of control signals between the I/O controller and an external device.

1.3.4.b Standard I/O Versus Memory-Mapped I/O

I/O ports are addressed using either standard I/O or memory-mapped I/O techniques. The standard I/O, also called isolated I/O, uses the $IO/\overline{M}$ control pin on the microprocessor chip. The processor outputs a HIGH on this pin to indicate to memory and the I/O chips that an I/O operation is taking place. A LOW output from the processor to this pin indicates a memory operation. Execution of IN or OUT instructions makes the $IO/\overline{M}$ HIGH, whereas memory-oriented instructions, such as LDA and STA, drive the IO/M to LOW. In standard I/O, the processor uses the $IO/\overline{M}$ pin to distinguish between I/O and memory. For 8-bit microprocessors, an 8-bit address is typically used for each I/O port. This is because 8 bits are the basic data unit for these processors. Eight-bit processors are usually capable of directly addressing 64K bytes of memory using 16 address lines. With an 8-bit I/O port address, these processors are capable of addressing 256 ports. However, in a typical application, there are usually four or five I/O ports required. Some of the address bits of the microprocessor are normally decoded to obtain the I/O port addresses. With memory-mapped I/O, the processor does not differentiate between I/O and memory and, therefore, does not use the microprocessor's $IO/\overline{M}$ control pin. The microprocessor uses the memory addresses (which may not exist in the microcomputer's physical memory) to represent I/O ports. The I/O ports are mapped into the microprocessor's main memory and, hence, are called *memory-mapped I/O*.

In memory-mapped I/O, the most significant bit (MSB) of the address may be used to distinguish between I/O and memory. If the MSB of address is 1, an I/O port is selected. If the MSB of address is 0, a memory location is accessed. This reduces the microprocessor's addressing memory (main memory) by 50%. Sixteen and thirty-two bit microprocessors provide special control signals for performing memory-mapped I/O. Thus, these processors do not use MSB of the address lines. Intel microprocessors can use either standard or memory-mapped I/O while Motorola microprocessors use only memory-mapped I/O. For example, with standard I/O, Intel 8086 uses IN AL, PortA, and OUT PortA, AL for inputing and outputing data. On the other hand, the 8086 uses memory-oriented instructions such as MOV AL, START, and MOV START, AL for inputing and outputing data. Note that START is an

I/O port mapped as a memory address. Motorola, on the other hand, does not have any IN or OUT instructions and uses memory-oriented instructions for I/O operation.

1.3.4.c Unconditional and Conditional Programmed I/O

The microprocessor can send data to an external device at any time during unconditional I/O. The external device must always be ready for data transfer. A typical example is when the processor outputs a 7-bit code through an I/O port to drive a seven-segment display connected to this port.

In conditional I/O, the microprocessor outputs data to an external device via handshaking. Data transfer occurs by the exchanging of control signals between the microprocessor and an external device. The microprocessor inputs the status of the external device to determine whether the device is ready for data transfer. Data transfer takes place when the device is ready.

The concept of conditional I/O will now be demonstrated by means of data transfer between a microprocessor and an analog-to-digital (A/D) converter. Consider, for example, the A/D converter shown in the accompanying figure.
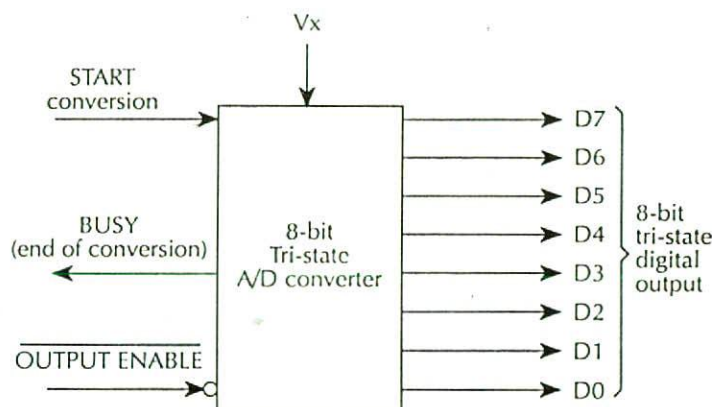
The A/D converter just shown transforms an analog voltage Vx into an 8-bit binary output at pins D7—D0. A pulse at the START conversion pin initiates the conversion. This drives the BUSY signal LOW. The signal stays LOW during the conversion process. The BUSY signal goes HIGH as soon as the conversion ends. Since the A/D converter's output is tristated, a LOW on the $\overline{\text{OUTPUT ENABLE}}$ transfers the converter's output. A HIGH on the $\overline{\text{OUTPUT ENABLE}}$ drives the converter's output to a high impedance state.

The concept of conditional I/O can be demonstrated by interfacing the A/D converter to an 8-bit processor. Figure 1.14 shows such an interfacing example.

The user writes a program to carry out the conversion process. When this program is executed, the processor sends a pulse to the START pin of the converter via bit 2 of port A. The microprocessor then checks the BUSY signal by bit 1 of port A to determine if the conversion is completed. If the BUSY signal is HIGH (indicating the end of conversion), the microprocessor sends a LOW to the $\overline{\text{OUTPUT ENABLE}}$ pin of the A/D converter. The microprocessor then inputs the converter's D0—D7 outputs via port B. If the conversion is not completed, the microprocessor waits in a loop checking for the BUSY signal to go HIGH.

1.3.4.d Typical Microcomputer Output Circuit

The microcomputer designer is often concerned with the output circuit because of the microcomputer's small output current drive capability. The tristate output circuit shown in Figure 1.15 typically is utilized by a microcomputer as the output circuit. This circuit uses totem pole-type output called a PUSH-PULL circuit providing low-output currents. Therefore, a current amplifier (buffer) is required to drive devices such as LEDs.
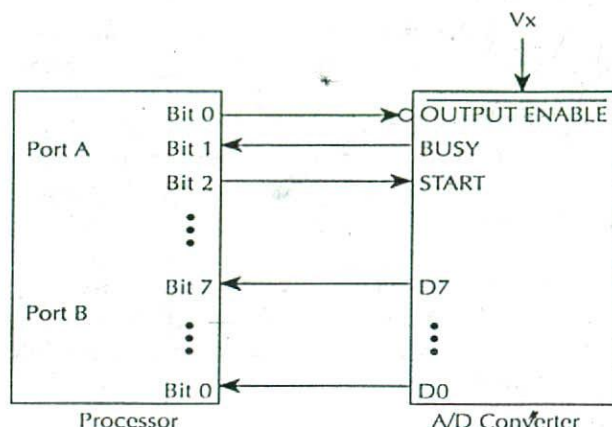
FIGURE 1.14   Interfacing an A/D converter to an 8-bit processor.
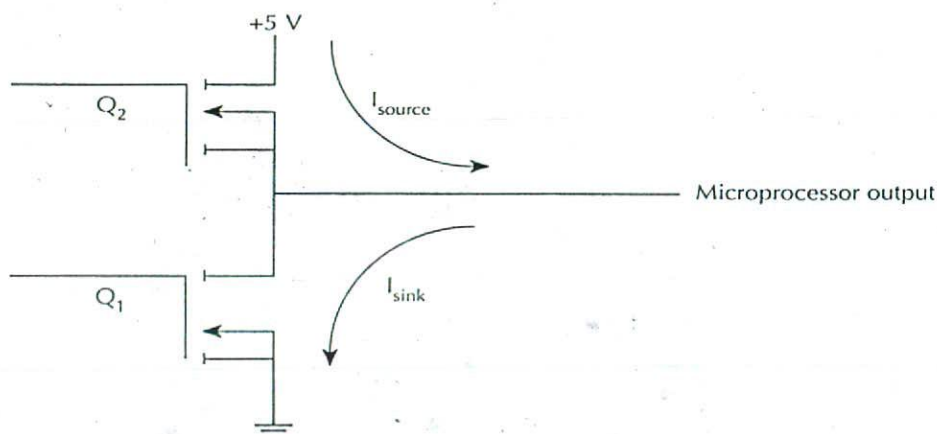
FIGURE 1.15   Digital microprocessor output circuit.

In the preceding figure, when Q_1 is ON, Q_2 is OFF and I_{sink} will flow from the external device into Q_1. Also, when Q_2 is ON, Q_1 is OFF and I_{source} will flow from Q_2 into the output device. Figure 1.16 shows a hardware interface circuit using the push-pull circuit for driving an external device such as LED.

Assume I_{source} to be 400 μA (usually represented by a negative sign such as $I_{OH} = -400$ μA; the negative sign indicates that the chip is losing current) with a minimum voltage V_A of 2.4 at point A, and that the LED requires 10 mA at 1.7 V. Therefore, a buffer such as a transistor is required at the output circuit to increase the current drive capability to drive the LED. In order to design the interface, the values of R_1, R_2 and minimum β of the transistor will be determined in the following:

$$R_1 = \frac{V_A - V_{BE}(Q_3)}{400 \text{ μA}} = \frac{2.4 - 0.7}{400 \text{ μA}} = 4.25 \text{ k}\Omega$$

Since $I_{LED} = 10$ μA at 1.7 V and assuming that V_{CE} (saturation) = 0 V,

$$R_2 = \frac{5 - 1.7 - V_{CE}(Q_3)}{10 \text{ mA}} = \frac{3.3}{10 \text{ mA}} = 330 \text{ }\Omega$$
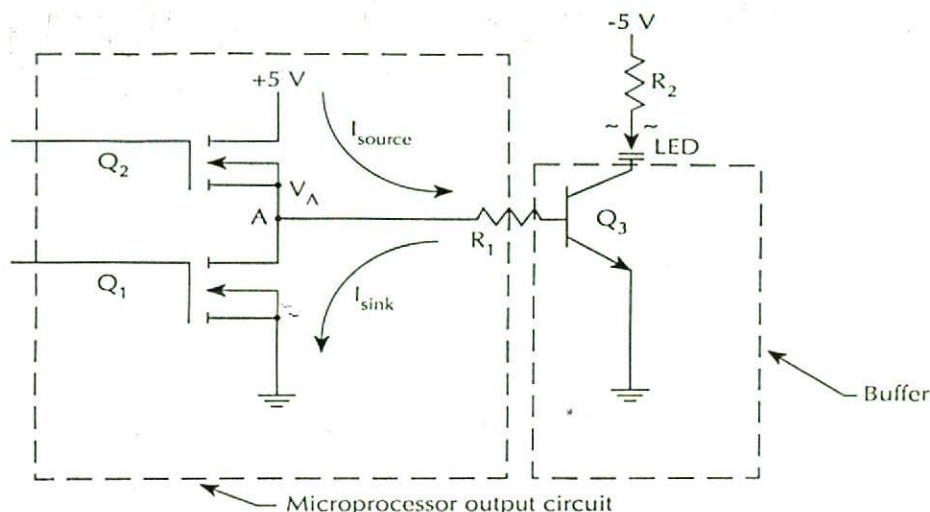
FIGURE 1.16   Microprocessor I/O interfaces for driving an LED.

Since $I_{source} = 400\ \mu A = I_B(Q_3)$, β for transistor Q_3 is

$$\beta = \frac{I_C(Q_3)}{I_B(Q_3)} = \frac{10\ mA}{400\ \mu A} = \frac{10 \times 10^{-3}}{400 \times 10^{-6}} = 25$$

Therefore, the interface design is complete, and a transistor with a mimimum of saturation β of 25 and $R_1 = 4.25\ k\Omega$ and $R_2 = 330\ \Omega$ is required. Note that a MOS outputs more sink current than source current. If sink current is used, the LED in Figure 1.16 can directly be connected to the microcomputer's output through an appropriate resistance. However, if the resistor value is not large enough, it may damage the transistor Q1.

1.3.4.e Interrupt Driven I/O

A disadvantage of conditional programmed I/O is that the microcomputer needs to check the status bit (BUSY signal for the A/D converter) by waiting in a loop. This type of I/O transfer is dependent on the speed of the external device. For a slow device, this waiting may slow down the capability of the microprocessor to process other data. The polled I/O and interrupt I/O techniques are efficient in this type of situation.

Interrupt I/O is a device-initiated I/O transfer. The external device is connected to a pin called the *interrupt* (INT) pin on the microprocessor chip. When the device needs an I/O transfer with the microcomputer, it activates the interrupt pin of the microprocessor chip. The microcomputer usually completes the current instruction and saves at least the contents of the current program counter on the stack.

The microcomputer then automatically loads an address into the program counter to branch to a subroutine-like program called the interrupt-service routine. This program is written by the user. The external device wants the microcomputer to execute this program to transfer data. The last instruction of the service routine is a RETURN, which is typically the same instruction used at the end of a subroutine. This instruction normally loads the return address (saved in the stack before going to the service routine) in the program counter. Then, the microcomputer continues executing the main program. Note that subroutines and interrupts are handled in a similar way: subroutines are initiated via execution of an instruction such as subroutine CALL while interrupts are initiated via activation of the microprocessor's interrupt pin by the interrupting device.

1.3.4.e.i *Interrupt Types.* There are typically three types of interrupts: external interrupts, traps or internal interrupts, and software interrupts.

External interrupts are initiated through the microcomputer's interrupt pins by external devices such as A/D converters. External interrupts can further be divided into two types: maskable and nonmaskable. A maskable interrupt is enabled or disabled by executing instructions such as EI or DI. If the microcomputer's interrupt is disabled, the microcomputer ignores the maskable interrupt. Some microprocessors, such as the Intel 8086, have an interrupt-flag bit in the processor status register. When the interrupt is disabled, the interrupt-flag bit is 1, so no maskable interrupts are recognized by the processor. The interrupt-flag bit resets to zero when the interrupt is enabled.

The nonmaskable interrupt has higher priority than the maskable interrupt. If both maskable and nonmaskable interrupts are activated at the same time, the processor will service the nonmaskable interrupt first. The nonmaskable interrupt is typically used as a power failure interrupt. Microprocessors normally use +5 V DC, which is transformed from 110 V AC. If the power falls below 90 V AC, the DC voltage of +5 V cannot be maintained. However, it will take a few milliseconds before the AC power can drop this low (below 90 V AC). In these few milliseconds, the power failure-sensing circuitry can interrupt the microprocessor. An interrupt service routine can be written to store critical data in nonvolatile memory such as battery-backed CMOS RAM. The interrupted program can continue without any loss of data when the power returns.

Some microprocessors are provided with a maskable handshake interrupt. This interrupt is usually implemented by using two pins — INTR and INTA. When the INTR pin is activated by an external device, the processor completes the current instruction, saves at least the current program counter onto stack, and generates an interrupt acknowledge $\left(\overline{INTA}\right)$. In response to the $\overline{INTA}$, the external device provides an instruction, such as CALL, using external hardware on the data bus of the microcomputer. This instruction is then read and executed by the microcomputer to branch to the desired service routine.

Internal interrupts, or traps, are activated internally by exceptional conditions such as overflow, division by zero, or execution of an illegal op-code. Traps are handled the same way as external interrupts. The user writes a service routine to take corrective measures and provide an indication to inform the user that an exceptional condition has occurred.

Many microprocessors include software interrupts, or system calls. When one of these instructions is executed, the microprocessor is interrupted and serviced similarly to external or internal interrupts. Software interrupt instructions are normally used to call the operating system. These instructions are shorter than subroutine calls, and no calling program is needed to know the operating system's address in memory. Software interrupt instructions allow the user to switch from user to supervisor mode. For some microprocessors, a software interrupt is the only way to call the operating system, since a subroutine call to an address in the operating system is not allowed.

1.3.4.e.ii *Interrupt Address Vector.* The technique used to find the starting address of the service routine (commonly known as the *interrupt address vector*) varies from one microprocessor to another. With some microprocessors, the manufacturers define the fixed starting address for each interrupt. Other manufacturers use an indirect approach by defining fixed locations where the interrupt address vector is stored.

1.3.4.e.iii *Saving the Microprocessor Registers.* When a microprocessor is interrupted, it saves at least the program counter on the stack so the microprocessor can return to the main program after executing the service routine. Some microprocessors save only one or two registers, such as the program counter and status register. Some other microprocessors save all microprocessor registers before going to the service routine. The user should know the specific registers the

microprocessor saves prior to executing the service routine. This will enable the user to use the appropriate return instruction at the end of the service routine to restore the original conditions upon return to the main program.

1.3.4.e.ivInterrupt Priorities. A microprocessor is typically provided with one or more interrupt pins on the chip. Therefore, a special mechanism is necessary to handle interrupts from several devices that share one of these interrupt lines. There are two ways of servicing multiple interrupts: polled and daisy chain techniques.

Polled interrupts are handled by software and therefore are slower when compared with daisy chaining. The processor responds to an interrupt by executing one general service routine for all devices. The priorities of devices are determined by the order in which the routine polls each device. The processor checks the status of each device in the general service routine, starting with the highest-priority device to service an interrupt. Once the processor determines the source of the interrupt, it branches to the service routine for the device.

In a daisy chain priority system, devices are connected in a daisy chain fashion to set up a priority system. Suppose one or more devices interrupt the processor. In response, the processor pushes at least the PC and generates an interrupt acknowledge (INTA) signal to the highest-priority device. If this device has generated the interrupt, it will accept the INTA. Otherwise it will pass the INTA onto the next device until INTA is accepted. Once accepted, the device provides a means for the processor to find an interrupt address vector by using external hardware. The daisy chain priority scheme is based on mostly hardware and is therefore faster than the polled interrupt.

1.3.4.f Direct Memory Access (DMA)

Direct Memory Access (DMA) is a technique that transfers data between a microcomputer's memory and an I/O device without involving the microprocessor. DMA is widely used in transferring large blocks of data between a peripheral device and the microcomputer's memory. The DMA technique uses a DMA controller chip for the data-transfer operation. The main functions of a typical DMA controller are summarized as follows:

- The I/O devices request DMA operation via the DMA request lines of the controller chip.
- The controller chip activates the microprocessor HOLD pin, requesting the CPU to release the bus.
- The processor sends HLDA (hold acknowledge) back to the DMA controller, indicating that the bus is disabled. The DMA controller places the current value of its internal registers, such as the address register and counter, on the system bus and sends a DMA acknowledge to the peripheral device. The DMA controller completes the DMA transfer and releases the buses.

There are three basic types of DMA: block transfer, cycle stealing, and interleaved DMA.

For block-transfer DMA, the DMA controller chip takes the bus from the microcomputer to transfer data between the memory and I/O device. The microprocessor has no access to the bus until the transfer is completed. During this time, the microprocessor can perform internal operations that do not need the bus. This method is popular with microprocessors. Using this technique, blocks of data can be transferred.

Data transfer between the microcomputer memory and an I/O device occurs on a word-by-word basis with cycle stealing. Typically, the microprocessor clock is enabled by ANDing an INHIBIT signal with the system clock. The system clock has the same frequency as the microprocessor clock.

The DMA controller controls the INHIBIT line. During normal operation, the INHIBIT line is HIGH, providing the microprocessor clock. When DMA operation is desired, the controller

makes the $\overline{\text{INHIBIT}}$ line LOW for one clock cycle. The microprocessor is then stopped completely for one cycle. Data transfer between the memory and I/O takes place during this cycle. This method is called cycle stealing because the DMA controller takes away or steals a cycle without microprocessor recognition. Data transfer takes place over a period of time.

With interleaved DMA, the DMA controller chip takes over the system bus when the microprocessor is not using it. For example, the microprocessor does not use the bus while incrementing the program counter or performing an ALU operation. The DMA controller chip identifies these cycles and allows transfer of data between the memory and I/O device. Data transfer takes place over a period of time for this method.

The DMA controller chip usually has at least three registers normally selected by the controller's register select (RS) line: an address register, a terminal count register, and a status register. Both the address and terminal count registers are initialized by the microprocessor. The address register contains the starting of the data to be transferred, and the terminal count register contains the desired block to be transferred. The status register contains information such as completion of DMA transfer.

It should be mentioned that while using either block transfer or cycle stealing DMA in systems with dynamic RAMs, circuitry must be included for refreshing the dynamic RAM's during DMA transfer.

1.3.4.g Summary of Microcomputer I/O Methods

Figure 1.17 summarizes the I/O structure (explained so far) of typical microcomputers.
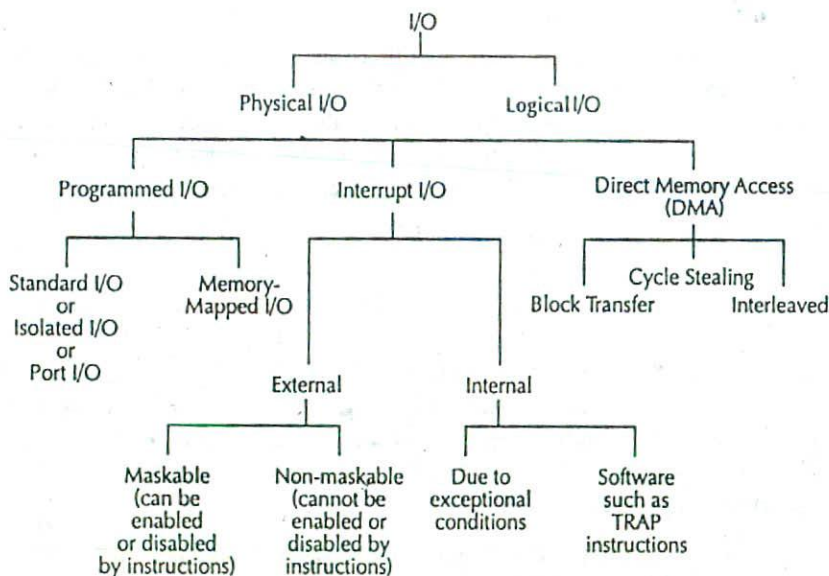
FIGURE 1.17   I/O structure of a typical microcomputer.

1.3.4.h Coprocessors

In typical 8-bit microprocessors such as the Intel 8085, technology places a limit on the chip area. As a consequence, these microprocessors include no hardware or firmware for performing scientific computations such as floating-point arithmetic, matrix manipulation, and graphic-data processing. Therefore, users of these systems must write these programs. Unfortunately,

this approach is unacceptable in high-speed applications, since program execution takes a significant amount of time. To eliminate this problem, coprocessors are used.

In this approach, a single chip is built for performing scientific computations at high speed. However, the chip is regarded as a companion to the original or host microprocessor. Typically, each special operation is encoded as an instruction that can be interpreted only by the companion processor. When the companion processor encounters one of these special instructions, it assumes the processing functions independent of the host microprocessor. The companion processor that operates in this manner is called the coprocessor. Therefore, this concept not only extends the capabilities of the host microprocessor, but also increases the processing rate of the system. The coprocessor concept is widely used with typical 32-bit microprocessors such as the Motorola 68020 and Intel 80386.

It is important to make the distinction between standard peripheral hardware and a coprocessor. A coprocessor is a device that has the capability of communicating with the main processor through the protocol defined as the coprocessor interface. As mentioned before, the coprocessor also adds additional instructions, registers, and data types that are not directly supported by the main microprocessor. The coprocessor provides capabilities to the user without appearing to be hardware external to the main microprocessor.

Standard peripheral hardware, on the other hand, is generally accessed through the use of interface registers mapped into the memory space of the main processor. The programmer uses standard processor instructions to access the peripheral interface registers and thus utilize the services provided by the peripheral. It should be pointed out that even though a peripheral can provide capabilities equivalent to a coprocessor for many applications, the programmer must implement the communication protocol between the main microprocessor and the peripheral necessary to use the peripheral hardware. Two main techniques may be used to pass commands to a coprocessor. These are intelligent monitor interface and coprocessors using special signals.

In the *intelligent monitor interface,* the coprocessor monitors the instruction stream by obtaining commands directly from the bus at the same time as the main microprocessor. The Intel 80387 floating-point coprocessor is of this type, as it monitors the instruction stream simultaneously with a main microprocessor such as the Intel 80386. This has the obvious advantage of requiring no additional bus cycles to pass the content of the instruction word to the coprocessor. One of the main disadvantages of this approach is that each coprocessor in the system must duplicate the bus monitoring circuitry and instruction queue, tracking all branches, wait states, operand fetches, and instruction fetches.

In the second type, the coprocessor may be explicitly addressed by certain instructions, which initiate a special sequence of microinstructions in the main microprocessor to effect command and operand transfer.

In this approach, when the main microprocessor executes a coprocessor instruction, it decodes the instruction and writes a command in the command register (one of the interface registers) specifying the operation required by the coprocessor. In response, the coprocessor writes data back in a register called the *response register* (one of the interface registers). The main microprocessor can read these data, and it tells the main microprocessor certain information such as whether additional information is required by the coprocessor to carry out the operation. If such data are required, the main microprocessor provides this; otherwise, the coprocessor carries out the operation concurrently with the main microprocessor and provides the result.

An advantage of this approach is that no special signals are required for the coprocessor interface.

One of the main disadvantages of this method is that once the main processor detects a coprocessor instruction, the main has to use bus bandwidth and timing to transmit the command to the appropriate coprocessor. The Motorola 68881 (floating-point coprocessor) is of this type.

Note that state-of-the-art 32-bit microprocessors such as the Intel 80486 and the Motorola 68040 implement coprocessor hardware such as floating point hardware and MMU on the microprocessor chip.

1.4 Microcomputer System Software and Programming Concepts

In this section, the basic concepts associated with system software and programming will be discussed.

1.4.1 System Software

Typical microcomputer system software provided in microcomputer development systems such as the HP64000 includes editors, assemblers, compilers, interpreters, debuggers, and an operating system.

The editor is used to create and change source programs. Source programs can be written in assembly language or a high-level language such as Pascal. The editor has commands to change, delete, or insert lines or characters. The text editor is a special type of editor that is used to enter and edit text in a general-purpose computer, whether the text is a report, a letter, or a program.

An assembler translates a source text that was created using the editor into a target machine language in object code (binary).

High-level languages contain English-like commands that are readily understandable by the programmer. High-level languages normally combine a number of assembly-level statements into a single high-level statement. A compiler is used to translate the high-level languages such as Pascal into machine language. The advantage of high-level languages over assembly language are ease of readability and maintainability.

Like a compiler, an interpreter usually processes a high-level language program. Unlike a compiler, an interpreter actually executes the high-level language program one statement at a time, rather than translating the whole program into a sequence of machine instructions to be run later.

The debugger provides an interactive method of executing and debugging the user's software, one or a few instructions at a time, allowing the user to see the effects of small pieces of the program and thereby isolate programming errors.

An operating system performs resource management and human-to-machine translation functions. A resource may be the microprocessor, memory, or an I/O device. Basically, an operating system is another program that tells the machine what to do under a variety of conditions. Major operating system functions include efficient sharing of memory, I/O peripherals, and the microprocessor among several users. An operating system is

1. The interface between hardware and users
2. The manager of system resources in accordance with system policy to achieve system objectives

Operating systems for microcomputers became available when microcomputers moved from process control applications to the general-purpose computer applications. It was appropriate to write a process control program in assembly language because the microcomputer was required to perform dedicated real-time control functions. But when the microcomputers evolved to the point of controlling several I/O devices (disks, printers), an organized operating system was needed.

Note that many laboratory trainers such as the Intel SDK-86 include primitive operating systems called monitors to provide functions such as keyboard/display interfaces and program debugging features.

1.4.2 Programming Concepts

In general, programs are developed using assembly and high-level languages.

1.4.2.a Assembly Language Programming

Program designers realized the importance of symbols in programming to improve readability and expedite the program development process. As a first step, they came up with the idea of giving a symbolic name for each instruction. These names are called mnemonics, and a program written using such mnemonics is called an assembly language program. Given below is a typical 8086 assembly language program:

```
MOV    AL,5 ; Load AL reg with 5
ADD    AL,3 ; (AL) ← (AL) + 3
HLT           ; Halt processing
```

From this example, it is clear that the usage of mnemonics (in our example MOV, ADD, HLT are the mnemonics) improves the readability of our program significantly.

An assembly language program cannot be executed by a machine directly, as it is not in binary form. Usually, we refer to a symbolic program as a source program. An assembler is needed in order to translate an assembly language (source) program into the machine language (object) executable by the machine. This is illustrated in Figure 1.18.
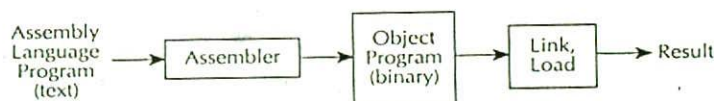
FIGURE 1.18   Assembly Process.

1.4.2.b High-Level Language Programming

Typical examples of high-level languages include FORTRAN, BASIC, Pascal, C and C++. The program shown below is written in FORTRAN, in order to obtain the sum of the first N natural numbers:

```
        READ (5,10) N
10      FORMAT (I4)
        NSUM = N*(N+1)/2
        WRITE (2,10) NSUM
        STOP
        END
```

A translator called a compiler is needed to translate a program written in a high-level language into binary form.
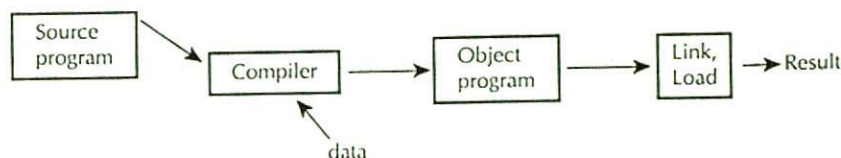
FIGURE 1.19   Compilation process.

1.4.2.c Which Programming Language to Choose?

Compilers normally provide inefficient machine codes because of the general guidelines which must be followed for designing them. It was found that a compiled high level language generates many more lines of machine code than an equivalent assembly language program. Therefore, the machine code program generated by the assembler will take up less memory space and also will execute much faster.

Therefore, although the C language includes I/O instructions, applications involving I/O are normally written in assembly language. One of the main uses of assembly language is in writing programs for real-time applications. Real time applications means that the task required by the application must be completed before any other input to the program can occur which will change its operation.

High level languages are normally used for applications which require extensive mathematical computations. These applications include optimizations and other non-real time control system applications.

1.5 Typical Microcomputer Addressing Modes and Instructions

In this chapter, some important characteristics and properties of microcomputer instruction sets are discussed. Topics include addressing modes and instruction types.

1.5.1 Introduction

An instruction manipulates the stored data, and a sequence of instructions constitutes a program. In general, an instruction has two components:

- Op-code field
- Operand field(s)

The op-code field specifies how data are to be manipulated. The op-code field may contain data or a microprocessor register or a memory address. Consider the following instruction:

<div align="center">

ADD   R1, R0

op-code field operand field

</div>

Assume that this microcomputer uses R1 as the source register and R0 as the destination register. The preceding instruction then adds the contents of registers R0 and R1 and saves the sum in register R0.

Depending on the number of addresses specified, one can have the following instruction formats:

- Three-operand
- Two-operand
- One-operand
- Zero-operand

The 8-bit microprocessors include mostly one-operand instructions along with some zero- and two-operand instructions. In 16-bit microprocessors, two-operand instructions are predominant although some zero- and one-operand instructions are present. The 32-bit microprocessors include all four instruction formats.

1.5.2 Addressing Modes

The sequence of operations that a microprocessor has to carry out while executing an instruction is called its *instruction cycle*. One of the activities in an instruction cycle is the determination of the addresses of the operands involved in that instruction.

The way in which a microprocessor accomplishes this task is by recognizing the *addressing mode* used in the instruction. Typical addressing modes supported by the instruction sets of popular processors will be examined.

An instruction is said to have an *inherent* addressing mode if it is a zero-operand instruction. As an example, consider the zero-operand instruction CLC which clears the carry flag to zero.

Whenever an instruction contains data in the operand field, it is called an *immediate mode* instruction. For example, consider the following instruction:

$$\text{ADD } \#25, \text{ R1; R1} \leftarrow \text{R1} + 25$$

In this instruction, the symbol # indicates that it is an immediate-mode instruction. This convention is adopted in the assemblers for processors such as MC68000.

An instruction is said to have an *absolute addressing mode* if it contains the address of the operand. For example, consider the following move instruction:

$$\text{MOVE } 5000, \text{ R2; R2} \leftarrow [5000]$$

This instruction copies the contents of memory location 5000 in the register R2.

An instruction is said to have a *register mode* if it contains a register in the operand field. For example, consider the following register mode intruction.

$$\text{ADD } \text{R2, R3; R3} \leftarrow \text{R2} + \text{R3}$$

This instruction uses register mode for both source and destination operands.

Whenever an instruction specifies a register that holds the address of an operand, the resulting addressing mode is known as the *register indirect* mode. From this definition, it follows that the Effective Address (EA) of an operand in the register-indirect mode is the contents of the register R. More formally, this result is written as follows:

$$\text{EA} = [\text{R}]$$

To illustrate this idea clearly, consider the following instruction:

$$\text{MOVE } (\text{R2}), (\text{R3}); [\text{R3}] \leftarrow [\text{R2}]$$

Assume that the following configuration exists:

$$[\text{R2}] = 5000_{16}$$
$$[\text{R3}] = 4000_{16}$$
$$[5000] = 1256_{16}$$
$$[4000] = 4629_{16}$$

This instruction copies the contents of the memory location, whose address is specified by the register R2, into the location whose address is specified by the register R3. Thus, after the execution of this instruction, the memory location 4000 will contain the value 1256_{16}.

1.5.3 Instruction Types

In general, instructions available in a processor may be broadly classified into five groups:

- Data transfer instructions
- Arithmetic instructions
- Logical instructions
- Program control instructions
- I/O instructions

Data transfer instructions are primarily concerned with data transfers between the microprocessor registers or between register and memory. An example is MOVE R0, R1 which transfers the contents of register R0 to register R1.

Typical arithmetic instructions include ADD and SUBTRACT instructions. For example, ADD R0, R1 adds the contents of R0 to R1 and stores the result in R1.

Logical instructions perform Boolean AND, OR, NOT, and EXCLUSIVE-OR operations on a bit-by-bit basis. An example is OR R0, R1 which logically ORs the contents of R0 with R1 and places the result in R1.

Typical program control instructions include unconditional and conditional branch and subroutine CALL instructions. For example, JMP 2035H unconditionally branches to the 16-bit address 2035H.

I/O instructions perform input and output operations. An example is IN PORTA which inputs the contents of an I/O port called Port A into a microprocessor register such as the accumulator.

1.6 Basic Features of Microcomputer Development Systems

A microcomputer development system is a tool that allows the designer to develop, debug, and integrate error-free application software in microprocessor systems.

Development systems fall into one of two categories: systems supplied by the device manufacturer (non-universal systems) and systems built by after-market manufacturers (universal systems). The main difference between the two categories is the range of microprocessors that a system will accommodate. Non-universal systems are supplied by the microprocessor manufacturer (Intel, Motorola, RCA) and are limited to use for the particular microprocessor manufactured by the supplier. In this manner, an Intel development system may not be used to develop a Motorola-based system. The universal development systems (Hewlett-Packard, Tektronix) can develop hardware and software for several microprocessors.

Within both categories of development systems, there are basically three types available: single-user systems, time-shared systems, and networked systems. A single-user system consists of one development station that can be used by one user at a time. Single-user systems are low in cost and may be sufficient for small systems development. Time-shared systems usually consist of a "dumb"-type terminal connected by data lines to a centralized microcomputer-based system that controls all operations. A networked system usually consists of a number of smart Cathode Ray Tubes (CRTs) capable of performing most of the development work and can be connected over data lines to a central microcomputer. The central microcomputer in a network system usually is in charge of allocating disk storage space and will download some programs into the user's work station microcomputer. A microcomputer development system is a combination of the hardware necessary for microprocessor design and the software to control the hardware. The basic components of the hardware are the central processor, the CRT terminal, mass storage device (floppy or hard disk), and usually an In-Circuit Emulator (ICE).

In a single-user system, the central processor executes the operating system software, handles the Input/Output (I/O) facilities, executes the development programs (editor, assembler, linker), and allocates storage space for the programs in execution. In a large multiuser networked system the central processor may be responsible for mass storage allocation, while a local processor may be responsible for the I/O facilities and execution of development programs.

The CRT terminal provides the interface between the user and the operating system or program under execution. The user enters commands or data via the CRT keyboard and the program under execution displays data to the user via the CRT screen.

Each program (whether system software or user program) is stored in an ordered format on disk. Each separate entry on the disk is called a *file*. The operating system software contains the

routines necessary to interface between the user and the mass storage unit. When the user requests a file by a specific *file name*, the operating system finds the program stored on disk by the file name and loads it into main memory. More advanced development systems contain *memory management* software that protects a user's files from unauthorized modification by another user. This is accomplished via a unique user identification code called USER ID. A user can only access files that have the user's unique code.

The equipment listed above comprises a basic development system, but most systems have other devices such as printers and PROM programmers attached. A printer is needed to provide the user with a hard copy record of the program under development.

After the application system software has been completely developed and debugged, it needs to be permanently stored for execution in the target hardware. The EPROM programmer takes the machine code and programs it into an EPROM. Erasable/Programmable Read Only Memories (EPROMs) are more generally used in system development as they may be erased and reprogrammed if the program changes. EPROM programmers usually interface to circuits particularly designed to program a specific PROM. These interface boards are called personality cards and are available for all the popular PROM configurations.

Most development systems support one or more in-circuit emulators (ICEs). The ICE is one of the most advanced tools for microprocessor hardware development. To use an ICE, the microprocessor chip is removed from the system under development (called the target processor) and the emulator plugged into the microprocessor socket. The ICE will functionally and electrically act identically to the target processor with the exception that the ICE is under the control of development system software. In this manner the development system may exercise the hardware that is being designed and monitor all status information available about the operation of the target processor. Using an ICE, processor register contents may be displayed on the CRT and operation of the hardware observed in a single-stepping mode. In-circuit emulators can find hardware and software bugs quickly that might take many hours using conventional hardware testing methods.

Architectures for development systems can be generally divided into two categories: the master/slave configuration and the single-processor configuration. In a master/slave configuration, the master (host) processor controls all development functions such as editing, assembling, and so on. The master processor controls the mass storage device and processes all I/O (CRT, printer).

The software for the development systems is written for the master processor which is usually not the same as the slave (target) processor. The slave microprocessor is typically connected to the user prototype via a 40-pin connector (the number varies with the processor) which links the slave processor to the master processor.

Some development systems such as the HP 64000 completely separate the system bus from the emulation bus and therefore use a separate block of memory for emulation. This separation allows passive monitoring of the software executing on the target processor without stopping the emulation process. A benefit of the separate emulation facilities allows the master processor to be used for editing, assembling, and so on, while the slave processor continues the emulation. A designer may therefore start an emulation running, exit the emulator program, and at some future time return to the emulation program.

Another advantage of the separate bus architecture is that an operating system needs to be written only once for the master processor and will be used no matter what type of slave processor is being emulated. When a new slave processor is to be emulated, only the emulator probe needs to be changed.

A disadvantage of the master/slave architecture is that it is expensive. In single-processor architecture, only one processor is used for system operation and target emulation. The single processor does both jobs of executing system software as well as acting as the target processor. Since there is only one processor involved, the system software must be rewritten for each type

of processor that is to be emulated. Since the system software must reside in the same memory used by the emulator, not all memory will be available to the emulation process, which may be a disadvantage when large prototypes are being developed. The single processor systems are inexpensive.

The programs provided for microprocessor development are the operating system, editor, assembler, linker, compiler, and debugger.

The operating system is responsible for executing the user's commands. The operating system (such as UNIX) handles I/O functions, memory management, and loading of programs from mass storage into RAM for execution.

The editor allows the user to enter the source code (either assembly language or some high-level language) into the development system.

Almost all current microprocessor development systems use the character-oriented editor, more commonly referred to as the screen editor. The editor is called a screen editor because the text is dynamically displayed on the screen and the display automatically updates any edits made by the user.

The screen editor uses the pointer concept to point to the character(s) that need editing. The pointer in a screen editor is called the cursor and special commands allow the user to position the cursor to any location displayed on the screen. When the cursor is positioned, the user may insert characters, delete characters, or simply type over the existing characters.

Complete lines may be added or deleted using special editor commands. By placing the editor in the *insert* mode, any text typed will be inserted at the cursor position when the cursor is positioned between two existing lines. If the cursor is positioned on a line to be deleted, a single command will remove the entire line from the file.

Screen editors implement the editor commands in different fashions. Some editors use dedicated keys to provide some cursor movements. The cursor keys are usually marked with arrows to show the direction of cursor movement.

More advanced editors (such as the HP 64000) use *soft keys*. A soft key is an unmarked key located on the keyboard directly below the bottom of the CRT screen. The mode of the editor decides what functions the keys are to perform. The function of each key is displayed on the screen directly above the appropriate key. The soft key approach is valuable because it frees the user from the problem of memorizing many different special control keys. The soft key approach also allows the editor to reassign a key to a new function when necessary.

The source code generated on the editor is stored as ASCII or text characters and cannot be executed by a microprocessor. Before the code can be executed, it must be converted to a form acceptable by the microprocessor. An assembler is the program used to translate the assembly language source code generated with an editor into object code (machine code) which may be executed by a microprocessor.

Assemblers recognize four *fields* on each line of source code. The fields consist of a variable number of characters and are identified by their position in the line. The fields, from left to right on a line, are the label field, the mnemonic or op-code field, the operand field, and the comment field. Fields are separated by characters called *delimiters* which serve as a flag to the assembler that one field is done and the next one is to start. Typical delimiters and their uses are

space	used to separate fields
TAB	used to separate fields
,	used between addresses or data in the operand field
;	used before a comment statement
:	used after a label

A few typical lines of 8085 source code are

LABEL FIELD	MNEMONIC FIELD	OPERAND FIELD	COMMENT FIELD
	MOV	AL,5	; LOAD 5 INTO AL
	ADD	AL,2	; ADD 5 AND 2, STORE RESULT IN AL

As can be seen in the above example, tab keys are used instead of spaces to separate the fields to give a more *spread out* line which is easier to read during debugging.

In order for the assembler to differentiate between numbers and labels, specific rules are set up which apply to all assemblers. A label must start with a letter. After the letter, a combination of letters and numbers (called alphanumerics) may be used. For example, when grouping lines of code by function, a common alphabetic string may be used followed by a unique number for the label: LOOP01, LOOP02, LOOP10, and so on.

A numeric quantity must start with a number, even though the number may be in hex (which may start with a letter). Most assemblers assume that a number is expressed in the decimal system and if another base is desired, a special code letter is used immediately following the number. The usual letter codes used are

B   binary
C   octal
H   hex (Motorola uses $ before the number)

To avoid confusion when hex quantities are used, a leading zero is inserted to tell the assembler that the quantity is a number and not a label (for example, the quantity FA in hex would be represented by 0FAH in the source code).

Assembler Directives

Assembler directives are instructions entered into the source code along with the assembly language. These directives do not get translated into object code but are used as special instructions to the assembler to perform some special functions. The assembler will recognize directives that assign memory space, assign addresses to labels, format the pages of the source code, and so on.

The directive is usually placed in the op-code field. If any labels or data are required by the directive, they are placed in the label or operand field as necessary.

. Some common directives will now be discussed in detail.

a. *ORIGIN (ORG).* The ORG statement is used by the programmer when it is necessary to place the program in a particular location in memory. As the assembler is translating the source code, it keeps an internal counter (similar to the microprocessor program counter) that keeps track of the address for the machine code. The counter is incremented automatically and sequentially by the assembler. If the programmer wishes to alter the locations where the machine code is going to be located, the ORG statement is used.

For example, if it is desired to have a subroutine at a particular location in memory, such as 2000H, the statement ORG 2000H, would be placed immediately before the subroutine to direct the assembler to alter the internal program counter.

Most assemblers will assume a starting address of zero if no ORG statement is given in the source code.

b. *EQUATE (EQU).* The EQU instruction is used to assign the data value or address in the operand field to the label in the label field. The EQU instruction is valuable because it allows the programmer to write the source code in symbolic form and not be concerned with the

numeric value needed. In some cases, the programmer is developing a program without knowing what addresses or data may be required by the hardware. The program may be written and debugged in symbolic form and the actual data added at a later time. Using the EQU instruction is also helpful when a data value is used several times in a program. If, for example, a counter value was loaded at ten different locations in the program, a symbolic label (such as COUNT) could be used and the label count defined at the end of the program. By using this technique, if it is found during debugging that the value in COUNT must be changed, it need only be changed at the EQU instruction and not at each of the ten locations where it is used in the program.

c. *DEFINE BYTE (DEFB or DB).*   The DB instruction is used to set a memory location to a specific data value. The DB instruction is usually used to create data tables or to preset a flag value used in a program. As the name implies, the DB instruction is used for creating an 8-bit value.

For example, if a table of four values, 44H, 34H, 25H, and 0D3H, had to be created at address 2000H, the following code could be written:

```
        ORG 2000H              ; SET TABLE ADDRESS
    TABLE DB  44H,34H,25H,0D3H; PRESET  TABLE  VALUES
```

The commas are necessary for the assembler to be able to differentiate between data values. When the code is assembled, the machine code would appear as follows:

```
              . . .
           2000    44
           2001    34
           2002    25
           2003    D3
              . . .
```

d. *DEFINE WORD (DEFW or DW).*   Similarly to DB, DW defines memory locations to specific values. As the name implies, the memory allotted is in word lengths which are usually 16 bits wide. When assigning a 16-bit value to memory locations, two 8-bit memory locations must be used. By convention, most assemblers store the least significant byte of the 16-bit value in the first memory location and the most significant byte of the 16-bit value in the next memory location. This technique is sometimes referred to as *Intel style*, because the first microprocessors were developed by Intel, and this storage method is how the Intel processors store 16-bit words.

Data tables may be created with the DW instruction, but care must be taken to remember the order in which the 16-bit words are stored. For example, consider the following table:

```
           ORG 2500H
       DATA DW 4000H,  2300H,  4BCAH
```

The machine code generated for this table would appear as follows:

```
              . . .
           2500    00
           2501    40
           2502    00
           2503    23
           2504    CA
           2505    4B
              . . .
```

e. *TITLE.*   TITLE is a formatting instruction that allows the user to name the program and have the name appear on the source code listing. Consider the following line:

TITLE 'MULTIPLICATION ROUTINE'

When the assembler generates the program listing, each time it starts a new page the title MULTIPLICATION ROUTINE appears at the top of each page.

Several types of assemblers are available, the most common types are discussed below.

a. *One-Pass Assembler.*   The one-pass assembler was the first type to be developed and is therefore the most primitive. Very few systems use a one-pass assembler because of the inherent problem that only *backward references* may be used.

In a one-pass assembler the source code is processed only once. As the source code is processed, any labels encountered are given an address and stored in a table. Therefore, when the label is encountered again, the assembler may look backward to find the address of the label. If the label has not been defined yet (for example, a jump instruction that jumps forward), the assembler issues an error message.

b. *Two-Pass Assembler.*   In the two-pass assembles, the source code is passed twice through the assembler. The first pass made through the source code is specifically for the purpose of assigning an address to all labels. When all labels have been stored in a table with the appropriate addresses, a second pass is made to actually translate the source code into machine code.

The two-pass style assembler is the most popular type of assembler currently in use.

c. *Macroassembler.*   A macroassembler is a type of two-pass assembler that allows the programmer to write the source code in *macros*. A macro is a sequence of instructions that the programmer gives a name. Whenever the programmer wishes to duplicate the sequence of instructions, the macro name is inserted into the source code.

d. *Cross Assemblers.*   A cross assembler may be of any of the types already mentioned. The distinguishing feature of a cross assembler is that it is not written in the same language used by the microprocessor that will execute the machine code generated by the assembler.

Cross assemblers are usually written in a high-level language such as FORTRAN which will make them machine independent. For example, an 8086 assembler may be written in FORTRAN and then the assembler may be executed on another machine such as the Motorola 6800.

e. *Metaassembler.*   The most powerful assembler is the metaassembler because it will support many different microprocessors. The programmer merely specifies at the start of the source code which microprocessor assembly language will be used and the metaassembler will translate the source code to the correct machine code.

The output file from most development system assemblers is an object file. The object file is usually relocatable code that may be configured to execute at any address. The function of the linker is to convert the object file to an *absolute* file which consists of the actual machine code at the correct address for execution. The absolute files thus created are used for debugging and finally for programming EPROMs.

Debugging a microprocessor-based system may be divided into two categories: software debugging and hardware debugging. Both debug processes are usually carried out separately from each other because software debugging can be carried out on an Out-Of-Circuit-Emulator (OCE) without having the final system hardware.

The usual software development tools provided with the development system are

- Single-step facility
- Breakpoint facility

A single-stepper simply allows the user to execute the program being debugged one instruction at a time. By examining the register and memory contents during each step, the debugger can detect such program faults as incorrect jumps, incorrect addressing, erroneous op codes, and so on.

A breakpoint allows the user to execute an entire section of a program being debugged.

There are two types of breakpoint systems: hardware and software. The hardware breakpoint uses hardware to monitor the system address bus and detect when the program is executing the desired breakpoint location. When the breakpoint is detected, the hardware uses the processor control lines to either halt the processor for inspection or cause the processor to execute an interrupt to a breakpoint routine. Hardware breakpoints can be used to debug both ROM- and RAM-based programs. Software breakpoint routines may only operate on a system with the program in RAM because the breakpoint instruction must be inserted into the program that is to be executed.

Single-stepper and breakpoint methods complement each other. The user may insert a breakpoint at the desired point and let the program execute up to that point. When the program stops at the breakpoint the user may use a single-stepper to examine the program one instruction at a time. Thus, the user can pinpoint the error in a program.

There are two main hardware debugging tools: the logic analyzer and the in-circuit emulator.

Logic analyzers are usually used to debug hardware faults in a system. The logic analyzer is the digital version of an oscilloscope because it allows the user to view logic levels in the hardware.

In-circuit emulators can be used to debug and integrate software and hardware.

PC-based workstations are extensively used as development systems.

1.7 System Development Flowchart

The total development of a microprocessor-based system typically involves three phases: software design, hardware design, and program diagnostic design. A systems programmer will be assigned the task of writing the application software, a logic designer will be assigned the task of designing the hardware, and typically both designers will be assigned the task of developing diagnostics to test the system. For small systems, one engineer may do all three phases, while on large systems several engineers may be assigned to each phase. Figure 1.20 shows a flowchart for the total development of a system. Notice that software and hardware development may occur in parallel to save time.

1.7.1 Software Development

The first step in developing the software is to take the system specifications and write a flowchart to accomplish the desired tasks that will implement the specifications.

The assembly language or high-level source code may now be written from the system flowchart.

The complete source code is then assembled. The assembler will check for syntax errors and print error messages to help in the correction of errors.

The normal output of an assembler is the object code and a program listing. The object code will be used later by the linker. The program listing may be sent to a disk file for use in debugging or it may be directed to the printer.
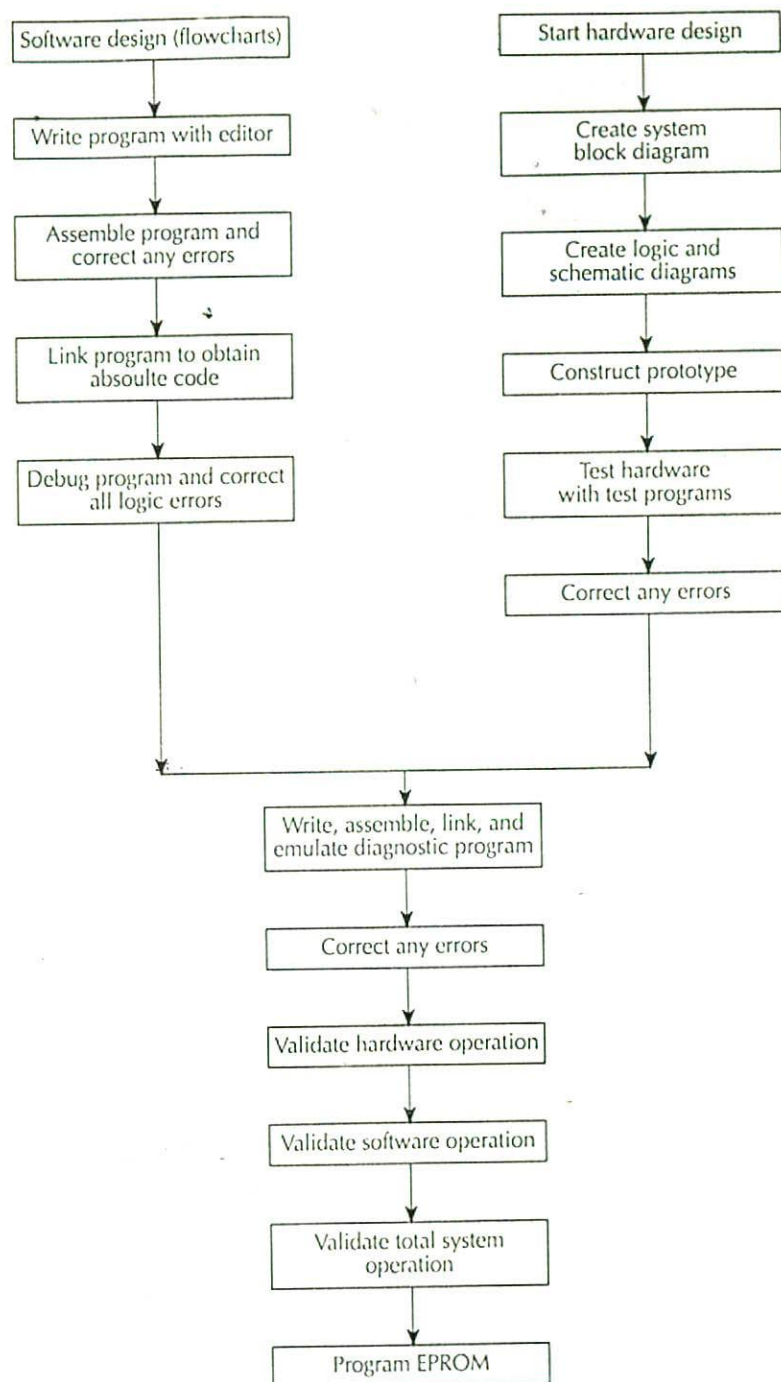
FIGURE 1.20   Microprocessor system development flowchart.

The linker can now take the object code generated by the assembler and create the final absolute code that will be executed on the target system. The emulation phase will take the absolute code and load it into the development system RAM. From here, the program may be debugged using breakpoints or single-stepping.

1.7.2 Hardware Development

Working from the system specifications, a block diagram of the hardware must be developed. The logic diagram and schematics may now be drawn using the block diagram as a guide. A prototype may now be constructed and tested for wiring errors.

When the prototype has been constructed it may be debugged for correct operation using standard electronic testing equipment such as oscilloscopes, meters, logic probes, and logic analyzers, all with test programs created for this purpose.

After the prototype has been debugged electrically, the development system in-circuit emulator may be used to check it functionally. The ICE will verify memory map, correct I/O operation, and so on.

The next step in the system development is to validate the complete system by running operational checks on the prototype with the finalized application software installed. The EPROM is then programmed with the error-free programs.

1.8 Typical Microprocessors

Intel and Motorola established the standard for future microprocessors.

Both Intel and Motorola introduced popular 8-bit microprocessors, namely, Intel 8080/8085, and Motorola 6800/6809 in the mid 1970's. These microprocessors are capable of directly addressing 64k of memory and were popular for several years. They are still being utilized in many inexpensive low-speed applications.

In 1978, Intel introduced its first 16-bit microprocessor, the Intel 8086. Several variations of the Intel 16-bit microprocessor family include the 8088, 80186/80188 and 80286. The 8088 is similar to the 8086 except that the 8088 has an 8-bit data bus. Both the 8086 and 8088 can run at 5 Mhz. They are packaged in 40 pins and can directly address one megabyte of memory. In 1982, Intel introduced an enhanced version of the 8086/8088 which is the Intel 80186/80188. The 80186 has a 16-bit data bus while the 80188 has an 8-bit data bus. The 80186/80188 can operate at 8 Mhz (80186), 6 Mhz (80186-6) and other frequencies. The 80186 integrates the 8086 and several new functional units into a single chip. The new on-chip components include a clock generator, two DMA channels, interrupt controller, address decoding, and three 16-bit programmable timers. Like the 8086, the 80186 can directly address one megabyte of memory. The 80186 is housed in a 68-pin leadless package.

The 80286 is Intel's other 16-bit microprocessor with on-chip memory protection capability primarily designed for multiuser/multitasking systems and introduced virtual memory.

Intel's 32-bit microprocessor family includes the 80386, 80486, and Pentium. The 80386 includes an on-chip memory management unit while the 80486 (DX and DX2) contains all the features of the 80386 along with floating-point hardware in the same chip. The Pentium, on the other hand, is a superscalar processor. This means that it includes dual pipelining and executes more than one instruction per cycle.

Soon after introduction of the Intel 8086, Motorola released its first 16-bit microprocessor, the 68000. It is housed in 64-pin or 68-pin package with a direct addressing capability of 16 megabytes. The 68008 is similar to the 68000 except that the 68000 has a 16-bit data bus while the 68008 has an 8-bit data bus. Also, the 68008 can directly address one megabyte of memory. The 68010 and 68012 are 16-bit microprocessors similar to the Intel 80286 and provide on-chip memory management support. Motorola CPU 32 is an enhanced 68000 along with many features of the 16-bit 68010 and the 32-bit 68020. The CPU 32 can support high level languages and is suited for controller applications. For example, in applications where power consumption is a consideration, the CPU32 forces the device into a low-power standby mode when immediate processing is not required upon execution of the LPSTOP instruction. Also, to maximize throughput for real-time applications, reference data is often precalculated and stored in memory for quick access upon execution of the TBL instruction.

Motorola's 32-bit microprocessor family includes the 68020, 68030, and 68040.

Since 1988, Intel, Morotola, and others have been introducing the RISC microprocessors. Some of these are the Intel 80960 family, the Motorola MC88100, the Apple/IBM/Motorola PowerPC, and Digital Equipment Corporation's Alpha 21164. The 80960 and 88100 are 32-bit microprocessors. They can directly address 4 gigabytes of memory. The PowerPC family, on the other hand, includes both 32-bit and 64-bit microprocessors. The 80960 and 88100 include a 32-bit data bus while the PowerPC contains a 64-bit data bus. The PowerPc is based on IBM PowerPC architecture and Motorola's 88100 bus interface design. These RISC microprocessors find extensive applications in embedded controls such as laser printers.

The PowerPC 601 outperforms and under-prices the Pentium but to be successful must build up its selection of software applications. Note that the Pentium had a flaw in its division algorithm caused by a problem with a Lookup table used in the division. Intel recently corrected this problem.

Tables 1.2a and 1.2b provide a brief description of Intel and Motorola's typical 16- and 32-bit microprocessors.

1.9 Typical Practical Applications

Microprocessors are being extensively used in a wide variety of applications. Typical applications include dedicated controllers, personal workstations, and real-time robotics control. Some of these applications are described in the following.

1.9.1 Personal Workstations

Personal workstations are designed using the high-performance 16- and 32-bit microprocessors. A dedicated single user (rather than multiple users sharing resources of a single microcomputer) can obtain significant computing power from these workstations.

The state-of-the-art workstations use 32-bit microprocessors to provide certain sophisticated functions such as IC layout, 3D graphics, and stress analysis.

1.9.2 Fault-Tolerant Systems

In many applications such as control of life-critical systems, control of nuclear waste, and unattended remote system operation, the reliability of the hardware is of utmost importance. The need for such reliable systems resulted in fault-tolerant systems. These systems use redundant computing units to provide reliable operation. However, the cost of fault-tolerant

Table 1.2a   Intel 16-bit microprocessors

	8086	8088	80286
Introduced	June 1978	June 1978	Feb. 1982
Maximum clock speed	10 Mhz	10 Mhz	20 Mhz
On-chip cache	No	No	No
MIPS (Millions of Instructions Per Second)	0.33	0.33	1.2
Addressing Modes*	8	8	8
Transistors	29,000	29,000	134,000
Data Bus	16-bit	8-bit	16-bit
Address Bus	20-bit	20-bit	24-bit
Directly Addressable memory	One megabyte	One megabyte	Sixteen megabytes
Number of pins	40	40	68
Virtual Memory	No	No	Yes; One gigabyte per task
On-Chip memory management and protection	No	No	Yes
Math Coprocessor Interface	8087	8087	80287XL/XLT

*Addressing modes include register, immediate, and memory modes.

Table 1.2a   Intel 32-bit microprocessors

	80386DX	80386SX	486DX	486SX	486DX2	Pentium
Introduced	October 1985	June 1988	April 1989	April 1991	March 1992	May 1993
Maximum clock speed	40 Mhz	33 Mhz	50 Mhz	25 Mhz	66 Mhz	100 Mhz
MIPS	6	2.5	20	16.5	40	112
Transistor	275,000	275,000	1.2 Millions	1.185 Millions	1.2 Millions	3.1 Millions
On-chip cache memory	Support chips avaliable	Support chips avaliable	Yes	Yes	Yes	Yes
Data Bus	32-bit	16-bit	32-bit	32-bit	32-bit	64-bit
Address Bus	32-bit	24-bit	32-bit	32-bit	32-bit	32-bit
Directly Addressable Memory	4 gigabytes	16 megabytes	4 gigabytes	4 gigabytes	4 gigabytes	4 gigabytes
Pins	132	100	168	168	168	273
Virtual Memory	Yes	Yes	Yes	Yes	Yes	Yes
On-chip Memory Management and Protection	Yes	Yes	Yes	Yes	Yes	Yes
Addressing Modes*	11	11	11	11	11	11
Floating-point	387DX	387SX	**	487SX	**	**

   * Addressing modes include register, immediate, and memory modes.

   ** Indicates on-chip floating point.

   Note that the 80386SL is also a 32-bit microprocessor with a 16-but data bus like the 80386SX (not listed in the table above). The 80386SL can run at a speed of up to 25 MHz and has a direct addressing capability of 32 megabytes. The 80386SL provides virtual memory support along with on-chip memory management and protection. It can be interfaced to the 80387SX to provide floating-point support. The 80386SL includes an on-chip disk controller hardware.

Table 1.2b   Motorola 16-bit Microprocessors

	MC68000	MC68010	CPU32
Maximum Clock Speed	25 Mhz	25 Mhz	16.67 Mhz
Pins	64, 68	64,68	64, 68
Data Bus	16-bit	16-bit	16-bit
Virtual Memory	No	Yes	Yes
Directly Addressable Memory	16 Megabytes	16 Megabytes	16 Megabytes
Control Registers	None	3	3
Stack Pointers	2	2	2
Cache	No	No	No
Addressing Modes	14	14	16
Coprocessor Interface	No on-chip hardware	No on-chip hardware	No on-chip hardware

Table 1.2b   Motorola 32-bit Microprocessors

	MC68020	MC68030	MC68040
Maximum Clock Speed	33 Mhz (8 Mhz min.)	33 Mhz (8 Mhz min.)	33 Mhz (8 Mhz min.)
Pins	114	118	179
Address Bus	32-bit	32-bit	32-bit
Addressing Modes	18	18	18
Maximum Addressable Memory	4 Gigabytes	4 Gigabytes	4 Gigabytes
Memory Management	By interfacing the 68851 MMU chip	On-chip MMU	On-chip MMU
Cache (on-chip)	Instruction cache	Instruction and data cache	Instruction and data cache
Floating Point	By interfacing 68881/68882 floating-point coprocessor chip	By interfacing 68881/68882 floating-point coprocessor chip	On-chip floating-point hardware

systems can be very high if the performance requirements of the application need high-performance VAX-type computers. Since the performance levels of 32-bit microprocessors are comparable to the VAX-type computer, multiple 32-bit microprocessors in a redundant configuration outperform the VAXs. Thus, the 32-bit microprocessors provide efficient fault-tolerant systems.

1.9.3 Real-Time Controllers

Real-time controllers such as flight-control systems for aircraft, flight simulators, and automobile engine control require high-performance computers. Some of these applications were handled in the past by using mainframe computers which resulted in high cost and the controllers occupied large spaces.

The state-of-the-art flight simulators use multiple 32-bit microprocessors to perform graphic manipulation, data gathering, and high-speed communications. Obviously, an application such as real-time automobile engine control using mainframe computers is not practical since these systems are not small enough to fit under a car hood. These controllers are currently being designed using the small-sized 32-bit microprocessors to perform high-speed data manipulation and calculation.

1.9.4 Robotics

The processing requirements of complex robots attempting to emulate human activities exceed the capabilities of 8- and 16-bit microprocessors. With 32-bit microprocessors, it is now feasible to design these controllers at low cost. In many cases, the microprocessor is used as the brain of the robot. In a typical application, the microprocessor will input the actual arm angle measurement from a sensor, compare it with the desired arm angle, and will then send outputs to a motor to position the arm.

Mitsubishi manufactured the first 68020-based system robot control system.

1.9.5 Embedded Control

Embedded Control microprocessors, also called embedded controllers, are designed to be programmed to manage specific tasks. Once programmed, the embedded controllers can manage the functions of a wide variety of electronic products. Since the microprocessors are embedded in the host system, their presence and operation are basically hidden from the host system. Typical embedded control applications include office automation products such as copiers, laser printers, fax machines, and consumer electronics like VCRs, microwave ovens, and automative systems.

There are two types of embedded control applications, event control (real-time) and data control. Even control applications require distributed embedded controllers dedicated to single functions.

These applications include motor engine or instrument control system. In these applications, the controllers offer real-time response to support programs with small amounts of data. Typically, 8- and 16-bit single chip microcontrollers such as the Intel 8751/8096 and Motorola HC11/HC16 are used.

The 80186/80188 based microcontrollers are used for data control applications.

RISC microprocessor-based embedded controllers perform many different functions while handling large programs using large amounts of data. Applications such as laser printers require a high performance microprocessor with on-chip floating-point hardware. The RISC microprocessors are ideal for these types of applications.

QUESTIONS AND PROBLEMS

1.1 What is the basic difference between the microprocessor and the microcomputer?

1.2 What is meant by an 8-bit microprocessor and a 16-bit microprocessor?

1.3 Interpret FE_{16} as unsigned, signed, and floating-point numbers.

1.4 Determine the carry, sign, overflow, and zero flags for the following operation:

$$ADD\ 82A1H,\ 231FH$$

1.5 What are the basic difference between EPROM, EAROM, and flash ROM?

1.6 What is the difference between static and dynamic RAM?

1.7 Assume the following ROM chip and the microprocessor:

ROM
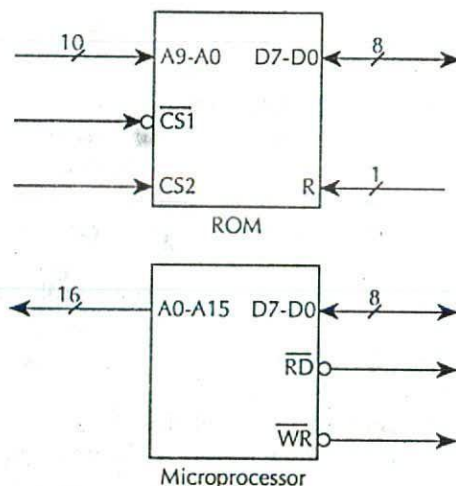
Microprocessor

Connect the microprocessor to the ROM to obtain the following memory map: 0000_{16} through $03FF_{16}$. Use only the signals shown. Draw a neat logic diagram and analyze the memory map.

1.8 Assume the following microprocessor and the RAM chip:

Microprocessor   RAM
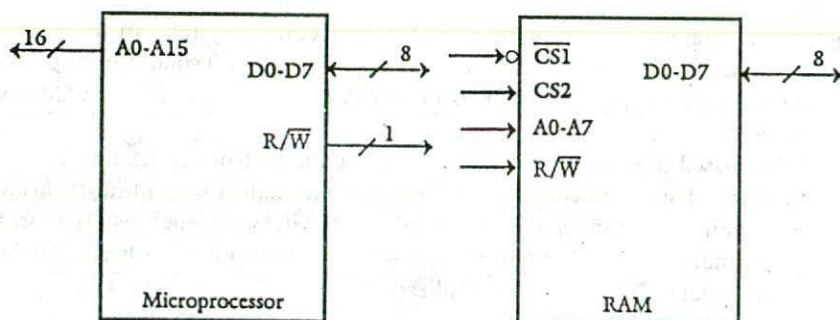
Draw a neat logic diagram showing connections between the above microprocessor and the RAM chip using only the signals shown to include the memory map 0200_{16} through $02FF_{16}$. Use linear decoding.

1.9   Use the following chips to design a microcomputer memory to include the map 3000_{16} through $33FF_{16}$.
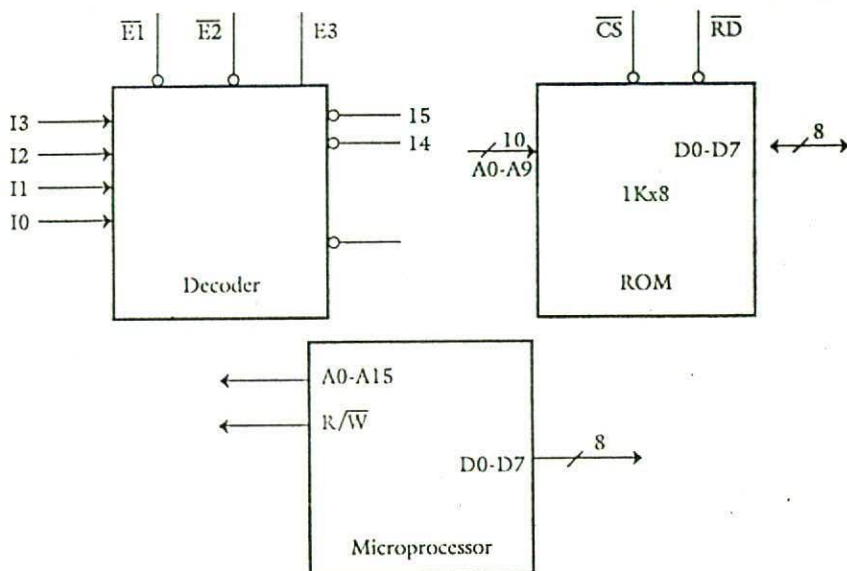
Draw a neat logic diagram.

1.10   What is meant by foldback in linear decoding?

1.11   Define the three basic types of I/O.

1.12   What is the difference between:

   i)   Standard and memory-mapped programmed I/O
   ii) Maskable and nonmaskable interrupts
   iii) Internal and external interrupts
   iv) Block transfer and interleaved DMA
   v)   Polled and daisy chain
   vi) One-pass and two-pass assemblers

1.13   Comment on the importance of the following architectural features in an operating system implementation:

   a)   Address translation
   b)   Protection
   c)   Program relocation

1.14   Explain clearly the differences between segmentation and paging. Can you think of a situation where it would be advantageous to define a virtual memory that is smaller than available physical memory?

1.15   What is the purpose of cache memory? Discuss briefly the various types of cache.

1.16 Discuss the main features of typical coprocessors.

1.17 What is the difference between software breakpoint and hardware breakpoint?

1.18 Discuss the basic features of microcomputer development systems.

1.19 Compare the typical features of the 16- and 32-bit microprocessors by Intel and Motorola.

1.20 What types of applications are the RISC microprocessors used for?

1.21 Discuss floating-point data formats supported by both Intel 80387 and Motorola 68881/6882.

2

INTEL 8085

This chapter describes hardware, software, and interfacing aspects of the Intel 8085. Topics include 8085 register architecture, addressing modes, instruction set, input/output, and system design.

2.1 Introduction

The Intel 8085 is an 8-bit microprocessor. The 8085 is designed using NMOS in 40-in DIP (Dual In-line Package). The 8085 can be operated from either 3.03 MHz maximum (8085A) or 5 MHz maximum (8085A-2) internal clock frequency.

The 8085 has three enhanced versions, namely, the 8085AH, 8085AH-2, and 8085AH-1. These enhanced processors are designed using the HMOS (High-density MOS) technology. Each is packaged in a 40-pin DIP like the 8085. These enhanced microprocessors consume 20% lower power than the 8085A. The internal clock frequencies of the 8085AH, 8085AH-2, and 8085AH-1 are 3, 5, 6 MHz, respectively. These HMOS 8-bit microprocessors are expensive compared to the NMOS 8-bit 8085A.

Figure 2.1 shows a simplified block diagram of the 8085 microprocessor. The accumulator connects to the data bus and the Arithmetic and Logic Unit (ALU). The ALU performs all data manipulation, such as incrementing a number or adding two numbers.

The temporary register feeds the ALU's other input. This register is invisible to the programmer and is controlled automatically by the microprocessor's control circuitry.

The flags are a collection of flip-flops that indicate certain characteristics of the result of the most recent operation performed by the ALU. For example, the zero flag is set if the result of an operation is zero. The zero flag is tested by the JZ instruction.

The instruction register, instruction decoder, program counter, and control and timing logic are used for fetching instructions from memory and directing their execution.

2.2 Register Architecture

The 8085 registers and status flags are shown in Figure 2.2.

The accumulator (A) is an 8-bit register. Most arithmetic and logic operations are performed using the accumulator. All I/O data transfers between the 8085 and the I/O devices are performed via the accumulator. Also, there are a number of instructions that move data between the accumulator and memory.

The B, C, D, E, H, and L are each 8 bits long. Registers H and L are the memory address register or data counter. This means that these two registers are used to store the 16-bit address
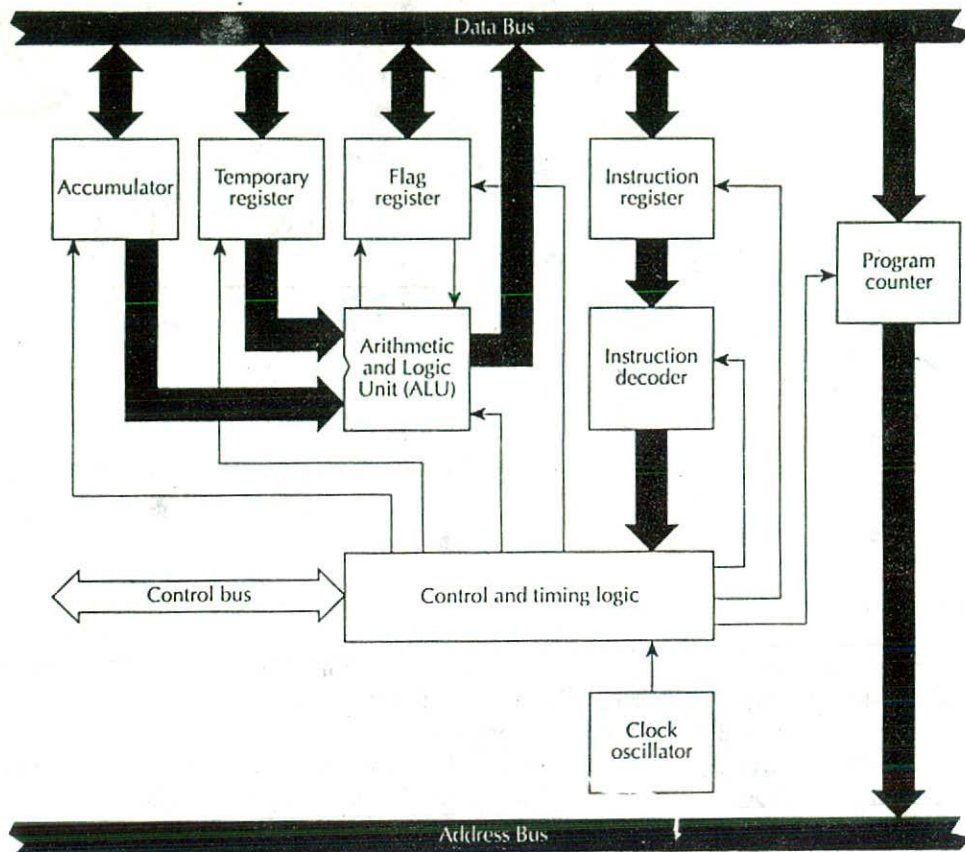
FIGURE 2.1   Simplified 8085 block diagram.

of 8-bit data being accessed from memory. This is the implied or register indirect addressing mode. There are a number of instructions, such as MOV reg, M, and MOV M, reg, which move data between any register and memory location addressed by H and L. However, using any other memory reference instruction, data transfer takes place between a memory location and the only 8085 register, the accumulator. The instruction LDAX B is a typical example.
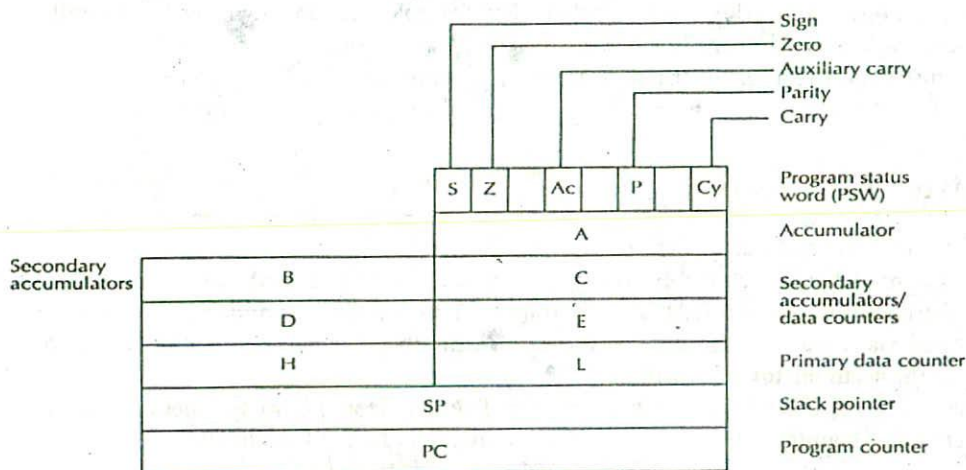
FIGURE 2.2   8085 microprocessor registers and status flags.

Registers B, C, D, and E are secondary accumulators or data counters. There are a number of instructions to move data between any two registers. There are also a few instructions that combine registers B and C or D and E, as a 16-bit data counter with high byte of a pair contained in the first register and low byte in the second. These instructions typically include LDAX B, LDAX D, STAX B, and STAX D, which transfer data between memory and the accumulator.

Each of these 8-bit registers can be incremented and decremented by a single byte instruction. There are a number of instructions which combine two of these 8-bit registers to form 16-bit register pairs as follows:

<div align="center">

A	and	PSW
B	and	C
D	and	E
H	and	L
high-order byte		low-order byte

</div>

The 16-bit register pair obtained by combining the accumulator and the program status word (PSW) is used only for stack operations. Sixteen bit arithmetic operations use B and C, D and E, or H and L as 16-bit data registers.

The program status word consists of five status flags. These are described below.

The carry flag (Cy) reflects the final carry out of the most significant bit of any arithmetic operation. Any logic instruction resets or clears the carry flag. This flag is also used by the shift and rotate instructions. The 8085 does not have any CLEAR CARRY instruction. One way of clearing the carry will be by ORing or ANDing the accumulator with itself.

The parity status flag (P) is set to 1 if an arithmetic or logic instruction generates an answer with even parity, that is, containing an even number of 1 bits. This flag is 0 if the arithmetic or logic instruction generates an answer with odd parity, that is, containing an odd number of 1s.

The auxiliary carry flag (Ac) reflects any carry from bit 3 to bit 4 (assuming 8-bit data with bit 0 as the least significant bit and bit 7 as the most significant bit) due to an arithmetic operation. This flag is useful for BCD operations.

The zero flag (Z) is set to 1 whenever an arithmetic or logic operation produces a result of 0. The zero flag is cleared to zero for a nonzero result due to arithmetic or logic operation.

The sign status flag (S) is set to the value of the most significant bit of the result in the accumulator after an arithmetic or logic operation. This provides a range of -128_{10} to $+127_{10}$ (with 0 being considered positive) as the 8085's data-handling capacity.

The 8085 does not have an overflow flag. Note that execution of arithmetic or logic instructions in the 8085 affects the flags. All conditional instructions in the 8085 instruction set use one of the status flags as the required condition.

The stack pointer (SP) is 16 bits long. All stack operations with the 8085 use 16-bit register pairs. The stack pointer contains the address of the last data byte written into the stack. It is decremented by 2 each time 2 bytes of data are written or pushed onto the stack and is incremented by 2 each time 2 bytes of data are read from or pulled (popped) off the stack, that is, the top of the stack has the lowest address in the stack that grows downward.

The program counter (PC) is 16 bits long to address up to 64K of memory. It usually addresses the next instruction to be executed.

2.3 Memory Addressing

When addressing a memory location, the 8085 uses either register indirect or direct memory addressing. With register indirect addressing, the H and L registers perform the function of the memory address register or data counter; that is, the H, L pair holds the address of the data.

With this mode, data transfer may occur between the addressed memory location and any one of the registers A, B, C, D, E, H, or L.

Also, some instructions, such as LDAX B, LDAX D, STAX B, and STAX D, use registers B and C or D and E to hold the address of data. These instructions transfer data between the accumulator and the memory location addressed by registers B and C or D and E using the register indirect mode.

There are also a few instructions, such as the STA ppqq, which use the direct-memory addressing mode to move data between the accumulator and the memory. These instructions use 3 bytes, with the first byte as the OP code followed by 2 bytes of address.

The stack is basically a part of the RAM. Therefore, PUSH and POP instructions are memory reference instructions.

All 8085 JUMP instructions use direct or absolute addressing and are 3 bytes long. The first byte of this instruction is the OP code followed by a 2-byte address. This address specifies the memory location to which the program would branch.

2.4 8085 Addressing Modes

The 8085 has five addressing modes:

1. **Direct** — Instructions using this mode specify the effective address as a part of the instruction. These instructions contain 3 bytes, with the first byte as the OP code followed by 2 bytes of address of data (the low-order byte of the address in byte 2, the high-order byte of the address in byte 3). Consider LDA 2035H. This instruction loads accumulator with the contents of memory location 2035_{16}. This mode is also called the absolute mode.
2. **Register** — This mode specifies the register or register pair that contains data. For example, MOV B, C moves the contents of register C to register B.
3. **Register Indirect** — This mode contains a register pair which stores the address of data (the high-order byte of the address in the first register of the pair, and the low-order byte in the second). As an example, LDAX B loads the accumulator with the contents of a memory location addressed by B, C register pair.
4. **Implied or Inherent** — The instructions using this mode have no operands. Examples include STC (Set the Carry Flag).
5. **Immediate** — For an 8-bit datum, this mode uses 2 bytes, with the first byte as the OP code, followed by 1 byte of data. On the other hand, for 16-bit data, this instruction contains 3 bytes, with the first byte as the OP code followed by 2 bytes of data. For example, MVI B, 05 loads register B with the value 5, and LXI H, 2050H loads H with 20H and L with 50H.

A JUMP instruction interprets the address that it would branch to in the following ways:

1. **Direct** — The JUMP instructions, such as JZ ppqq, use direct addressing and contain 3 bytes. The first byte is the OP code, followed by 2 bytes of the 16-bit address where it would branch to unconditionally or based on a condition if satisfied. For example, JMP 2020 unconditionally branches to location 2020H.
2. **Implied or Inherent Addressing** — This JUMP instruction using this mode is 1 byte long. A 16-bit register pair contains the address of the next instruction to be executed. The instruction PCHL unconditionally branches to a location addressed by the H, L pair.
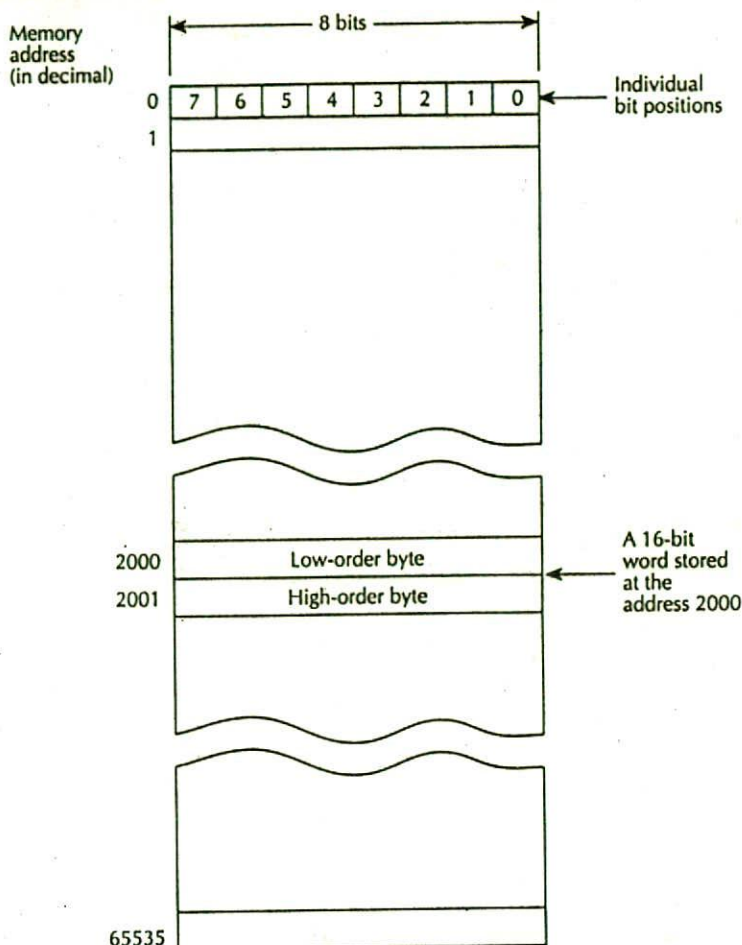
FIGURE 2.3   8085 addressing structures.

2.5 8085 Instruction Set

As mentioned before, the 8085 uses a 16-bit address. Since the 8085 is a byte-addressable machine, it follows that it can directly address 65,536 (2^{16}) distinct memory locations. The addressing structure of the 8085 processor is shown in Figure 2.3.

From this figure, we notice that two consecutive memory locations may be used to represent a 16-bit data item. However, according to the Intel convention, the high-order byte of a 16-bit quantity is always assigned to the high memory address.

The 8085 instructions are 1 to 3 bytes long and these formats are shown in Figure 2.4. The 8085 instruction set contains 74 basic instructions and supports conventional addressing modes such as immediate, register, absolute, and register indirect addressing modes.

Table 2.1 lists the 8085 instructions in alphabetical order; the object codes and instruction cycles are also included. When two instruction cycles are shown, the first is for "condition not met", while the second is for "condition met". Table 2.2 provides the 8085 instructions affecting the status flags. Note that not all 8085 instructions affect the status flags. The 8085 arithmetic and logic instructions normally affect the status flags.

In describing the 8085 instruction set, we will use the symbols in Table 2.3.

The 8085 move instruction transfers 8-bit data from one register to another, register to memory, and vice versa. A complete summary of these instructions is presented in Table 2.4.

(a) Single byte instruction

(b) Two byte instructions

(c) Three byte instructions
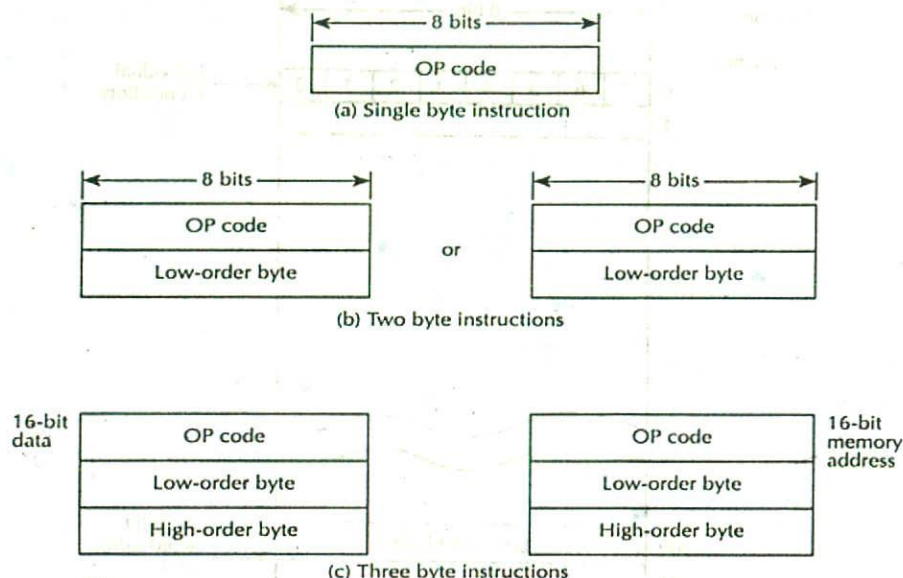
FIGURE 2.4   8085 instructions format.

Table 2.1   Summary of 8085 Instruction Set

Instruction	OP Code	Bytes	Cycles	Operations performed
ACI DATA	CE	2	7	$[A] \leftarrow [A]$ + second instruction byte + $[Cy]$
ADC A	8F	1	4	$[A] \leftarrow [A] + [A] + [Cy]$
ADC B	88	1	4	$[A] \leftarrow [A] + [B] + [Cy]$
ADC C	89	1	4	$[A] \leftarrow [A] + [C] + [Cy]$
ADC D	8A	1	4	$[A] \leftarrow [A] + [D] + [Cy]$
ADC E	8B	1	4	$[A] \leftarrow [A] + [E] + [Cy]$
ADC H	8C	1	4	$[A] \leftarrow [A] + [H] + [Cy]$
ADC L	8D	1	4	$[A] \leftarrow [A] + [L] + [Cy]$
ADC M	8E	1	7	$[A] \leftarrow [A] + [[H\,L]] + [Cy]$
ADD A	87	1	4	$[A] \leftarrow [A] + [A]$
ADD B	80	1	4	$[A] \leftarrow [A] + [B]$
ADD C	81	1	4	$[A] \leftarrow [A] + [C]$
ADD D	82	1	4	$[A] \leftarrow [A] + [D]$
ADD E	83	1	4	$[A] \leftarrow [A] + [E]$
ADD H	84	1	4	$[A] \leftarrow [A] + [H]$
ADD L	85	1	4	$[A] \leftarrow [A] + [L]$
ADD M	86	1	7	$[A] \leftarrow [A] + [[H\,L]]$
ADI DATA	C6	2	7	$[A] \leftarrow [A]$ + second instruction byte
ANA A	A7	1	4	$[A] \leftarrow [A] \wedge [A]$
ANA B	A0	1	4	$[A] \leftarrow [A] \wedge [B]$
ANA C	A1	1	4	$[A] \leftarrow [A] \wedge [C]$
ANA D	A2	1	4	$[A] \leftarrow [A] \wedge [D]$
ANA E	A3	1	4	$[A] \leftarrow [A] \wedge [E]$
ANA H	A4	1	4	$[A] \leftarrow [A] \wedge [H]$
ANA L	A5	1	4	$[A] \leftarrow [A] \wedge [L]$
ANA M	A6	1	4	$[A] \leftarrow [A] \wedge [[H\,L]]$
ANI DATA	E6	2	7	$[A] \leftarrow [A] \wedge$ second instruction byte
CALL ppqq	CD	3	18	Call A subroutine addressed by ppqq
CC ppqq	DC	3	9/18	Call a subroutine addressed by ppqq if $Cy = 1$
CM ppqq	FC	3	9/18	Call a subroutine addressed by ppqq if $S = 1$
CMA	2F	1	4	$[A] \leftarrow$ 1's complement of $[A]$
CMC	3F	1	4	$[Cy] \leftarrow$ 1's complement of $[Cy]$
CMP A	BF	1	4	$[A] - [A]$ and affects flags
CMP B	B8	1	4	$[A] - [B]$ and affects flags

Table 2.1 Summary of 8085 Instruction Set (*continued*)

Instruction	OP Code	Bytes	Cycles	Operations performed
CMP C	B9	1	4	[A] – [C] and affects flags
CMP D	BA	1	4	[A] – [D] and affects flags
CMP E	BB	1	4	[A] – [E] and affects flags
CMP H	BC	1	4	[A] – [H] and affects flags
CMP L	BD	1	4	[A] – [L] and affects flags
CMP M	BE	1	7	[A] – [[H L]] and affects flags
CNC ppqq	D4	3	9/18	Call a subroutine addressed by ppqq if Cy = 0
CNZ ppqq	C4	3	9/18	Call a subroutine addressed by ppqq if Z = 0
CP ppqq	F4	3	9/18	Call a subroutine addressed by ppqq if S = 0
CPE ppqq	EC	3	9/18	Call a subroutine addressed by ppqq if P = 1
CPI DATA	FE	2	7	[A] – second instruction byte and affects flags
CPO ppqq	E4	3	9/18	Call a subroutine addressed by ppqq if P = 0
CZ ppqq	CC	3	9/18	Call a subroutine addressed by ppqq if Z = 1
DAA	27	1	4	Decimal adjust accumulator
DAD B	09	1	10	[HL] ← [HL] + [BC]
DAD D	19	1	10	[HL] ← [HL] + [DE]
DAD H	29	1	10	[HL] ← [HL] + [HL]
DAD SP	39	1	10	[HL] ← [HL] + [SP]
DCR A	3D	1	4	[A] ← [A] – 1
DCR B	05	1	4	[B] ← [B] – 1
DCR C	0D	1	4	[C] ← [C] – 1
DCR D	15	1	4	[D] ← [D] – 1
DCR E	1D	1	4	[E] ← [E] – 1
DCR H	25	1	4	[H] ← [H] – 1
DCR L	2D	1	4	[L] ← [L] – 1
DCR M	35	1	4	[[HL]] ← [[HL]] – 1
DCX B	0B	1	6	[BC] ← [BC] – 1
DCX D	1B	1	6	[DE] ← [DE] – 1
DCX H	2B	1	6	[HL] ← [HL] – 1
DCX SP	3B	1	6	[SP] ← [SP] – 1
DI	F3	1	4	Disable interrupts
EI	FB	1	4	Enable interrupts
HLT	76	1	5	Halt
IN PORT	DB	2	10	[A] ← [specified port]
INR A	3C	1	4	[A] ← [A] + 1
INR B	04	1	4	[B] ← [B] + 1
INR C	0C	1	4	[C] ← [C] + 1
INR D	14	1	4	[D] ← [D] + 1
INR E	1C	1	4	[E] ← [E] + 1
INR H	24	1	4	[H] ← [H] + 1
INR L	2C	1	4	[L] ← [L] + 1
INR M	34	1	4	[[HL]] ← [[HL]] +1
INX B	03	1	6	[BC] ← [BC] +1
INX D	13	1	6	[DE] ← [DE] + 1
INX H	23	1	6	[HL] ← [HL] + 1
INX SP	33	1	6	[SP] ← [SP] + 1
JC ppqq	DA	3	7/10	Jump to ppqq if Cy = 1
JM ppqq	FA	3	7/10	Jump to ppqq if S = 1
JMP ppqq	C3	3	10	Jump to ppqq
JNC ppqq	D2	3	7/10	Jump to ppqq if Cy = 0
JNZ ppqq	C2	3	7/10	Jump to ppqq if Z = 0
JP ppqq	F2	3	7/10	Jump to ppqq if S = 0
JPE ppqq	EA	3	7/10	Jump to ppqq if P = 1
JPO ppqq	E2	3	7/10	Jump to ppqq if P = 0
JZ ppqq	CA	3	7/10	Jump to ppqq if Z = 1
LDA ppqq	3A	3	13	[A] ← [ppqq]
LDAX B	0A	1	7	[A] ← [[BC]]
LDAX D	1A	1	7	[A] ← [[DE]]
LHLD ppqq	2A	3	16	[L] ← [ppqq],[H] ← [ppqq + 1]

Table 2.1   Summary of 8085 Instruction Set (*continued*)

Instruction	OP Code	Bytes	Cycles	Operations performed
LXI B	01	3	10	[BC] ← second and third instruction bytes
LXI D	11	3	10	[DE] ← second and third instruction bytes
LXI H	21	3	10	[HL] ← second and third instruction bytes
LXI SP	31	3	10	[SP] ← second and third instruction bytes
MOV A,A	7F	1	4	[A] ← [A]
MOV A,B	78	1	4	[A] ← [B]
MOV A,C	79	1	4	[A] ← [C]
MOV A,D	7A	1	4	[A] ← [D]
MOV A,E	7B	1	4	[A] ← [E]
MOV A,H	7C	1	4	[A] ← [H]
MOV A,L	7D	1	4	[A] ← [L]
MOV A,M	7E	1	7	[A] ← [[HL]]
MOV B,A	47	1	4	[B] ← [A]
MOV B,B	40	1	4	[B] ← [B]
MOV B,C	41	1	4	[B] ← [C]
MOV B,D	42	1	4	[B] ← [D]
MOV B,E	43	1	4	[B] ← [E]
MOV B,H	44	1	4	[B] ← [H]
MOV B,L	45	1	4	[B] ← [L]
MOV B,M	46	1	7	[B] ← [[HL]]
MOV C,A	4F	1	4	[C] ← [A]
MOV C,B	48	1	4	[C] ← [B]
MOV C,C	49	1	4	[C] ← [C]
MOV C,D	4A	1	4	[C] ← [D]
MOV C,E	4B	1	4	[C] ← [E]
MOV C,H	4C	1	4	[C] ← [H]
MOV C,L	4D	1	4	[C] ← [L]
MOV C,M	4E	1	7	[C] ← [[HL]]
MOV D,A	57	1	4	[D] ← [A]
MOV D,B	50	1	4	[D] ← [B]
MOV D,C	51	1	4	[D] ← [C]
MOV D,D	52	1	4	[D] ← [D]
MOV D,E	53	1	4	[D] ← [E]
MOV D,H	54	1	4	[D] ← [H]
MOV D,L	55	1	4	[D] ← [L]
MOV D,M	56	1	7	[D] ← [[HL]]
MOV E,A	5F	1	4	[E] ← [A]
MOV E,B	58	1	5	[E] ← [B]
MOV E,C	59	1	4	[E] ← [C]
MOV E,D	5A	1	4	[E] ← [D]
MOV E,E	5B	1	4	[E] ← [E]
MOV E,H	5C	1	4	[E] ← [H]
MOV E,L	5D	1	4	[E] ← [L]
MOV E,M	5E	1	7	[E] ← [[HL]]
MOV H,A	67	1	4	[H] ← [B]
MOV H,B	60	1	4	[H] ← [A]
MOV H,C	61	1	4	[H] ← [C]
MOV H,D	62	1	4	[H] ← [D]
MOV H,E	63	1	4	[H] ← [E]
MOV H,H	64	1	4	[H] ← [H]
MOV H,L	65	1	4	[H] ← [L]
MOV H,M	66	1	7	[H] ← [[HL]]
MOV L,A	6F	1	4	[L] ← [A]
MOV L,B	68	1	4	[L] ← [B]
MOV L,C	69	1	4	[L] ← [C]
MOV L,D	6A	1	4	[L] ← [D]
MOV L,E	6B	1	4	[L] ← [E]
MOV L,H	6C	1	4	[L] ← [H]
MOV L,L	6D	1	4	[L] ← [L]

Table 2.1 Summary of 8085 Instruction Set (*continued*)
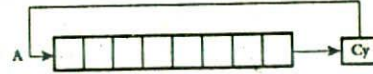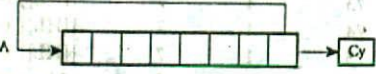
Instruction	OP Code	Bytes	Cycles	Operations performed
MOV L,M	6E	1	7	[L] ← [[HL]]
MOV M,A	77	1	7	[[HL]] ← [A]
MOV M,B	70	1	7	[[HL]] ← [B]
MOV M,C	71	1	7	[[HL]] ← [C]
MOV M,D	72	1	7	[[HL]] ← [D]
MOV M,E	73	1	7	[[HL]] ← [E]
MOV M,H	74	1	7	[[HL]] ← [H]
MOV M,L	75	1	7	[[HL]] ← [L]
MVI A, DATA	3E	2	7	[A] ← second instruction byte
MVI B, DATA	06	2	7	[B] ← second instruction byte
MVI C, DATA	0E	2	7	[C] ← second instruction byte
MVI D, DATA	16	2	7	[D] ← second instruction byte
MVI E, DATA	1E	2	7	[E] ← second instruction byte
MVI H, DATA	26	2	7	[H] ← second instruction byte
MVI L, DATA	2E	2	7	[L] ← second instruction byte
MVI M, DATA	36	2	10	[[HL]] ← second instruction byte
NOP	00	1	4	No operation
ORA A	B7	1	4	[A] ← [A] ∨ [A]
ORA B	B0	1	4	[A] ← [A] ∨ [B]
ORA C	B1	1	4	[A] ← [A] ∨ [C]
ORA D	B2	1	4	[A] ← [A] ∨ [D]
ORA E	B3	1	4	[A] ← [A] ∨ [E]
ORA H	B4	1	4	[A] ← [A] ∨ [H]
ORA L	B5	1	4	[A] ← [A] ∨ [L]
ORA M	B6	1	7	[A] ← [A] ∨ [[HL]]
ORI DATA	F6	2	7	[A] ← [A] ∨ second instruction byte
OUT PORT	D3	2	10	[specified port] ← [A]
PCHL	E9	1	6	[PCH]* ← [H], [PCL]* ← [L]
POP B	C1	1	10	[C] ← [[SP]], [SP] ← [SP] + 2 [B] ← [[SP] + 1]
POP D	D1	1	10	[E] ← [[SP]], [SP] ← [SP] + 2 [D] ← [[SP] + 1]
POP H	E1	1	10	[L] ← [[SP]], [SP] ← [SP] + 2 [H] ← [[SP] + 1]
POP PSW	F1	1	10	[A] ← [[SP] + 1], [PSW] ← [[SP]], [SP] ← [SP] + 2
PUSH B	C5	1	12	[[SP] − 1] ← [B], [SP] ← [SP] − 2 [[SP] − 2] ← [C]
PUSH D	D5	1	12	[[SP] − 1] ← [D], [[SP] − 2] ← [E] [SP] ← [SP] − 2
PUSH H	E5	1	12	[[SP] − 1] ← [H], [SP] ← [SP] − 2 [[SP] − 2] ← [L]
PUSH PSW	F5	1	12	[[SP] − 1] ← [A], [SP] ← [SP] − 2 [[SP] − 2] ← [PSW]
RAL	17	1	4	
RAR	1F	1	4	
RC	D8	1	6/12	Return if carry; [PC] ← [SP]]
RET	C9	1	10	[PCL]* ← [[SP]], [SP] ← [SP] + 2 [PCH]* ← [[SP] + 1]
RIM	20	1	4	Read interrupt mask
RLC	07	1	4	
RM	F8	1	6/12	Return if minus; [PC] ← [[SP]]

Table 2.1   Summary of 8085 Instruction Set (*continued*)

Instruction	OP Code	Bytes	Cycles	Operations performed
RNC	D0	1	6/12	Return if no carry; $[PC] \leftarrow [[SP]]$
RNZ	C0	1	6/12	Return if result not zero; $[PC] \leftarrow [[SP]]$
RP	F0	1	6/12	Return if positive; $[PC] \leftarrow [[SP]]$, $[SP] \leftarrow [SP] + 2$
RPE	E8	1	6/12	Return if parity even; $[PC] \leftarrow [[SP]]$, $[SP] \leftarrow [SP] + 2$
RPO	E0	1	6/12	Return if parity odd; $[PC] \leftarrow [[SP]]$, $[SP] \leftarrow [SP] + 2$
RRC	0F	1	4	
RST0	C7	1	12	Restart
RST1	CF	1	12	Restart
RST2	D7	1	12	Restart
RST3	DF	1	12	Restart
RST4	E7	1	12	Restart
RST5	EF	1	12	Restart
RST6	F7	1	12	Restart
RST7	FF	1	12	Restart
RZ	C8	1	6/12	Return if zero; $[PC] \leftarrow [[SP]]$
SBB A	9F	1	4	$[A] \leftarrow [A] - [A] - [Cy]$
SBB B	98	1	4	$[A] \leftarrow [A] - [B] - [Cy]$
SBB C	99	1	4	$[A] \leftarrow [A] - [C] - [Cy]$
SBB D	9A	1	4	$[A] \leftarrow [A] - [D] - [Cy]$
SBB E	9B	1	4	$[A] \leftarrow [A] - [E] - [Cy]$
SBB H	9C	1	4	$[A] \leftarrow [A] - [H] - [Cy]$
SBB L	9D	1	4	$[A] \leftarrow [A] - [L] - [Cy]$
SBB M	9E	1	7	$[A] \leftarrow [A] - [[HL]] - [Cy]$
SBI DATA	DE	2	7	$[A] \leftarrow [A] -$ second instruction byte $- [Cy]$
SHLD ppqq	22	3	16	$[ppqq] \leftarrow [L], [ppqq + 1] \leftarrow [H]$
SIM	30	1	4	Set interrupt mask
SPHL	F9	1	6	$[SP] \leftarrow [HL]$
STA ppqq	32	3	13	$[ppqq] \leftarrow [A]$
STAX B	02	1	7	$[[BC]] \leftarrow [A]$
STAX D	12	1	7	$[[DE]] \leftarrow [A]$
STC	37	1	4	$[Cy] \leftarrow 1$
SUB A	97	1	4	$[A] \leftarrow [A] - [A]$
SUB B	90	1	4	$[A] \leftarrow [A] - [B]$
SUB C	91	1	4	$[A] \leftarrow [A] - [C]$
SUB D	92	1	4	$[A] \leftarrow [A] - [D]$
SUB E	93	1	4	$[A] \leftarrow [A] - [E]$
SUB H	94	1	4	$[A] \leftarrow [A] - [H]$
SUB L	95	1	4	$[A] \leftarrow [A] - [L]$
SUB M	96	1	7	$[A] \leftarrow [A] - [[HL]]$
SUI DATA	D6	2	7	$[A] \leftarrow [A] -$ second instruction byte
XCHG	EB	1	4	$[D] \leftrightarrow [H], [E] \leftrightarrow [L]$
XRA A	AF	1	4	$[A] \leftarrow [A] \oplus [A]$
XRA B	A8	1	4	$[A] \leftarrow [A] \oplus [B]$
XRA C	A9	1	4	$[A] \leftarrow [A] \oplus [C]$
XRA D	AA	1	4	$[A] \leftarrow [A] \oplus [D]$
XRA E	AB	1	4	$[A] \leftarrow [A] \oplus [E]$
XRA H	AC	1	4	$[A] \leftarrow [A] \oplus [H]$
XRA L	AD	1	4	$[A] \leftarrow [A] \oplus [L]$
XRA M	AE	1	7	$[A] \leftarrow [A] \oplus [[HL]]$
XRI DATA	EE	2	7	$[A] \leftarrow [A] \oplus$ second instruction byte
XTHL	E3	1	16	$[[SP]] \leftrightarrow [L], [[SP] + 1]] \leftrightarrow [H]$

*PCL — program counter low byte; PCH — program counter high byte.

Table 2.2   8085 Instructions Affecting the Status Flags

Instructions[a]	Cy	Ac	Z	S	P
	\multicolumn status flags[b]				
ACI DATA	+	+	+	+	+
ADC reg	+	+	+	+	+
ADC M	+	+	+	+	+
ADD reg	+	+	+	+	+
ADD M	+	+	+	+	+
ADI DATA	+	+	+	+	+
ANA reg	0	1	+	+	+
ANA M	0	1	+	+	+
ANI DATA	0	1	+	+	+
CMC	+				
CMP reg	+	+	+	+	+
CMP M	+	+	+	+	+
CPI DATA	+	+	+	+	+
DAA	+	+	+	+	+
DAD rp	+				
DCR reg		+	+	+	+
DCR M		+	+	+	+
INR reg		+	+	+	+
INR M		+	+	+	+
ORA reg	0	0	+	+	+
ORA M	0	0	+	+	+
ORI DATA	0	0	+	+	+
RAL	+				
RAR	+				
RLC	+				
RRC	+				
SBB reg	+	+	+	+	+
SBB M	+	+	+	+	+
SBI DATA	+	+	+	+	+
STC	+				
SUB reg	+	+	+	+	+
SUB M	+	+	+	+	+
SUI DATA	+	+	+	+	+
XRA reg	0	0	+	+	+
XRA M	0	0	+	+	+
XRI DATA	0	0	+	+	+

[a]reg — 8-bit register; M — memory; rp — 16-bit register pair.
[b]Note that instructions which are not shown in the table do not affect the flags; + indicates that the particular flag is affected; 0 or 1 indicates that these flags are always 0 or 1 after the corresponding instructions are executed.

All mnemonics copyright Intel Corporation 1976.

Table 2.3   Symbols to be Used in 8085 Instruction Set

Symbol	Interpretation
r1, r2	8-bit register
rp	Register pair
data8	8-bit data
data16	16-bit data
M	Memory location indirectly addressed through the register pair H,L
addr16	16-bit memory address

Table 2.4   8085 MOVE Instructions

Instruction	Symbolic description	Addressing mode		Example	Comments
		Source	Destination		Illustration
MOV r1, r2	(r1) ← (r2)	Register	Register	MOV A, B	Copy the contents of the register B into the register A
MOV r, M	(r) ← M((HL))	Register indirect	Register	MOV A, M	Copy the contents of the memory location whose address is specified in the register pair H,L into the A register
MVI r, data8	(r) ← data8	Immediate	Register	MVI A, 08	Initialize the A register with the value 08
MOV M, r	M((HL)) ← (r)	Register	Register indirect	MOV M, B	Copy the contents of the B register into the memory location addressed by H,L pair
MVI M, data8	M((HL)) ← data8	Immediate	Register indirect	MVI M, 07	Initialize the memory location whose address is specified in the register pair H,L with the value 07

Table 2.5　8085 Load and Store Instructions

Instruction	Symbolic description	Restriction	Addressing mode	Example	Illustration	
					Comments	
LDA addr16	$(A) \leftarrow M(addr16)$	The destination is always the accumulator register	Absolute	LDA 2000H	Load the accumulator with the contents of the memory location whose address is 2000H	
LHLD addr16	$(L) \leftarrow M(addr16)$ $(H) \leftarrow M(addr16+1)$	The destination is always the register pair H,L	Absolute	LHLD 2000H	Load the H and L registers with contents of the memory locations 2001H and 2000H, respectively	
LXI rp, data16	$(rp) \leftarrow data16$	rp may be HL, DE, BC, or SP	Immediate	LXI H, 2024H	$H \leftarrow 20_{16}, L \leftarrow 24_{16}$	
LDAX rp	$(A) \leftarrow M((rp))$	Destination is always the accumulator; also rp may be either B,C or D,E	Register indirect	LDAX B	Load the accumulator with the contents of the memory location whose address is specified with the register pair B,C	
STA addr16	$M(addr16) \leftarrow (A)$	Source is always the accumulator register	Absolute	STA 2001H	Save the contents of the accumulator into the memory location whose address is 2001H	
SHLD addr16	$M(addr16) \leftarrow (L)$ $M(addr16+1) \leftarrow (H)$	The source is always the register pair H,L	Absolute	SHLD 2000	$M(2000) \leftarrow (L)$ $M(2001) \leftarrow (H)$	
STAX rp	$M((rp)) \leftarrow (A)$	The source is always the accumulator and the register pair may be B,C or D,E	Register indirect	STAX D	Save the contents of the accumulator register into the memory location whose address is specified with the register pair D,E	
XCHG	$[D] \leftrightarrow [H], [E] \leftrightarrow [L]$		Inherent	XCHG	Exchanges [DE] with [HL]	

Table 2.6   8085 Arithmetic Instructions

Operation	Instruction	Interpretation	Addressing mode	Example	Comments
					Illustration
8-bit addition	ADD r	(A) ← (A) + (r)	Register	ADD B	(A) ← (A) + (B)
	ADI data8	(A) ← (A) + data8	Immediate	ADI 05	(A) ← (A) + 05
	ADD M	(A) ← (A) + M((HL))	Register indirect	—	—
8-bit addition with	ADC r	(A) ← (A) + (r) + Cy	Register	ADC C	(A) ← (A) + (C) + Cy
a carry	ACI data8	(A) ← (A) + data8 + Cy	Immediate	ACI 07	(A) ← (A) + 07 + Cy
	ADC M	(A) ← (A) + M((HL)) + Cy	Register indirect	—	—
8-bit subtraction	SUB r	(A) ← (A) – (r)	Register	SUB C	(A) ← (A) – (C)
	SUI data8	(A) ← (A) – data8	Immediate	SUI 03	(A) ← (A) – 03
	SUB M	(A) ← (A) – M((HL))	Register indirect	—	—
8-bit subtraction	SBB r	(A) ← (A) – (r) – Cy	Register	SBB D	(A) ← (A) – (D) – Cy
with a borrow	SBI data8	(A) ← (A) – data8 – Cy	Immediate	SBI 04	(A) ← (A) – 04 – Cy
	SBB M	(A) ← (A) – M((HL)) – Cy	Register indirect	—	—
16-bit addition	DAD rp	(HL) ← (HL) + (rp)	Register	DAD B	(HL) ← (HL) + (BC)
Decimal adjust	DAA	Convert the 8-bit number stored in the accumulator into BCD	Inherent	—	—
8-bit increment	INR r	(r) ← (r) + 1	Register	INR B	(B) ← (B) + 1
	INR M	M((HL)) ← M((HL)) + 1	Register indirect	—	—
16-bit increment*	INX rp	(rp) ← (rp) + 1	Register	INX D	(DE) ← (DE) + 1
8-bit decrement	DCR r	(r) ← (r) – 1	Register	DCR B	(B) ← (B) – 1
	DCR M	M((HL)) ← M((HL)) – 1	Register indirect	—	—
16-bit decrement*	DCX rp	(rp) ← (rp) – 1	Register	—	

*rp = BC, DE, HL, or SP.

The 8085 instruction set also accomplishes the 8- and 16-bit data transfers using the load and store instructions. These instructions are summarized in Table 2.5.

From Table 2.5 notice that we adopt the following convention when we specify a register pair in the instruction.

Symbol	Register pairs used
B	B,C
D	D,E
H	H,L

Also, observe that in Table 2.5 the 8085 processor does not provide LDAX H instruction. This is because the same result can be obtained by using the MOV A,M instruction. The 8085 includes a one-byte exchange instruction, namely, XCHG. The XCHG exchanges the contents of DE with HL. That is, it performs the following operation:

$$[D] \leftrightarrow [H]$$
$$[E] \leftrightarrow [L]$$

The arithmetic instructions provided by the 8085 processor allow one to add (or subtract) two 8-bit data with or without carry (or borrow). The subtraction operation is realized by adding the two's complement of the subtrahend to the minuend. During the subtraction operation, the carry flag will be treated as the borrow flag. Table 2.6 lists these instructions. For some instructions such as ADD M, examples and comments are not included. This is due to limited space in the table.

As far as logical operations are concerned, the 8085 includes some instructions to perform traditional Boolean operations such as AND, OR, EXCLUSIVE-OR. In addition, instructions are available to complement the accumulator and to set the carry flag. The 8085 COMPARE instructions subtract the specified destination from the contents of the

Table 2.7   8085 Logical Instructions

Operation	Instruction	Interpretation	Addressing mode	Illustration Example	Comments
Boolean AND	ANA r	$(A) \leftarrow (A) \wedge (r)$	Register	ANA B	$(A) \leftarrow (A) \wedge (B)$
	ANI data8	$(A) \leftarrow (A) \wedge \text{data8}$	Immediate	ANI 0FH	$(A) \leftarrow (A) \wedge 00001111_2$
	ANA M	$(A) \leftarrow (A) \wedge M((HL))$	Register indirect	—	—
Boolean OR	ORA r	$(A) \leftarrow (A) \vee (r)$	Register	ORA C	$(A) \leftarrow (A) \vee (C)$
	ORI data8	$(A) \leftarrow (A) \vee \text{data8}$	Immediate	ORI 08H	$(A) \leftarrow (A) \vee 00001000_2$
	ORA M	$(A) \leftarrow (A) \vee M((HL))$	Register indirect	—	—
Boolean EXCLUSIVE-OR	XRA r	$(A) \leftarrow (A) \oplus (r)$	Register	XRA A	$(A) \leftarrow (A) \oplus (A)$
	XRI data8	$(A) \leftarrow (A) \oplus \text{data8}$	Immediate	XRI 03H	$(A) \leftarrow (A) \oplus 00000011_2$
	XRA M	$(A) \leftarrow (A) \oplus M((HL))$	Register indirect	—	—
Compare	CMP r	$(A) - (r)$ and affect flags	Register	CMP D	Compare (A) register with (D) register
	CPI data8	$(A) - \text{data8}$ and affect flags	Immediate	CPI 05	Compare (A) with 05
	CMP M	$(A) - M((HL))$ and affect flags	Register indirect	—	—
Complement	CMA	$(A) \leftarrow (\overline{A})$	Inherent	—	—
Bit manipulation	STC	$Cy \leftarrow 1$ (set carry to 1)	Inherent	—	—
	CMC	$Cy \leftarrow Cy'$ (complement carry flag)	Inherent	—	—

accumulator and affect the status flags according to the result. However, in this case the result of the subtraction is not provided in the accumulator. All 8085 logical instructions are specified in Table 2.7. For some instructions in this figure, examples and comments are not provided. This is due to limited space in the table.

The AND instruction can be used to perform a masking operation. If the bit value in a particular bit position is desired in a word, the word can be logically ANDed with appropriate masking data to accomplish this. For example, the bit value at bit 3 of the word $1011\ X011_2$ can be determined as follows:
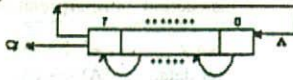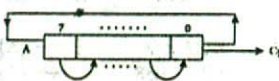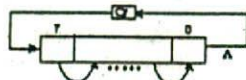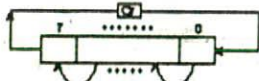
$$
\begin{array}{lll}
 & 1011 \quad \text{x}011 & \text{Word} \\
\text{AND} & 0000 \quad 1000 & \text{Masking Data} \\
\hline
 & 0000 \quad \text{x}000 & \text{Result}
\end{array}
$$

If the bit value X at bit 3 is 1, then the result is nonzero ($Z = 0$); otherwise the result is zero ($Z = 1$). The Z-flag can be tested using JZ (jump if $Z = 1$) or JNZ (jump if $Z = 0$) to determine whether $X = 0$ or 1. The AND instruction can also be used to determine whether a binary number is odd or even by checking the least significant bit (LSB) of the number (LSB = 0 for even and LSB = 1 for odd). XRA instruction can be used to find ones compliment of a binary number by exclusive-ORing the number with all ones as follows:

$$
\begin{array}{lll}
 & 1010 \quad 1001 & \text{Original number} \\
\text{XOR} & 1111 \quad 1111 & \\
\hline
 & 0101 \quad \text{x}110 & \text{Ones complement}
\end{array}
$$

One of the applications of the compare instruction is to find a match in an array. The number to be matched can be loaded in the accumulator and compared with each element in the array. The JZ (jump if $Z = 1$) can then be executed to find the match. If the subtraction instruction is used in place of the COMPARE, the number to be matched in the accumulator

TABLE 2.8 8085 Rotate Instructions

Instruction	Interpretation	Illustration
RLC	Rotate left accumulator by one position without the carry flag Cy	
RRC	Rotate right accumulator by one position without the carry flag Cy	
RAR	Rotate right accumulator by one position through the carry flag	
RAL	Rotate left accumulator by one position through the carry flag	

will be lost after each subtraction, and therefore the number needs to be loaded for the next subtraction.

The 8085 instruction set includes rotating the contents of the A register to the left or right without or through the carry flag. These instructions are listed in Table 2.8.

In the 8085, only the absolute mode branch instruction is of the form JMP addr16. There is also a one-byte implied unconditional jump instruction, namely, PCHL. The PCHL loads [H] into PC high byte and [L] into PC low byte. That is, PCHL performs an unconditional JUMP to a location addressed by the contents of H and L. The general format of an 8085 conditional branch instruction is J <condition code> addr16 where the condition code may represent one of the following conditions:

Conditional jumps	Condition	Comment
JZ	$Z = 1$	Z flag is set (result equal to zero)
JNZ	$Z = 0$	Z flag is reset (result not equal to zero)
JC	$Cy = 1$	Cy flag is set
JNC	$Cy = 0$	Cy flag is reset
JPO	$P = 0$	The parity is odd
JPE	$P = 1$	The parity is even
JP	$S = 0$	S flag is reset (or the number is positive)
JM	$S = 1$	S flag is set (or the number is negative, or minus)

For example, the following instruction sequence causes a branch to the memory address 2000_{16} only if the contents of the A and B registers are equal:

```
CMP B
JZ 2000H
```

In the 8085, the subroutine call instruction is of the form

```
CALL addr16
```

The instruction RET transfers the control to the caller, and it should be the last instruction of the subroutine. Both instructions use the PC and SP for subroutine linkage.

For example, consider the following:

Main program		Subroutine
	SUB	—
—		—
—		—
	CALL SUB	—
START	—	—
—		—
—		**RET**
—		

The call SUB instruction pushes or saves the current PC contents (START which is the address of the next instruction) onto the stack and loads PC with the starting address of the Subroutine (SUB) specified with the CALL instruction. The RET instruction at the end of the subroutine pops or reads the address START (saved onto the stack by the CALL instruction) into PC and transfers control to the right place in the main program.

There are a number of conditional call instructions. These include:

```
CC  addr (call if Cy = 1)
CNC addr (call if Cy = 0)
CZ  addr (call if Z = 1)
CNZ addr (call if Z = 0)
CM  addr (call if S = 1)
CP  addr (call if S = 0)
CPE addr (call if P = 1)
CPO addr (call if P = 0)
```

Also, there are a number of conditional return instructions. These include:

```
RC  (Return if Cy = 1)
RNC (Return if Cy = 0)
RZ  (Return if Z = 1)
RNZ (Return if Z = 0)
RPE (Return if P = 1)
RPO (Return if P = 0)
RM  (Return if S = 1)
RP  (Return if S = 0)
```

There are eight one-byte call instructions (RST 0 to 7) which have predefined addresses. The format for these instructions is

```
11  XXX  111  =  000 for RST0
              =  001 for RST1
              =  010 for RST2
              =  011 for RST3
              =  100 for RST4
              =  101 for RST5
              =  110 for RST6
              =  111 for RST7
```

RSTs are one-byte call instructions used mainly with interrupts. Each RST has a predefined address. However, RST0 and the hardware reset vector have the same address 0000_{16}. Therefore, use of RST0 is not usually recommended. The RSTs cause the 8085 to push the PC onto the stack. The 8085 then loads the PC with a predefined address based on the particular RST being used.

	7	6	5	4	3	2	1	0	— bit number
OP code for RST	1	1	X	X	X	1	1	1	

	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
PC after RST	0	0	0	0	0	0	0	0	0	0	X	X	X	0	0	0
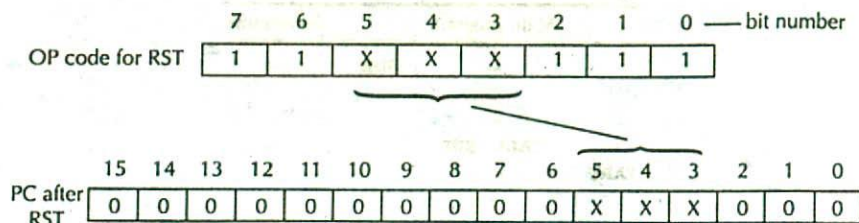
FIGURE 2.5 Execution of the RST instruction.

Table 2.9 RST0–RST7 Vector Addresses

Instruction	OP code (hexadecimal)	Vector address (hexadecimal)
RST0	C7	0000
RST1	CF	0008
RST2	D7	0010
RST3	DF	0018
RST4	E7	0020
RST5	EF	0028
RST6	F7	0030
RST7	FF	0038

A 3-bit code in the OP code for a particular RST determines the address to which the program would branch. This is shown in Figure 2.5.

The vector addresses for the RSTs are listed in Table 2.9. Note that a limited number of locations are available for each RST instruction. This may not provide enough locations for writing a complete subroutine. Therefore, one may place an unconditional JUMP at the vector address to JUMP to an address in RAM where the actual subroutine is written.

The 8085 stack manipulation instructions allow one to save and retrieve the contents of the register pairs into and from the stack, respectively. For example, the following instruction saves the register pair B,C into the stack:

```
PUSH  B ; (SP) ← (SP) - 1
        ; M((SP)) ← (B)
        ; (SP) ← (SP) - 1
        ; M((SP)) ← (C)
```

Similarly, the instruction POP D retrieves the top two words of the stack and places them into the registers E and D in that order as follows:

```
POP  D ; (E) ← M((SP))
       ; (SP) ← (SP) + 1
       ; (D) ← M((SP))
       ; (SP) ← (SP) + 1
```

This means that all 8085 registers can be saved onto the stack using the following instruction sequence:

```
PUSH PSW ; save the A and flags register
PUSH B   ; save the D,E pair
PUSH D   ; save the D,E pair
PUSH H   ; save the H,L pair
```

Similarly, the saved status can be restored by using the following sequence of POP instructions:

```
POP H   ; restore H,L pair
POP D   ; restore D,E pair
POP B   ; restore B,C pair
POP PSW ; restore A and flags register
```

There are two other stack instructions: SPHL and XTHL. SPHL is a one-byte instruction. It moves the [L] to SP high byte and [H] to SP low byte. XTHL is also a one-byte instruction. It exchanges the [L] with the top of the stack addressed by SP and [H] with the next stack addressed by SP + 1. That is, XTHL performs the following:

$$[[SP]] \leftrightarrow [L]$$
$$[[SP+1]] \leftrightarrow [H]$$

The 8085 can use either standard or memory-mapped I/O. Using standard I/O, the input and output instructions have the following format:

```
IN  (8-bit port address) ;      input instruction
OUT (8-bit port address) ;      output instruction
```

For example, the instruction IN 02H transfers the contents of the input port with address 02_{16} into the accumulator. Similarly, the instruction OUT 00H transfers the contents of the accumulator to the output port with address 00_{16}. Using memory-mapped I/O, LDA addr and STA addr can be used as input and output instructions, respectively.

The 8085 HLT instruction forces the 8085 to enter into the halt state. Similarly, the dummy instruction NOP neither achieves any result nor affects any CPU registers. This is a useful instruction for producing software delay routines and to insert diagnostic messages.

The 8085 8-bit increment (INR) and decrement (DCR) instructions affect the status flags. However, the 16-bit increment (INX) and decrement (DCX) instructions do not affect the flags. Therefore, while using these instructions in a loop counter value greater than 256_{10}, some other instructions must be used with DCX or INX to affect the flags after their execution. For example, the following instruction sequence will affect the flags for DCX:

```
          LXI B,16-BIT DATA ;  Load
Loop          —            ;  initial
              —            ;  16-bit loop
              —            ;  count to BC
              —
              —
              —            ;  Decrement
          DCX B            ;  counter
          MOV A,B          ;  Move B to A to
                              test for zero
          ORA C            ;  Logically or with A
          JNZ Loop         ;  Jump if not zero
              —
              —
              —
```
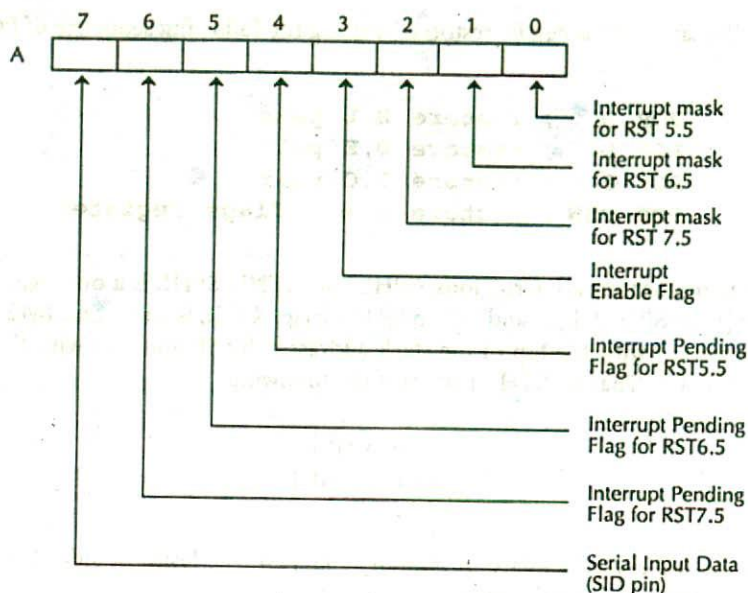
FIGURE 2.6   Accumulator data format after execution of RIM.

There are four one-byte 8085 interrupt instructions. These are DI, EI, RIM, and SIM.

DI disables the 8085's maskable interrupt capability. EI, on the other hand, enables the 8085 maskable interrupt capability.

RIM is a one-byte instruction. It loads the accumulator with 8 bits of data as shown in Figure 2.6.

Bits 0, 1, and 2 provide the values of the RST 5.5, RST 6.5, and RST 7.5 mask bits, respectively. If the mask bit corresponding to a particular RST is one, the RST is disabled; a zero in a specific RST (bits 0, 1, and 2) means that RST is enabled.

If the interrupt enable bit (bit 3) is 0, the 8085's maskable interrupt capability is disabled; the interrupt is enabled if this bit is one.

A "one" in a particular interrupt pending bit indicates that an interrupt is being requested on the identified RST line; if this bit is zero, no interrupt is waiting to be serviced. The serial input data (bit 7) indicate the value of the SID pin.

The SIM instruction outputs the contents of the accumulator to define interrupt mask bits and the serial output data line. The bits in the accumulator before execution of the SIM are defined as shown in Figure 2.7.
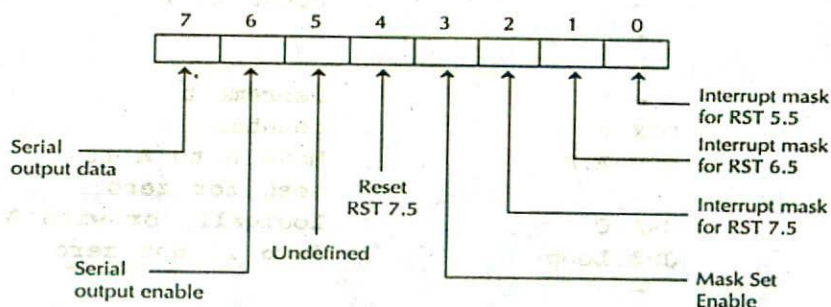
FIGURE 2.7   Accumulator data format before execution of the SIM.

If the mask set enable bit is set to one, interrupt mask bits for RST 7.5, RST 6.5, and RST 5.5 are sent out; a zero value at the mask set enable does not affect the interrupt mask bits. A one at a particular interrupt mask disables that interrupt and a zero enables it.

The RESET RST 7.5, if set to one, resets as internal flip-flop to zero in order to clear the 7.5 interrupt.

If the serial output enable is one, the serial output data are sent to the SOD pin.

The interrupt instructions will be covered in detail during discussion of the 8085 interrupts.

Example 2.1

Write an 8085 assembly language program to add a 16-bit number in locations 5000H (high byte) and 5001H (low byte) with another 16-bit number stored in 5002H (high byte) and 5003H (low byte). Store result in BC.

Solution

```
        ORG    2000H
        LDA    MEM2      ;  Load low byte of
        LXI    H,5003H   ;  number 1
        ADD    M         ;  Add with low
                         ;  byte of number 2
        MOV    C, A      ;  Store result in C
        LDA    MEM1      ;  Load high byte
        DCX    H         ;  of number 1
        ADC    M         ;  Add with
                         ;  high byte and Cy
        MOV    B, A      ;  Store result
                         ;  in B
        HLT              ;  stop
        ORG    5000H
MEM1 DB    DATA1
MEM2 DB    DATA2
MEM3 DB    DATA3
MEM4 DB    DATA4
```

Example 2.2

Write an 8085 assembly language program to perform a parity check on an 8-bit word in location 5000_{16}. If the parity is odd, store DD_{16} in location 5000_{16}. However, if the parity is even, store EE_{16} in location 5000_{16}.

Solution

```
        ORG 3000H
        LDA PARITY    ;  Load 8-bit
                      ;  data into A
        ADI 00H       ;  Add with
                      ;  00H to affect
        MVI A, 0DDH   ;  flags
        JPO ODD       ;  Check for
                      ;  odd parity
        MVI A,0EEH    ;  If parity
                      ;  even, store
                      ;  EEH
```

```
ODD        STA PARITY   ;  in 5000H
           HLT          ;  Stop
           ORG 5000H
PARITY     DB DATA
```

Note that address 3000H at ORG 3000H is arbitrarily chosen.

Example 2.3

Write a program in 8085 assembly language to perform an unsigned 8-bit by 8-bit multiplication via repeated addition. Assume that the multiplicand is in "B" register and the multiplier is in "C" register. Store the product in HL.

Solution

```
           ORG 5000H
           LXI H,0000H  ;  Initialize the
                        ;  16-bit product to zero
           MVI C, MULT  ;  Move multiplier to C
           MOV E,B      ;  Move multiplicand
                        ;  to register E
           MVI D,OOH    ;  Convert multiplicand
                        ;  to unsigned 16-bit
START      DAD D        ;  Perform 16-bit addition
           DCR C        ;  Decrement multiplier
           JNZ START    ;  Jump if multiplier not zero
           HLT          ;  Stop. Product
                        ;  is in HL
```

Example 2.4

Write a subroutine in 8085 assembly language to check whether an 8-bit number in location 3000_{16} is odd or even. If the number is odd, store DD_{16} in location 3000_{16}; if the number is even, store EE_{16} in location 3000_{16}.

Also, write the main program to initialize stack pointer 5020_{16}, load the 8-bit number to be checked for odd or even into A, store DD or EE depending on the result, and stop.

Solution

`Main Program`

```
           ORG 4000H
           LXI SP, 5020H ;  Initialize SP
           LDA START     ;  Load number
           CALL CHECK    ;  Call subroutine
           STA START     ;  Store DD or EE
           HLT           ;  stop
           ORG 3000H
START      DB DATA
```

`Subroutine`

```
           ORG 7000H
CHECK      RAR           ;  Rotate 'A' for checking
           MVI A,OEEH    ;  Jump if even
```

```
              JNC  RETURN      ;  Store DD16 if
              MVI  A,0DDH      ;  odd and return
                              ;  Store EE16
      RETURN  RET             ;  if even and return
```

Example 2.5

Write an 8085 assembly language program to clear 100_{10} consecutive bytes starting at 2050_{16}.

Solution

```
              ORG  2000H
              MVI  A, 64H    ;  Load 'A' with number of
                            ;  bytes to be cleared
              LXI  H,2050H  ;  Load HL with 2050H
      LOOP    MVI  M, 00H   ;  Clear memory location
              INX  H        ;  Increment address
              DCR  A        ;  Decrement loop counter
              JNZ  LOOP     ;  Jump to loop if
                            ;  loop counter zero
              HLT           ;  Stop
```

2.6 Timing Methods

Timing concepts are very important in microprocessor applications. Typically, in sequential process control the microprocessor is required to provide time delays for on-off devices such as pumps or motor-operated valves. DELAY routines are used to provide such time delays.

A delay program typically has an input register that contains the initial count. The register pair D,E is used for this purpose. A typical delay subroutine is given below:

```
      DELAY   DCX  D        ;  Decrement the D,E contents
              MOV  A,D      ;
              ORA  E        ;  Are the contents zero?
              JNZ  DELAY    ;  Jump if not zero
              RET           ;  Return
```

We now calculate the total time required by the DELAY routine using the following data:

Instruction	Number of cycles
CALL	18
DCX D	6
MOV A,D	4
ORA E	4
JNZ	7/10
RET	10

Note that in the above, if the JNZ condition is met ($Z = 0$), ten cycles are required and the program branches back to the DCX D instruction. However, if the JNZ condition is not met ($Z = 1$), the seven cycles are required, and the program executes the next instruction, that is, the RET instruction. Also, note that the CALL instruction is used in the main program written by the user, and the 3-byte instruction CALL DELAY is used.

For each iteration in which the JNZ condition is met ($Z = 0$), the number of cycles is equal to cycles for DCX D + cycles for MOV A,D + cycles for ORA E + cycles for JNZ = $6 + 4 + 4 + 10 = 24$ cycles.

These 24 cycles will be performed $(y - 1)$ times, where y is the initial contents of D,E. For the final iteration in which no jump is performed and the JNZ condition is not satisfied ($Z \neq 0$), the number of cycles is equal to cycles for DCX D + cycles for MOV A,D + cycles for ORA E + cycles for JNZ + cycles for RET = 6 + 4 + 4 + 7 + 10 = 31 cycles. Therefore, the time used, including a CALL instruction, is

$$18 + 31 + 24(y - 1) = 49 + 24(y - 1) \text{ clock cycles}$$

Suppose that in a program a delay time of 1/3 ms is desired. The DELAY routine can be used to accomplish this in the following way. Each cycle of the 8085 clock is 1/3 ms (3 MHz). The number of cycles required in the DELAY routine is

$$\frac{1/3 \text{ ms}}{1/3 \text{ μs}} = \frac{10^{-3}}{10^{-6}} = 1000 \text{ cycles}$$

Therefore, the initial counter value y of the D,E register pair can be calculated:

$$49 + 24 \, (y - 1) = 1000$$
$$24 \, (y - 1) = 951$$
$$y = \frac{951}{24} + 1 \cong 40_{10} = 28_{16}$$

Therefore, in the program the D,E register pair can be loaded with 0028_{16} and the DELAY routine can be called to obtain 1/3 ms of time delay. Table 2.10 shows initial counts for various time delays.

Table 2.10   Time Intervals along with Initial Counts

3-MHz clock milliseconds	Initial count (hexadecimal)
1/3	0028
1	007C
2	00F9
10	04E1
100	30D3

The following program produces a delay of 10 ms:

```
LXI SP, 5000H ;   Set stack pointer
LXI D, 04E1H  ;   Load D,E with initial count value of 04E1
              ;   to provide 10 ms of delay
CALL DELAY    ;   Call DELAY routine
HLT           ;   STOP
```

The delay times can be increased by using a counter. Suppose that a delay of 5 s is desired in a program. From Table 2.10 an initial count of $30D3_{16}$ produces a 100-ms delay. We can use a counter along with the 100-ms delay to obtain the 5-s delay as follows

$$(100 \text{ ms}) \times X = 5 \text{ s}$$

where X is the value of the counter. Then

$$X = \frac{5}{100 \times 10^{-3}} = \frac{5}{10^{-1}} = 50$$

Therefore, a counter of 50_{10} or 32_{16} is required. Now the program for the 5-s delay can be written as follows:

```
            LXI SP, 5000H ;  Set stack pointer
            MVI C, 32H    ;  Do DELAY loop 50₁₀ times by loading C
                          ;  with count 32₁₆
START       LXI D, 30D3H  ;  Load initial count
            CALL DELAY    ;  Call DELAY loop
            DCR C         ;  Decrement C and check if zero: if
                          ;  not, do another delay
            JNZ START     ;  Loop back
            HLT           ;  STOP
```

In the above, since execution times of DCR C and JNZ START are very small compared to 5 s, they are not considered in computing the delay.

2.7 8085 Pins and Signals

The 8085 is housed in a 40-pin dual in-line package (DIP). Figure 2.8 shows the 8085 pins and signals.

The low-order address byte and data lines AD0 to AD7 are multiplexed. These lines are bidirectional. The beginning of an instruction is indicated by the rising edge of the ALE signal. At the falling edge of ALE, the low byte of the address is automatically latched by some of the 8085 support chips such as 8155 and 8355: AD0 to AD7 lines can then be used as data lines. Note that ALE is an input to these support chips. However, if the support chips do not latch AD0 to AD7, then external latches are required to generate eight separate address lines A7 to A0 at the falling edge of ALE.

Pins A8 to A15 are unidirectional and contain the high byte of the address.

Table 2.11 lists the 8085 pins along with a brief description of each.

The $\overline{RD}$ pin signal is output LOW by the 8085 during a memory or I/O READ operation. Similarly, the $\overline{WR}$ pin signal is output LOW during a memory or I/O WRITE.

Next, we explain the purpose of IO/$\overline{M}$, S0, and S1 signals. The IO/$\overline{M}$ signal is output HIGH by the 8085 to indicate execution of an I/O instruction such as IN or OUT. This pin is output LOW during execution of a memory instruction such as LDA 2050H.

The IO/$\overline{M}$, S0, and S1 are output by the 8085 during its internal operations, which can be interpreted as follows:

IO/$\overline{M}$	S1	S0	Operation performed by the 8085
0	0	1	Memory WRITE
0	1	0	Memory READ
1	0	1	I/O WRITE
1	1	0	I/O READ
0	1	1	OP code fetch
1	1	1	Interrupt acknowledge
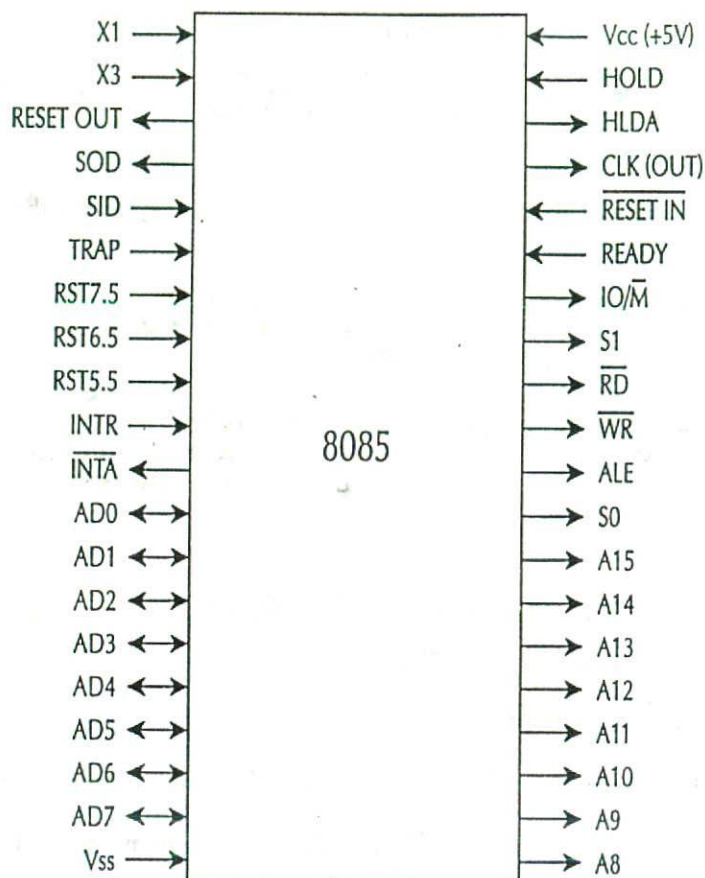
FIGURE 2.8   8085 microprocessor signals and pin assignments.

Table 2.11   8085 Signal Description Summary

Pin name	Description	Type
AD0–AD7	Address/data bus	Bidirectional, tristate
A8–A15	Address bus	Output, tristate
ALE	Address latch enable	Output, tristate
RD	Read control	Output, tristate
WR	Write control	Output, tristate
IO / M	I/O or memory indicator	Output, tristate
S0, S1	Bus state indicators	Output
READY	Wait state request	Input
SID	Serial data input	Input
SOD	Serial data output	Output
HOLD	Hold request	Input
HLDA	Hold acknowledge	Output
INTR	Interrupt request	Input
TRAP	Nonmaskable interrupt request	Input
RST5.5	Hardware vectored	Input
RST6.5	Hardware vectored interrupt request	Input
RST7.5	Hardware vectored	Input
INTA	Interrupt acknowledge	Output
RESET IN	System reset	Input
RESET OUT	Peripherals reset	Output
X1, X2	Crystal or RC connection	Input
CLK (OUT)	Clock signal	Output
Vcc, Vss	Power, ground	
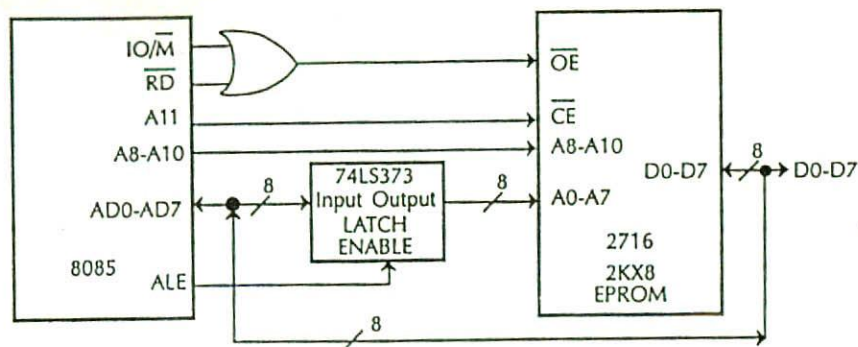
(a) 8155 - 2716 interface using internal latches.
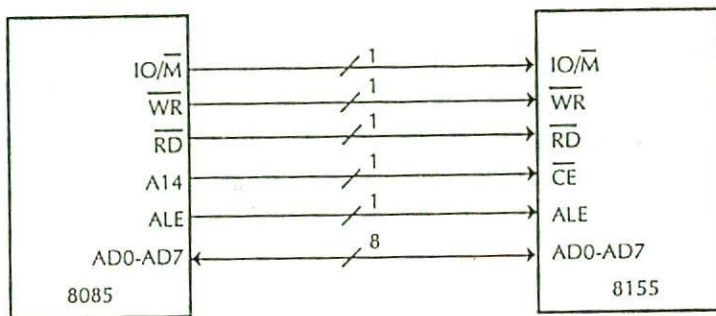
(b) 8085 - 8155 interface using ALE and AD0-AD7

FIGURE 2.9   8085's interface to external device using ALE and the multiplexed AD0 to AD7 pins.

Figure 2.9 illustrates the utilization of ALE and AD0 to AD7 signals for interfacing an EPROM and a RAM.

The 2716 is a 2K × 8 EPROM with separate address and data lines without any built-in latches. This means that a separate latch such as the 74LS373 must be used to isolate the 8085 low byte address and D0-D7 data lines at the falling edge of ALE (Figure 2.9a).

The 8155 contains 256-byte static RAM, three user ports, and a 14-bit timer. The 8155 is designed for 8085 in the sense that it has built-in latches with ALE as input along with multiplexed address (low byte) and data lines, AD0 to AD7. Therefore, as shown in Figure 2.9b, external latches are not required.

The READY input can be used by the slower external devices for obtaining extra time in order to communicate with the 8085. The READY signal (when LOW) can be utilized to provide wait-state clock periods in the 8085 machine. If READY is HIGH during a read or write cycle, it indicates that the memory or peripheral is ready to send or receive data. If not used, it must be tied high.

The Serial Input Data (SID) and Serial Output Data (SOD) lines are associated with the 8085 serial I/O transfer. The SOD line can be used to output the most significant bit of the accumulator. The SID signal can be input into the most significant bit of the accumulator.

The HOLD and HLDA signals are used for the Direct Memory Access (DMA) type of data transfer. The external devices place a HIGH on HOLD line in order to take control of the system bus. The HOLD function is acknowledged by the 8085 by placing a HIGH output on the HLDA pin and driving the tristate outputs high impedance.

The signals on the TRAP, RST7.5, RST6.5, RST5.5, INTR, and INTA are related to the 8085 interrupt signals. TRAP is a nonmaskable interrupt; that is, it cannot be enabled or disabled
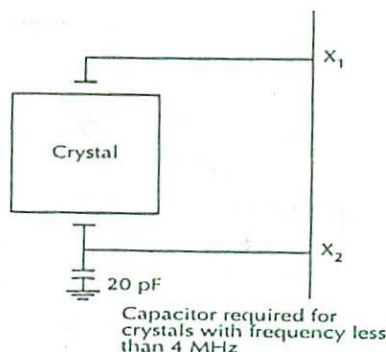
FIGURE 2.10   Crystal connection to X1 and X2 pins.

by an instruction. The TRAP has the highest priority. RST7.5, RST6.5, and RST5.5 are maskable interrupts used by the devices whose vector addresses are generated automatically. INTA is an interrupt acknowledge signal which is pulsed LOW by the 8085 in response to the interrupt INTR request. In order to service INTR, one of the eight OP codes (RST0 to RST7) has to be provided on the 8085 AD0-AD7 bus by external logic. The 8085 then executes this instruction and vectors to the appropriate address to service the interrupt.

All unused control pins such as interrupts and HOLD must be disabled by grounding them. (READY must be tied high).

The 8085 has the clock generation circuit on the chip and, therefore, no external oscillators need to be designed. The 8085A can operate with a maximum clock frequency of 3.03 MHz and the 8085A-2 can be driven with a maximum of 5 MHz clock. The 8085 clock frequency can be generated by a crystal, an LC tuned circuit, or an external clock circuit. The frequency at X1X2 is divided by 2 internally. This means that in order to obtain 3.03 MHz, a clock source of 6.06 MHz must be connected to X1X2. For crystals of less than 4 MHz, a capacitor of 20 pF should be connected between X2 and a ground to ensure the starting up of the crystal at the right frequency (Figure 2.10).

There is a TTL signal which is output on pin 37, called the CLK (OUT) signal. This signal can be used by other external microprocessors or support chips.

The RESET IN signal, when pulsed LOW then high, causes the 8085 to execute the first instruction at the 0000_{16} location. In addition, the 8085 resets instruction register, interrupt mask (RST5.5, RST6.5, and RST7.5) bits, and other registers. The RESET IN must be held LOW for at least three clock periods. A typical 8085 reset circuit is shown in Figure 2.11. In this circuit, when the switch is activated, RESET IN is driven to LOW with a large time constant providing adequate time to reset the system.

The 8085 requires a minimum operating voltage of 4.75 V. Upon applying power, the 8085A attains this voltage after 500 µs. The reset circuit of Figure 2.11 resets the 8085 upon activation
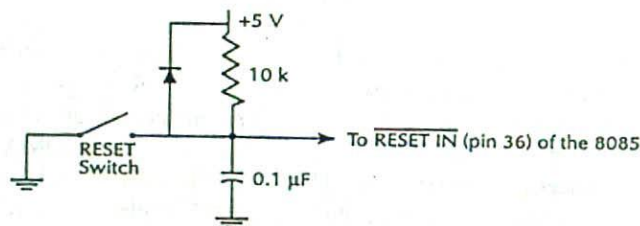
FIGURE 2.11   8085 reset circuit.

FIGURE 2.12   8085 machine cycles.

of the switch. The voltage across the 0.1-μF capacitor is zero on power-up. The capacitor then charges Vcc after a definite time determined by the time constant RC. The chosen values of RC in the figure will drive the RESET IN pin to low for at least three clock periods. In this case, after activating the switch, RESET IN will be low (assuming capacitance charge time is equal to the discharge time) for 10K*0.1 μF = 1 ms, which is greater than three clock periods (3*1/ 3 μs = 1 μs) of the 3-MHz 8085A. During normal operation of the 8085, activation of the switch will short the capacitor to ground and will discharge it. When the switch is opened, the capacitor charges and the RESET IN pin becomes HIGH. Upon hardware reset, the 8085 clears PC, IR, HALT flip-flop, and some other registers; the 8085 registers PSW, A, B, C, D, E, H, and L are unaffected. Upon activation of the RESET IN to low, the 8085 outputs HIGH at the RESET OUT pin which can be used to reset the memory and I/O chips connected to the 8085. Note that since hardware reset initializes PC to 0, the 8085 fetches the first instruction for address 0000_{16} after reset.

2.8 8085 Instruction Timing and Execution

An 8085's instruction execution consists of a number of machine cycles (MCs). These cycles vary from one to five (M1 to M5) depending on the instruction. Each machine cycle contains a number of 320-ns clock periods. The first machine cycle will be executed in either four or six clock periods, and the machine cycles that follow will have three clock periods. This is shown in Figure 2.12.

The shaded MCs indicate that these machine cycles are required by certain instructions. Similarly, the shaded clock periods (T5 and T6) mean that they are needed in M1 by some instructions.

The clock periods within a machine cycle can be illustrated as shown in Figure 2.13. Note that the beginning of a new machine cycle is indicated on the 8085 by outputting the Address Latch Enable (ALE) signal HIGH. During this time, lines AD0 to AD7 are used for placing the low byte of the address.

When the ALE signal goes LOW, the low byte of the address is latched so that the AD0 to AD7 lines can be used for transferring data.

We now discuss the timing diagrams for instruction fetch, READ, and WRITE.

FIGURE 2.13   Clock period within a machine cycle.

2.8.1 Basic System Timing

Figure 2.14 shows the 8085 basic system timing. An instruction execution includes two operations: OP code fetch and execution.

The OP code fetch cycle requires either four (for one-byte instructions such as MOV A,B) or six cycles (for 3 byte instructions such as LDA 2030H). The machine cycles that follow will need three clock periods.

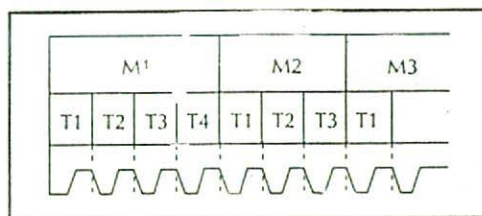The purpose of an instruction fetch is to read the contents of a memory location containing an instruction addressed by the program counter and to place it in the instruction register. The 8085 instruction fetch timing diagram shown in Figure 2.14 can be explained in the following way:

1. The 8085 puts a LOW on the $IO/\overline{M}$ line of the system bus, indicating a memory operation.
2. The 8085 sets S0 = 1 and S1 = 1 on the system bus, indicating the memory fetch operation.
3. The 8085 places the program counter high byte on the A8 to A15 lines and the program counter low byte on the AD0 to AD7 lines of the system bus. The 8085 also sets the ALE signal to HIGH. As soon as the ALE signal goes to LOW, the program counter low byte on the AD0 to AD7 is latched automatically by some 8085 support chips such as 8155 (if 8085 support chips are not used, these lines must be latched using external latches), since these lines will be used as data lines for reading the OP code.
4. At the beginning of T2 in M1, the 8085 puts the $\overline{RD}$ line to LOW indicating a READ operation. After some time, the 8085 loads the OP code (the contents of the memory location addressed by the program counter) into the instruction register.
5. During the T4 clock period in M1, the 8085 decodes the instruction.

The Machine Cycle M2 of Figure 2.14 shows a memory (or I/O) READ operation as seen by the external logic, and the status of the S0 and S1 signals indicates whether the operation is instruction fetch or memory READ; for example, S1 = 1, S0 = 1 during instruction fetch and S1 = 1, S0 = 0 during memory READ provided $IO/\overline{M}$ = 0.

The purpose of the memory READ is to read the contents of a memory location addressed by a register pair, such as the H,L pair, or a memory location specified with the instruction and the data placed in a microprocessor register such as the accumulator. In contrast, the purpose of the memory fetch is to read the contents of a memory location addressed by PC into IR. The machine cycle M3 of Figure 2.14 indicates a memory (or I/O) write operation. In this case, S1 = 0 and S0 = 1 indicate a memory write operation when $IO/\overline{M}$ = 0 and an I/O write operation when $IO/\overline{M}$ = 1.

2.8.2 8085 Memory READ ($IO/\overline{M}$ = 0, $\overline{RD}$ = 0) and I/O READ ($IO/\overline{M}$ = 1, $\overline{RD}$ = 0)

Figure 2.15a shows an 8085A clock timing diagram. The machine cycle of M2 of Figure 2.14 shows a memory READ timing diagram.

The purpose of the memory READ is to read the contents of a memory location addressed by a register pair, such as HL. Let us explain the 8085 memory READ timing diagram of Figure 2.15b along with the READ timing signals of Figure 2.14:

1. The 8085 uses machine cycle M1 to fetch and decode the instruction. It then performs the memory READ operation in M2.
2. The 8085 continues to maintain $IO/\overline{M}$ at LOW in M2 indicating a memory READ operation (or $IO/\overline{M}$ = 1 for I/O READ).
3. The 8085 puts S1 = 1, S0 = 0, indicating a READ operation.
4. The 8085 places the contents of the high byte of the memory address register, such as the contents of the H register, on lines A8 to A15.
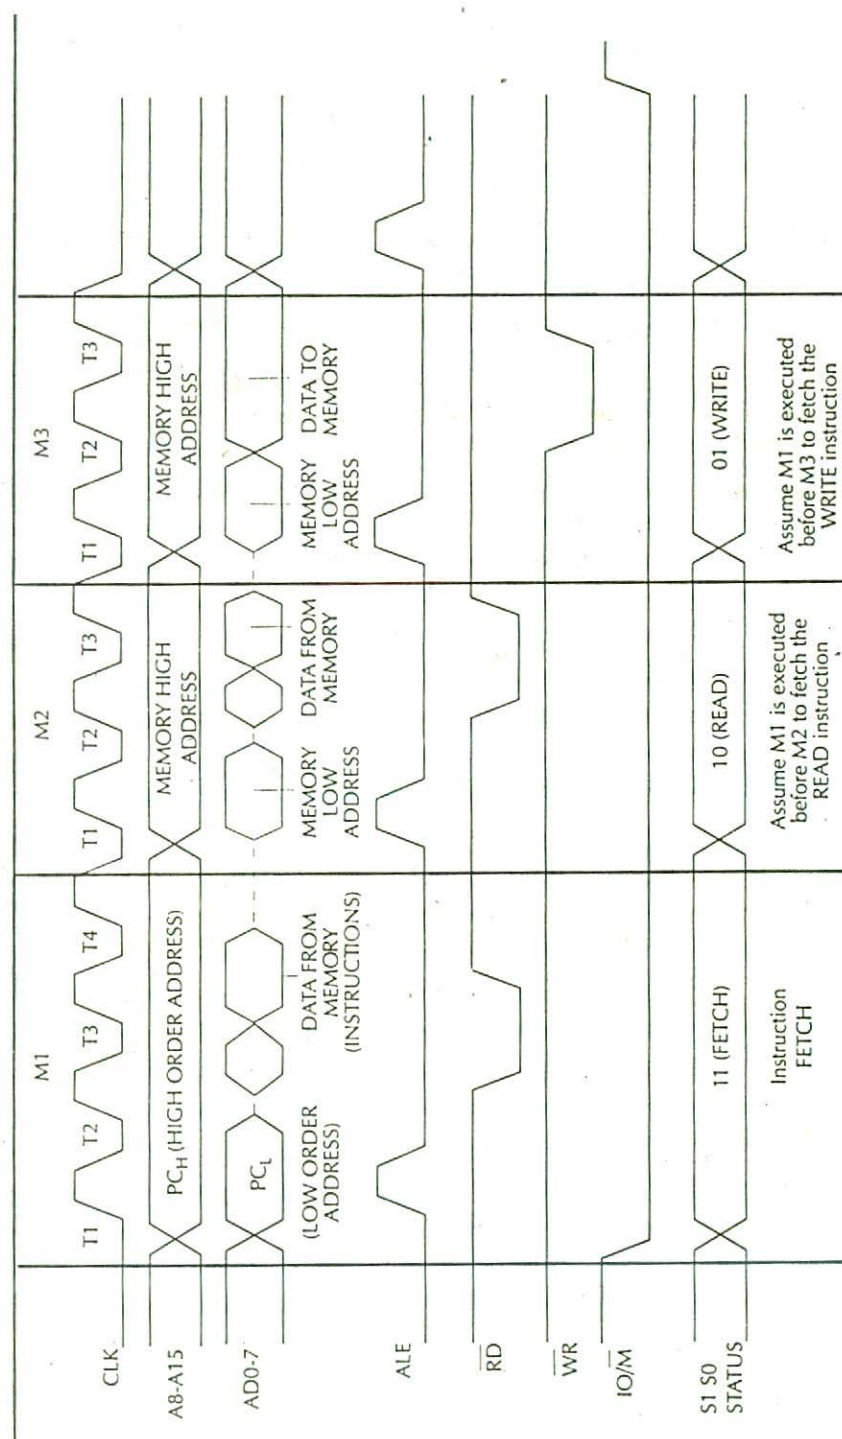
FIGURE 2.14   8085 Basic system timing.

5. The 8085 places the contents of the low byte of the memory address register, such as the contents of the L register, in lines AD0 to AD7.
6. The 8085 sets ALE to high, indicating the beginning of M2. As soon as ALE goes to low, the memory or support chip must latch the low byte of the address lines, since the same lines are going to be used as data lines.
7. The 8085 puts the $\overline{RD}$ signal to LOW, indicating a READ operation.
8. The 8085 gets the data from the memory location addressed by the memory address register, such as the H,L pair, and places the data into a register such as the accumulator. In case of I/O, the 8085 inputs data from the I/O port into the accumulator.

2.8.3 8085 Memory WRITE (IO/$\overline{M}$ = 0, $\overline{WR}$ = 0) and I/O WRITE (IO/$\overline{M}$ = 1, $\overline{WR}$ = 0)

The machine cycle M3 of Figure 2.14 shows a memory WRITE timing diagram. As seen by the external logic, the signals S1 = 0, S0 = 1, and $\overline{WR}$ = 0 indicate a memory WRITE operation.

The purpose of a memory WRITE is to store the contents of the 8085 register, such as the accumulator, into a memory location addressed by a pair, such as H,L.

The WRITE timing diagram of Figure 2.14 can be explained as follows:

1. The 8085 uses machine cycle M1 to fetch and decode the instruction. It then executes the memory WRITE instruction in M3.
2. The 8085 continues to maintain IO/$\overline{M}$ at LOW, indicating a memory operation (or IO/$\overline{M}$ = 1 for I/O WRITE).
3. The 8085 puts S1 = 0, S0 = 1, indicating a WRITE operation.
4. The 8085 places the Memory Address Register high byte, such as the contents of the H register, on lines A8 to A15.
5. The 8085 places the Memory Address Register low byte, such as the contents of L register, on lines AD0 to AD7.
6. The 8085 sets ALE to HIGH, indicating the beginning of M3. As soon as ALE goes to LOW, the memory or support chip must latch the low byte of the address lines, since the same lines are going to be used as data lines.
7. The 8085 puts the $\overline{WR}$ signal to LOW, indicating a WRITE operation.
8. It also places the contents of the register, say, accumulator, on data lines AD0 to AD7.
9. The external logic gets data from the lines AD0 to AD7 and stores the data in the memory location addressed by the Memory Address Register, such as the H,L pair. In case of I/O, the 8085 outputs [A] to an I/O port.

Figures 2.15a through c show the 8085A clock and read and write timing diagrams.

2.9 8085 Input/Output (I/O)

The 8085 I/O transfer techniques are discussed. The 8355/8755 and 8155/8156 I/O ports and 8085 SID and SOD lines are also included.

2.9.1 8085 Programmed I/O

There are two I/O instructions in the 8085, namely, IN and OUT. These instructions are 2 bytes long. The first byte defines the OP code of the instruction and the second byte specifies the I/O port number. Execution of the IN PORT instruction causes the 8085 to receive one byte of data into the accumulator from a specified I/O port. On the other hand, the OUT PORT instruction, when executed, causes the 8085 to send one byte of data from the accumulator into a specified I/O port.
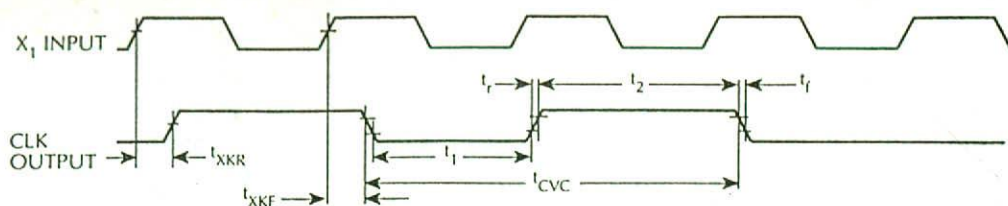
FIGURE 2.15a   8085A clock.
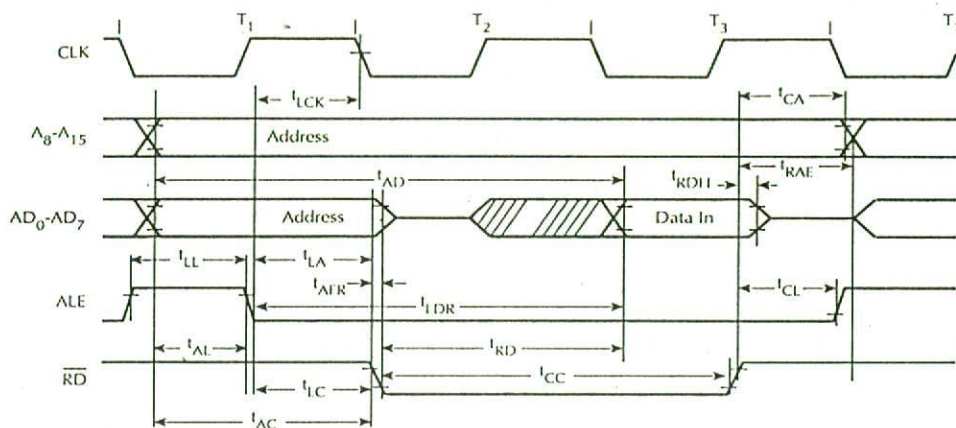
FIGURE 2.15b   8085 Read Timing diagram.
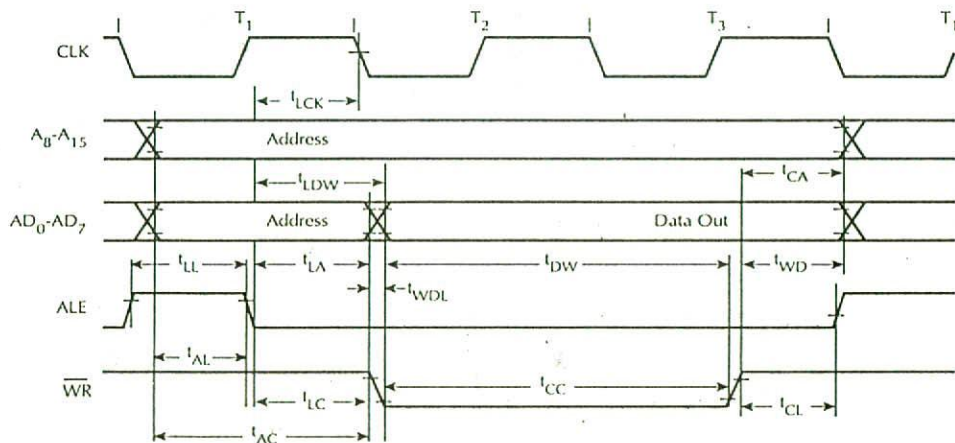
FIGURE 2.15c   8085 Write Timing diagram.

The 8085 can access I/O ports using either standard I/O or memory-mapped I/O.

In standard I/O, the 8085 inputs or outputs data using IN or OUT instructions.

In memory-mapped I/O, the 8085 maps I/O ports as memory addresses. Hence, LDA addr or STA addr instructions are used to input or output data to or from the 8085. The 8085's programmed I/O capabilities are obtained via the support chips, namely, 8355/8755 and 8155/8156. The 8355/8755 contains a 2K-byte ROM/EPROM and two 8-bit I/O ports (ports A and B).

The 8155/8156 contains 256-byte RAM, two 8-bit and one 6-bit I/O ports, and a 14-bit programmable timer. The only difference between the 8155 and 8156 is that chip enable is LOW on the 8155 and HIGH on the 8156.

2.9.1.a 8355/8755 I/O Ports

Two 8-bit ports are included in the 8355/8755. These are ports A and B. Another 8-bit port, called the data direction register, is associated with each one of these ports. These registers (DDRA and DDRB) can be used to configure each bit in ports A or B as either input or output. For example, a "0" written into a bit position of the data direction register sets up the corresponding bit in the I/O port as input. On the other hand, a "1" written in a particular bit position in the data direction register sets up the corresponding bit in the I/O port as output. For example, consider the following instruction sequence:

```
MVI   A,  05H
OUT   DDRA
```

The above instruction sequence assumes DDRA as the data direction register for port A. The bits of port A are configured as follows:

	7	6	5	4	3	2	1	0
DDRA	0	0	0	0	0	1	0	1

	7	6	5	4	3	2	1	0
Port A	IN	IN	IN	IN	IN	OUT	IN	OUT

The 8355/8755 uses the $IO/\overline{M}$ pin on the chip in order to distinguish between standard and memory-mapped I/O. This pin is controlled by the 8085 as shown in Figure 2.16.

The 8085 outputs a HIGH on the IO/M pin when it executes either an IN or OUT instruction. This means that $IO/\overline{M}$ in the 8355/8755 becomes HIGH during execution of IN or OUT. This, in turn, tells the 8355/8755 to decode the AD1 and AD0 lines in order to obtain the 8-bit address of various ports in the chip as follows:

AD1	AD0	
0	0	Port A
0	1	Port B
1	0	Data Direction Register A
1	1	Data Direction Register B

The other 6 bits of each 8-bit port address are don't care conditions. This means that these bits can be either one or zero. The 8085/8355/8755 standard I/O is illustrated in Figure 2.17.

Since the 8355/8755 is provided with 2K bytes of memory, 11 address lines A0 to A10 are required for memory addressing. Since the 8085 has 16 address pins A0 to A15, A11 to A15 will not be used for memory addressing. Note that the 8355/8755 includes two chip enables $\overline{CE}$ and CE. In Figure 2.17 these two chip enables are connected to Vcc and A11, respectively. It should be pointed out that the 8085 duplicates low and high bytes of the 16-bit address lines with the port address when it executes an IN or OUT instruction. This means that if the 8085 executes IN 01 instruction, it puts 0101_{16} on the 16 address lines. Note that in Figure 2.17, A11 = 0 for both memory and port addressing. This is because A11 = 0 enables this chip. When the 8085 executes an LDA addr or STA addr instruction, $IO/\overline{M}$ becomes LOW. This tells the 8355/8755 to interpret A0 to A10 as memory addresses. On the other hand, when the 8085 executes an IN PORT or OUT PORT instruction, the 8085 drives $IO/\overline{M}$ to HIGH. This tells the 8355/8755 to decode AD1 and AD0 for I/O port addresses. The port addresses are as follows:
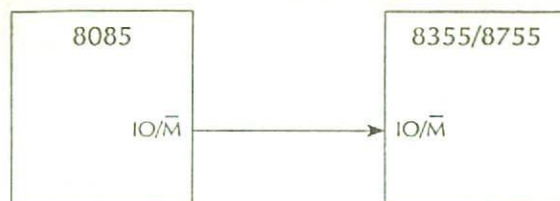
FIGURE 2.16   Interfacing the 8085 with 8355/8755 via the $IO/\overline{M}$ pin.
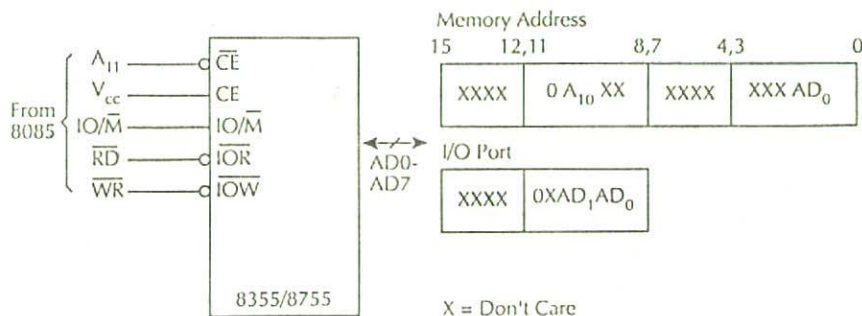
FIGURE 2.17   8355/8755 standard I/O.

| | A15 | A14 | A13 | A12 | A11 | A10 | A9 | A8 | |
	AD7	AD6	AD5	AD4	AD3	AD2	AD1	AD0	Address
Port A =	X	X	X	X	X	0	X	0	$= 00_{16}$
Port B =	X	X	X	X	0	X	0	1	$= 01_{16}$
DDRA =	X	X	X	X	0	X	1	0	$= 02_{16}$
DDRB =	X	X	X	X	0	X	1	1	$= 03_{16}$

X is don't care. Assume X is zero in the above.

Let us now discuss 8355/8755 memory-mapped I/O. Figure 2.18 provides such an example. In Figure 2.18, A11 must be zero for selecting the 8355/8755 and A15 is connected to $IO/\overline{M}$ of the 8355/8755. When A15 = 1, $IO/\overline{M}$ becomes HIGH. This tells the 8355/8755 chip to decode AD1 and AD0 for obtaining I/O port addresses. For example, if we assume all don't cares in the I/O port address are 1, then the I/O port addresses will be mapped into memory locations as follows:

Port name	16-bit memory address
I/O port A	$F7FC_{16}$
I/O port B	$F7FD_{16}$
DDRA	$F7FE_{16}$
DDRB	$F7FF_{16}$

Note that in Figure 2.18 memory addresses are mapped as 7000_{16} through $77FF_{16}$ assuming all don't cares to be ones.

Note that the above port addresses may not physically exist in memory. However, input or output operations with these ports can be accomplished by generating the necessary signals by executing LDA or STA instructions with the above addresses. For example, outputting to DDRA or DDRB can be accomplished via storing to locations $F7FE_{16}$ or $F7FF_{16}$, respectively. The instructions STA F7FEH or STA F7FFH will generate all of the required signals for OUT
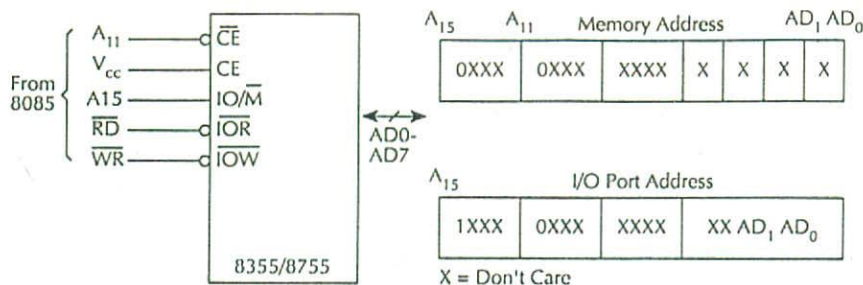
FIGURE 2.18   8355/8755 memory-mapped I/O.

DDRA or OUT DDRB, respectively. For example, upon execution of the STA F7FEH, the 8085 sends a LOW to the $\overline{WR}$ pin and places F7FEH on the address bus. This will make A15 = 1, A11 = 0, AD1 = 1, AD0 = 0, and thus will make the $IO/\overline{M}$ = 1, $\overline{CE}$ = 0, $\overline{IOW}$ = 0 on the 8355/8755 in Figure 2.18.

When ALE goes to LOW, the 8085 places the contents of the accumulator on the AD7 to AD0 pins. The 8355/8755 then takes this data and writes into DDRA. Therefore, STA F7FEH is equivalent to OUT DDRA instruction.

2.9.1.b 8155/8156 I/O Ports

The 8155 or 8156 includes 256 bytes of static RAM and three parallel I/O ports. These ports are port A (8-bit), port B (8-bit), and port C (6-bit). By parallel it is meant that all bits of the port are configured as either all input or all output. Bit-by-bit configuration like the 8355/8755 is not permitted. The only difference between the 8155 and 8156 is that the 8155 has LOW chip enable ($\overline{CE}$), while the 8156 includes a HIGH chip enable (CE). The 8155/8156 ports are configured by another port called the command status register (CSR). When data are output to the CSR via the accumulator, each bit is interpreted as a command bit to set up ports and control timer as shown in Figure 2.19. Port C can be used as a 6-bit parallel port or as a control port to support data transfer between the 8085 and an external device via ports A and B using handshaking. Note that handshaking means data transfer via exchange of control signals. Two bits (bits 2 and 3) are required in CSR to configure port C. Note that port A interrupt and port B interrupt are associated with handshaking and are different from the 8085 interrupts. For example, port A interrupt is HIGH when data are ready to be transferred using handshaking signals such as port A buffer full and port A strobe. The port A interrupt (PC0 in ALT3) can be connected to an 8085 interrupt pin and data can be transferred to or from the 8085 via port A by executing appropriate instructions in the interrupt service routine.

When the 8085 reads the CSR, it accesses the status register and information such as status of handshaking signals and timer interrupt is obtained.

Three bits are used to decode the 8155 six ports (CSR, port A, port B, port C, timer high port, timer low port) as follows:

AD2	AD1	AD0	Port selected
0	0	0	CSR
0	0	1	Port A
0	1	0	Port B
0	1	1	Port C
1	0	0	Timer-low port
1	0	1	Timer-high port
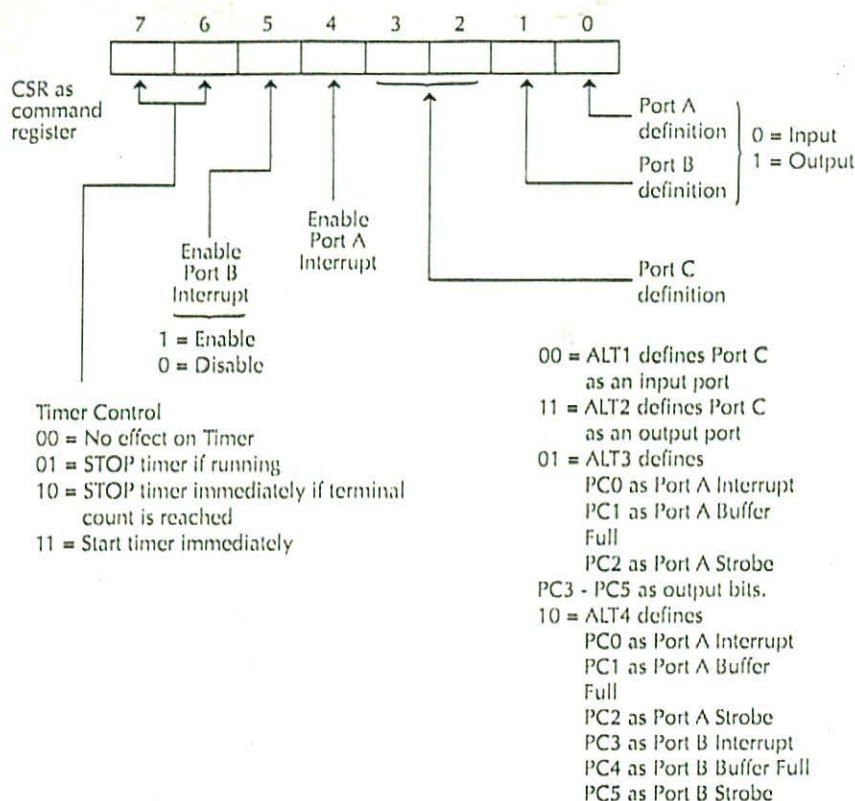
FIGURE 2.19 CSR format as command register.

A typical interface between the 8085 and 8155 is shown in Figure 2.20. Consider Figure 2.20a. Since the 8085 duplicates low byte address bus with the high byte address (i.e., AD7 to AD0 same as A15 to A8 for 8085 standard I/O), the address pins AD2 to AD0 will be the same as A10 to A8. This means that the pins A10 to A8 must not be used as chip enables since they will be used for decoding of port addresses. Therefore, A14 is used as chip enable in the figure. The unused address lines A11 to A13 and A15 are don't cares and are assumed to be zero in the following. Therefore, the port addresses are

A15 AD7	A14 AD6	A13 AD5	A12 AD4	A11 AD3	A10 AD2	A9 AD1	A8 AD0	Address
0	0	0	0	0	0	0	0	= 00H = CSR
0	0	0	0	0	0	0	1	= 01H = Port A
0	0	0	0	0	0	1	0	= 02H = Port B
0	0	0	0	0	0	1	1	= 03H = Port C

For memory-mapped I/O, consider Figure 2.20b. In this case, the 8085 low byte and high byte address bus are not duplicated. The ports will have 16-bit addresses as follows. Assume the unused address pins A8 to AD3 to be zeros.
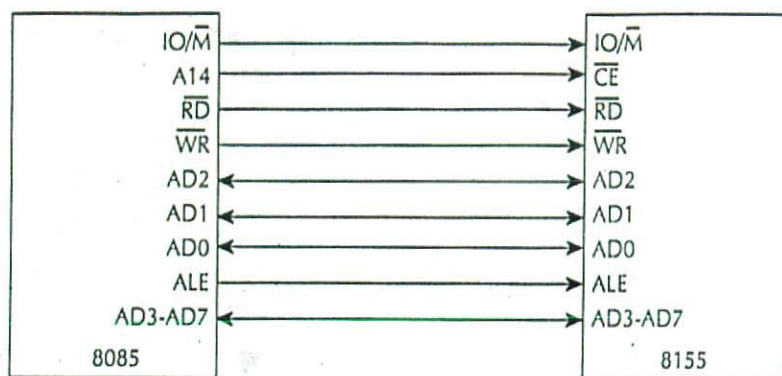
(a) 8085 – 8155 Interface using standard I/O
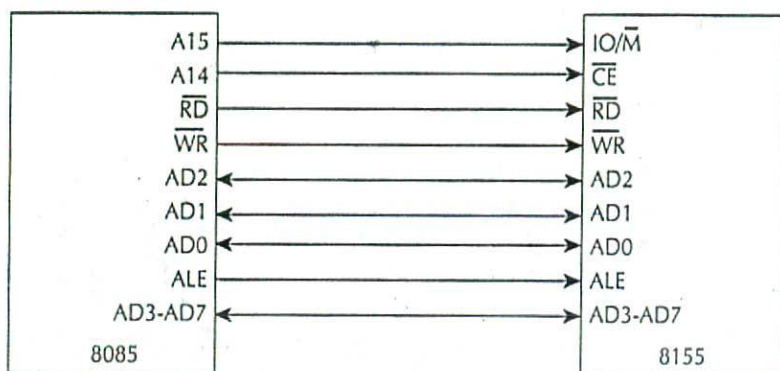
(b) 8085 – 8155 Interface using memory-mapped I/O

FIGURE 2.20   8085-8155 interface for I/O ports.

A15	A14	A13	A12	A11	A10	A9	A8	AD7	AD6	AD5	AD4	AD3	AD2	AD1	AD0	Port address
1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	$= 8000_{16}$ CSR
1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	$= 8001_{16}$ (Port A)
1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	$= 8002_{16}$ (Port B)
1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	$= 8003_{16}$ (Port C)

Like the 8355/8755 memory-mapped I/O, the above port addresses may not physically exist. However, read or write operations with them will generate the necessary signals for input or output transfer with the ports. Assuming all don't cares to be zeros, the memory map of either configuration of Figure 2.20a or b includes addresses 0000H through 00FFH.

Example 2.6

An 8085-8355-based microcomputer is required to drive an LED connected to bit 0 of port A based on the input conditions set by a switch on bit 1 of port A. The input/output conditions are as follows: if the input to bit 1 of port A is HIGH then the LED will be turned ON; otherwise

the LED will be turned OFF. Assume that a HIGH will turn the LED ON and a LOW will turn it OFF. Write an 8085 assembly language program starting at 5000H.

Solution

```
        ORG 5000H
PORT A  EQU 00H
DDRA    EQU 02H
        MVI A, 01H   ;  Configure Port A
        OUT DDRA
START   IN PORT A    ;  Input Port A
        RAR          ;  Rotate switch to LED position
        OUT PORT A   ;  Output to LED
        JMP START    ;  Endless loop
```

Example 2.7

An 8085-8155-based microcomputer is required to drive an LED connected to bit 0 of port A based on two switch inputs connected to bits 6 and 7 of port A. If both switches are either HIGH or LOW, turn the LED on; otherwise turn it OFF. Assume that a HIGH will turn the LED ON and a LOW will turn it OFF. Use port addresses of CSR and port A as 20_{16} and 21_{16}, respectively.

Write an 8085 assembly language program to accomplish this starting at address 3000H.

Solution

```
        ORG   3000H
CSR     EQU   20H
PORT A  EQU   21H
START   MVI   A, 00H  ;  Configure Port A
        OUT   CSR     ;  as input
        IN    PORT A  ;  Input Port A
        ANI   0C0H    ;  Retain bits 6 and 7
                      ;  turn LED ON
        MVI   A, 01H  ;  Otherwise, configure
        OUT   CSR     ;  Port A as output
        JPE   LEDON
        MVI   A, 00H  ;  and turn LED OFF.
LEDON   OUT   PORT A  ;
        JMP   START   ;  Jump to START
```

Example 2.8

Write an 8085 assembly language program starting at address 5000H to turn on an LED connected to bit 4 of the 8155 I/O port B. Use address of port B as 22_{16}.

Solution

```
        ORG   5000H
Port B  EQU   22H
        MVI   A,02H   ;  Configure
        OUT   CSR     ;  Port B as output
        MVI   A,10H   ;  Output HIGH
        OUT   PORT B  ;  to LED
        HLT
```

Example 2.9

An 8085-8355-based microcomputer is required to drive a common anode seven-segment display connected to port A as follows:
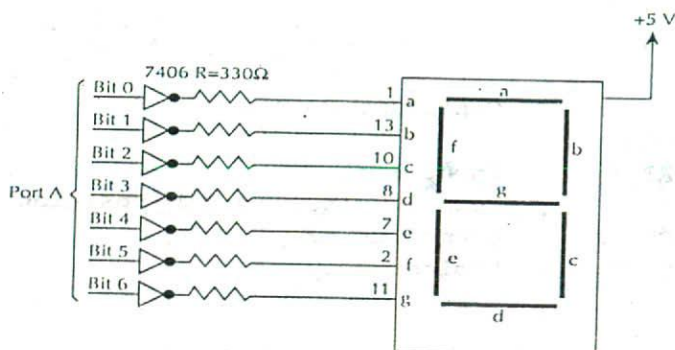
Write an 8085 assembly language program starting at address at 5000H to display a single hexidecimal digit (0 to F) from location 4000_{16}. Use a look-up table. Note that the 7406 shown in the figure contains six inverting buffers on one 7406. Therefore, two 7406 or one 7406 and one transistor are required for the seven segments. Use port addresses of DDRA and port A as 42H and 40H, respectively.

Solution

The decode table can be obtained as follows:

Hex digit	Bits								Decode byte
	7	6	5	4	3	2	1	0	
		g	f	e	d	c	b	a	
0	0	0	1	1	1	1	1	1	3F
1	0	0	0	0	0	1	1	0	06
2	0	1	0	1	1	0	1	1	5B
3	0	1	0	0	1	1	1	1	4F
4	0	1	1	0	0	1	1	0	66
5	0	1	1	0	1	1	0	1	6D
6	0	1	1	1	1	1	0	1	7D
7	0	0	0	0	0	1	1	1	07
8	0	1	1	1	1	1	1	1	7F
9	0	1	1	0	0	1	1	1	67
A	0	1	1	1	0	1	1	1	77
B	0	1	1	1	1	1	0	0	7C
C	0	0	1	1	1	0	0	1	39
D	0	1	0	1	1	1	1	0	5E
E	0	1	1	1	1	0	0	1	79
F	0	1	1	1	0	0	0	1	71

```
        ORG      5000H
DDRA    EQU      42H
Port A  EQU      40H
MVI     A,7FH
OUT     DDRA     ; Port A
LXI     H,4015H  ; Load HL with
                 ; starting address of table
LXI     D,0000H  ; Load 0000H to DE
```

```
        LDA       DIGIT    ; Load digit to be
                           ; Displayed into A
        MOV       E,A      ; Move digit to E
        DAD       D        ; Determine digit address
        MOV       A,M      ; Load decode byte to A
        OUT       Port A   ; Outport decode byte
        HLT                ; to display
        ORG       4015H
TABLE   DB        3FH, 06H, 5BH, 4FH
        DB        66H, 6DH, 7DH, 07H
        DB        7FH, 67H, 77H, 7CH
        DB        39H, 5EH, 79H, 71H
        ORG       4000H
DIGIT   DB        DATA
```

2.9.2 8085 Interrupt System

The 8085 chip has five interrupt pins, namely, TRAP, RST7.5, RST6.5, RST5.5, and INTR. If the signals on these interrupt pins go to HIGH simultaneously, then TRAP will be serviced first (i.e., highest priority) followed by RST7.5, RST6.5, RST5.5, and INTR. Note that once an interrupt is serviced, all the interrupts except TRAP are disabled. They can also be enabled or disabled simultaneously by executing the EI or DI instruction respectively. The 8085 interrupts are

1. TRAP — TRAP is a nonmaskable interrupt. That is, it cannot be enabled or disabled by an instruction. In order for the 8085 to service this interrupt, the signal on the TRAP pin must have a sustained HIGH level with a low to high transition. If this condition occurs, then the 8085 completes execution of the current instruction, pushes the program counter onto the stack, and branches to location 0024_{16} (interrupt address vector for the TRAP). Note that the TRAP interrupt is cleared by the falling edge of the signal on the pin.

2. RST7.5 — RST7.5 is a maskable interrupt. This means that it can be enabled or disabled using the SIM or EI/DI instruction. The 8085 responds to the RST7.5 interrupt when the signal on the RST7.5 pin has a low to high transition. In order to service RST7.5, the 8085 completes execution of the current instruction, pushes the program counter onto the stack, and branches to $003C_{16}$. The 8085 remembers the RST7.5 interrupt by setting an internal D flip-flop by the leading edge.

3. RST6.5 — RST6.5 is a maskable interrupt. It can be enabled or disabled using the SIM or EI/DI instruction. RST6.5 is HIGH level sensitive. In order to service this interrupt, the 8085 completes execution of the current instruction, saves the program counter onto the stack, and branches to location 0034_{16}.

4. RST5.5 — RST5.5 is a maskable interrupt. It can be enabled or disabled by the SIM or EI/DI instruction. RST5.5 is HIGH level sensitive. In order to service this interrupt, the 8085 completes execution of the current instruction, saves the program counter onto the stack, and branches to $002C_{16}$.

5. INTR — INTR is a maskable interrupt. It can be enabled or disabled by EI or DI instruction. This is also called the handshake interrupt. INTR is HIGH level sensitive. When no other interrupts are active and the signal on the INTR pin is HIGH, the 8085 completes execution of the current instruction, and generates an interrupt acknowledge, INTA, LOW pulse on the control bus. The 8085 then expects either a 1-byte CALL (RST0 through RST7) or a 3-byte CALL on the data lines. This instruction must be provided by external hardware. In other words, the INTA can be used to enable a tristate

buffer. The output of this buffer can be connected to the 8085 data lines. The buffer can be designed to provide the appropriate op code on the data lines. Note that the occurrence of $\overline{INTA}$ turns off the 8085 interrupt system in order to avoid multiple interrupts from a single device. Also note that there are eight RST instructions (RST0 through RST7). Each of these RST instructions has a vector address. These were shown in Table 2.9.

In response to a HIGH on the INTR, the 8085 proceeds with the sequence of events described below. If INTR is the only interrupt and if the 8085 system interrupt is enabled by executing the EI instruction, the 8085 will turn off the system interrupt and then make the $\overline{INTA}$ LOW for about two cycles. This $\overline{INTA}$ signal can be used to enable an external hardware to provide an op code on the data bus. The 8085 can then read this op code. Typically, the 1-byte RST or 3-byte CALL instruction can be used as the op code. If the 3-byte CALL is used, then the 8085 will generate two additional $\overline{INTA}$ cycles in order to fetch all 3 bytes of the instruction. However, on the other hand, if RST is used, then no additional $\overline{INTA}$ is required. The call op code is normally placed on the data bus by the 8259 programmable interrupt controller. At this point, only the op code for the CALL (CD16) is fetched by the 8085. The 8085 executes this instruction and determines that it needs two more bytes (the address portion of the 3-byte instruction). The 8085 then generates a second $\overline{INTA}$ cycle followed by a third $\overline{INTA}$ cycle in order to fetch the address portion of the CALL instruction from the 8259. The 8085 executes the CALL instruction and branches to the interrupt service routine located at an address specified in the CALL instruction. Note that the recognition of any maskable interrupt (RST7.5, RST6.5, RST5.5, and INTR) disables all maskable interrupts to avoid multiple interrupts from the same device.

Therefore, in order that the 8085 can accept another interrupt, the last two instructions of the interrupt service routine will be EI followed by RET.

One can produce a single RST instruction, say RST7 (op code FFH in HEX), using 74LS244. The inputs I0 to I7 of the 74LS244 (Figure 2.21) are connected to HIGH and its enable line ($\overline{OE}$) is tied to an $\overline{INTA}$. In response to INTA LOW, the 74LS244 places FF in hex (RST7) on the data bus. Figure 2.21 shows a typical circuit. Figure 2.22 provides a schematic for eight priority interrupts.
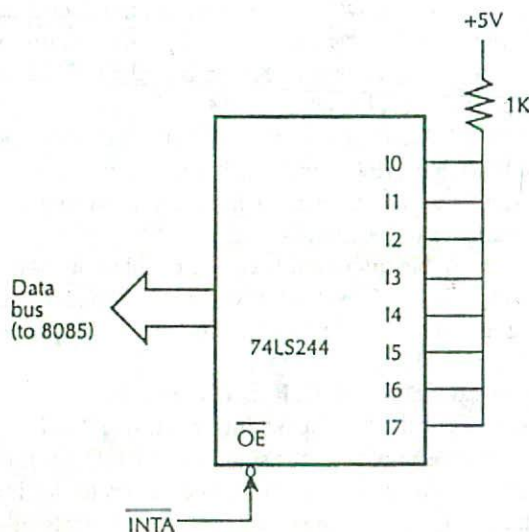
FIGURE 2.21   Using an octal buffer to provide RST7 instruction.
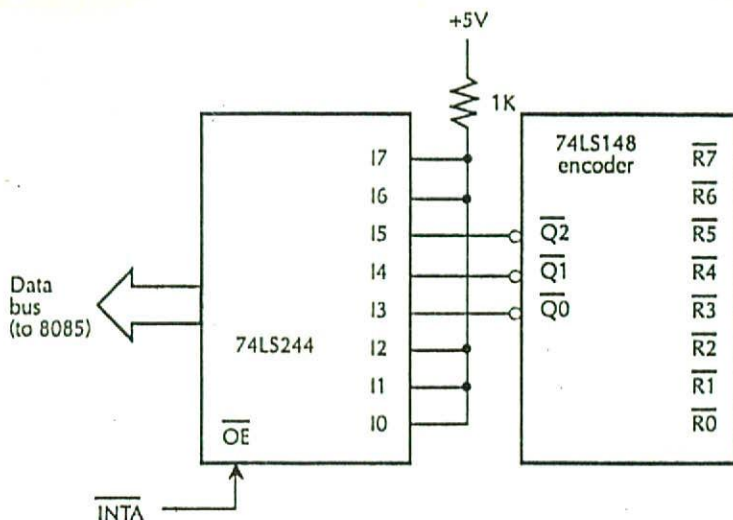
FIGURE 2.22   Forming eight RST instructions with a priority encoder.

Let us elaborate on Figure 2.23. Execution of the EI instruction sets the RS flip-flop of Figure 2.23 and makes one of the inputs to the AND gates #1 through #4 HIGH. Hence, in order for all the interrupts (except TRAP) to work, the interrupt system must be enabled. Execution of Disable Interrupts (DI) clears the RS flip-flop and disables all interrupts except TRAP. The SIM instruction outputs the contents of the accumulator which can be interpreted as shown in Figure 2.24.

The interrupt mask function is only executed if the mask set enable bit is 1. Suppose that if 06_{16} is stored in the accumulator and the SIM instruction is executed. A 1 will be sent to the interrupt mask for RST7.5 and RST6.5, and 0 will be sent to RST5.5. That is, in Figure 2.23, a 1 will be sent to the inputs of the AND gates #1 and #2, and a 0 will be sent to the AND gate #3, then inverted at the AND gate inputs (shown by circles), giving two LOW outputs disabling RST7.5, RST6.5, and a HIGH input to AND gate #3. Therefore, in order to enable RST7.5, RST6.5, or RST5.5, the interrupt system must be enabled by executing EI, the appropriate interrupt mask bit must be LOW by executing SIM, and the appropriate interrupt signal (leading edge or high level) at the respective pins must be available. For example, consider the RST7.5 interrupt. When the EI and SIM instructions are executed, the interrupt system can be enabled and also the interrupt mask bit for RST7.5 can be set to LOW, making the two inputs to AND gate #1 HIGH. The third input to this AND gate can be set to a HIGH by a leading edge at the RST7.5 pin. This sets the D flip-flop, thus making the output of the AND gate #1 HIGH, enabling RST7.5. The 8085 branches to location 003C16 where a 3-byte JMP instruction takes the program to the service routine. The RST5.5 and RST6.5 can similarly be explained from Figure 2.23.

The RIM instruction can be used to check whether one or more of the RST7.5, RST6.5, and RST5.5 interrupts are waiting to be serviced. The RIM instruction also provides the status of the mask bits for RST7.5, RST6.5, and RST5.5 and the status of the SID pin (HIGH or LOW). After execution of the RIM instruction, the 8085 loads the accumulator with 8 bits of data which can be interpreted as shown in Figure 2.25.

The bits are interpreted as follows:

Mask bits — Bits 0, 1, and 2 are the status of the mask bits for RST5.5, RST6.5, and RST7.5.
Interrupt Enable bit — This bit indicates whether one maskable interrupt capability is enabled (if this bit = 1) or disabled (if this bit = 0).
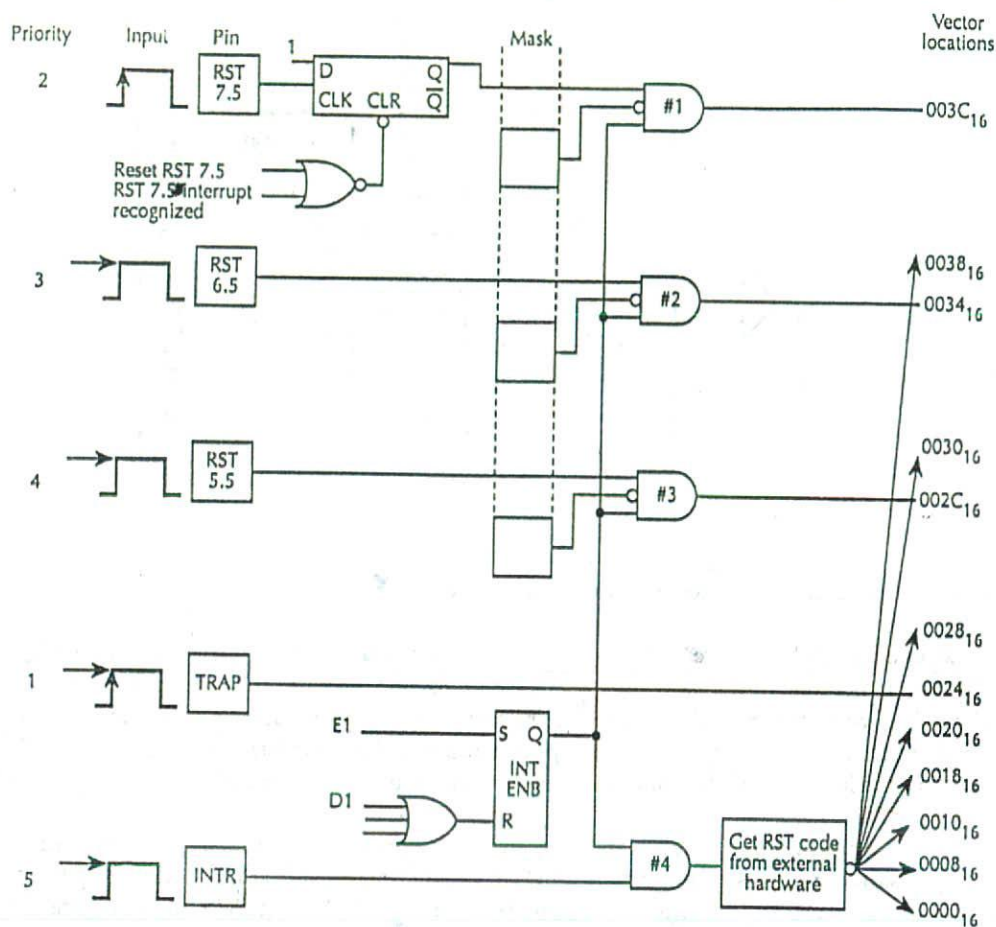
FIGURE 2.23   8085 interrupt structure.
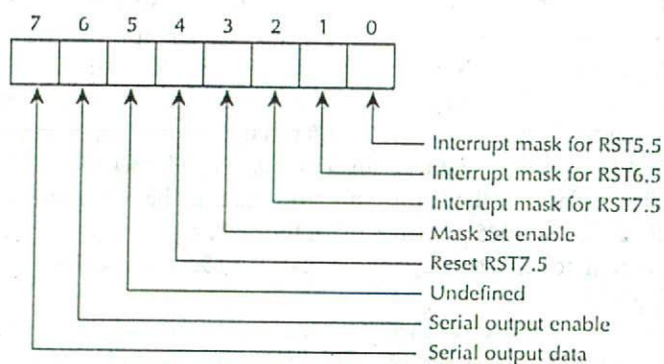
FIGURE 2.24   Interpretation of data output by the SIM instruction.

Interrupt Pending bits — A one in a particular bit position indicates that the particular RST is waiting to be serviced, while a 0 indicates that no interrupt is pending.
Serial Input data — These data indicate the status of the SID pin.

The RIM instruction can be used to read the interrupt pending bits in the accumulator. These bits can then be checked by software to determine whether any higher priority interrupts
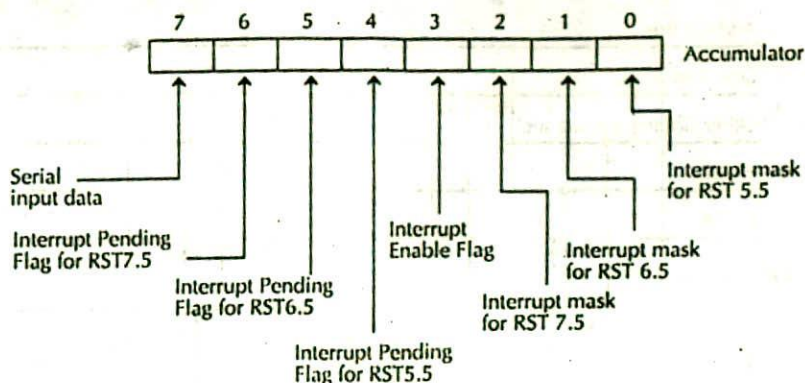
FIGURE 2.25   Execution of RIM instruction.

are pending. For example, suppose that RST6.5 and RST5.5 interrupts are used in a system. If RST5.5 interrupt first occurs, assuming no other higher level interrupts, the 8085 will execute the RST5.5 service routine. Note that while in this service routine, the RST6.5 interrupt (higher priority than RST5.5 interrupt) may occur. Since the 8085's maskable interrupt capability is disabled, the user can check whether the RST6.5 is pending while executing the RST5.5 service routine in one of two ways:

- By executing EI at the beginning of the RST5.5 service routine provided that the level at RST5.5 pin is LOW (RST5.5 disabled). This will service the RST6.5 immediately if pending and will then complete RST5.5 service routine (nested interrupts).
- By executing RIM and using rotate instruction to check whether RST6.5 is pending. The RIM instruction can be executed at several places in the RST5.5 service routine. If the RST6.5 is found to be pending, the interrupt can be enabled by executing EI in the RST5.5 service routine.

2.9.3 8085 DMA

The Intel 8257 DMA controller chip is a 40-pin DIP and is programmable. It is compatible with the 8085 microprocessor. The 8257 is a four-channel DMA controller with priority logic built into the chip. This means that the 8257 provides for DMA transfers for a maximum of up to four devices via the DMA request lines DRQ0 to DRQ3 (DRQ0 has the highest priority and DRQ3 the lowest). Associated with each DRQ is a DMA acknowledge (DACK0 to DACK3 for four DMA requests DRQ0 to DRQ3). Note that the DACK signals are active LOW. The 8257 uses the 8085 HOLD pin in order to take over the system bus. After initializing the 8257 by the 8085, the 8257 performs the DMA operation in order to transfer a block of data of up to 16,384 bytes between the memory and a peripheral without involving the microprocessor. A typical 8085-8257 interface is shown in Figure 2.26. An I/O device, when enabled by the 8085, can request a DMA transfer by raising the DMA request (DRQ) line of one of the channels of the 8257. In response, the 8257 will send a HOLD request (HRQ) to the 8085. The 8257 waits for the HOLD acknowledge (HLDA) from the 8085. On receipt of HLDA from the 8085, the 8257 generates a LOW on the DACK lines for the I/O device. Note that DACK is used as a chip select bit for the I/O device. The 8257 sends the READ or WRITE control signals, and data are transferred between the I/O and memory. On completion of the data transfer, the DACK0 is set to HIGH, and the HRQ line is reset to LOW in order to transfer control of the bus to the 8085. The 8257 utilizes four clock cycles in order to transfer 8 bits of data.

The 8257 has three main registers. These are a 16-bit DMA address register, a terminal count register, and a status register. Both address and terminal count registers must be initialized
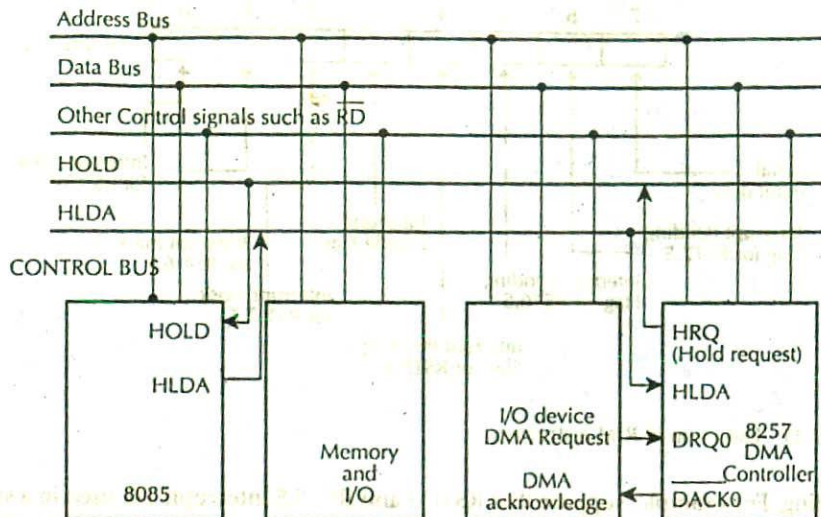
FIGURE 2.26   An 8085-8257 interface.

before a DMA operation. The DMA address register is initialized with the starting address of the memory to be written into or read from. The low-order 14 bits of the terminal count register are initialized with the value $(n - 1)$, where n is the desired number of DMA cycles. A Terminal Count (TC) pin on the 8257 is set to HIGH in order to indicate to the peripheral device that the present DMA cycle is the last cycle. An 8-bit status register in the 8257 is used to indicate which channels have attained a terminal count.

2.9.4   8085 SID and SOD Lines

Serial I/O is extensively used for data transfer between a peripheral device and the microprocessor. Since microprocessors perform internal operations in parallel, conversion of data from parallel to serial and vice versa is required to provide communication between the microprocessor and the serial I/O. The 8085 provides serial I/O capabilities via SID (Serial Input Data) and SOD (Serial Output Data) lines.

One can transfer data to or from the SID or SOD lines using the instruction RIM and SIM. After executing the RIM instruction, the bits in the accumulator are interpreted as follows:

1. Serial input bit is bit 7 of the accumulator.
2. Bits 0 to 6 are interrupt masks, the interrupt enable bit, and pending interrupts.

The SIM instruction sends the contents of the accumulator to the interrupt mask register and serial output line. Therefore, before executing the SIM, the accumulator must be loaded with proper data. The contents of the accumulator are interpreted as follows:

1. Bit 7 of the accumulator is the serial output bit.
2. The SOD enable bit is bit 6 of the accumulator. This bit must be 1 in order to output bit 7 of the accumulator to the SOD line.
3. Bits 0 to 5 are interrupt masks, enables, and resets.

Example 2.10

An 8085/8155-based microcomputer is required to input a switch via the SID line and output the switch status to an LED connected to the SOD line. Write an 8085 assembly language program to accomplish this.

Solution

```
        ORG   5000H
START   RIM           ; Bit 7 of 'A' is SID
        ORI   40H     ; Set SOD Enable to one
        SIM           ; output to LED
        JMP   START   ; REPEAT
```

Example 2.11

Write an 8085 assembly language program to implement the following requirements. i) if V1 > V2, the 8085/8155-based microprocessor system will read the switch input from port B. If switch is open (input high), turn the LED off. If switch is closed (input low), turn the LED on. ii) Repeat i) by using a) TRAP, b) RST6.5 with SID/SOD, c) INTR.

Assume

```
    CSR   = 00H
    Port A = 01H
    Port B = 02H
    Port C = 03H
```

Write all programs starting at 2000H and service routines at 3000H. Intialize SP to 20C0H.

i) *Solution*

```
            ORG    2000H
CSR     EQU    00H     ; Define CSR address.
PORTA   EQU    01H     ; Define Port A address.
PORTB   EQU    02H     ; Define Port B address.
PORTC   EQU    03H     ; Define Port C address.
        MVI    A,0CH   ; Set Port A/Port B input and Port
                         C output.
        OUT    CSR     ; Write to CSR.
        OUT    PORTC   ; Initialize bit 5 of port C. Turn
                         LED off.
SCAN    IN     PORTA   ; Get the data from Port A.
        RAR            ; Move the bit0 of A to Cy.
        JNC    SCAN    ; Go to SCAN if V1 < V2. Go to
                         next if V1>V2.
        IN     PORTB   ; Get the switch input from Port B.
        XRI    02H     ; Invert the input data.
        RAL            ; Move the data in bit 1 of Port A
                         to bit 5.
        RAL            ;
        RAL            ;
        RAL            ;
        OUT    PORTC   ; Write to Port C.

        JMP    SCAN    ; Loop
```
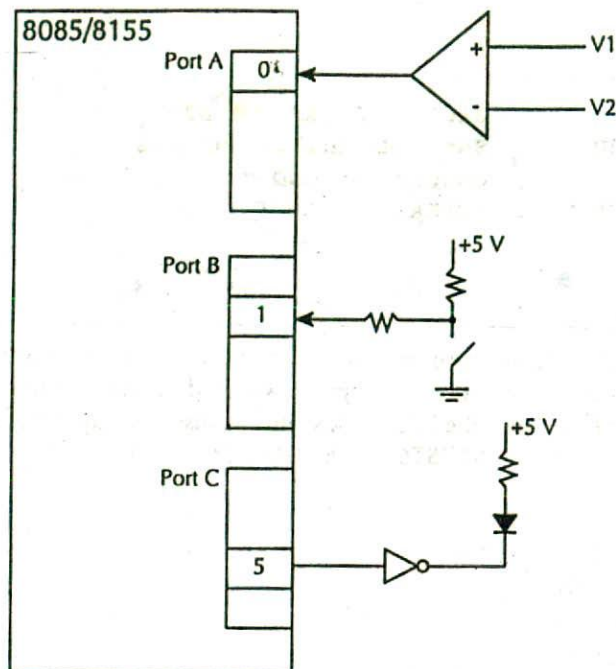
ii) *Solution*

```
a)      ORG    0024H      ; Trap vector address entry.
        JMP    3000H      ; Go to trap service routine.
        ORG    2000H
CSR     EQU    00H        ; Define CSR address.
PORTB   EQU    02H        ; Define Port B address.
PORTC   EQU    03H        ; Define Port C address.
START   LXI    SP,20C0H   ; Initialize SP.
        MVI    A,0CH      ; Set Port B input and Port C
                            output.
        OUT    CSR        ; Write to CSR.
        OUT    PORTC      ; Initialize bit 5 of Port C. Turn
                            LED off.
LOOP    NOP               ; Wait for TRAP request.
        JMP    Loop
        ORG    3000H      ; TRAP service routine entry.
        IN     PORTB      ; Get the switch input from Port
                            B.
        XRI    02H        ; Invert the input data.
        RAL               ; Move the data in bit 1 of A to
                            bit 5.
        RAL               ;
        RAL               ;
        RAL               ;
        OUT    PORTC      ; Write to Port C.
        JMP    START
```

ii) *Solution*

```
b)      ORG    0034H      ; RST6.5 vector address entry.
```
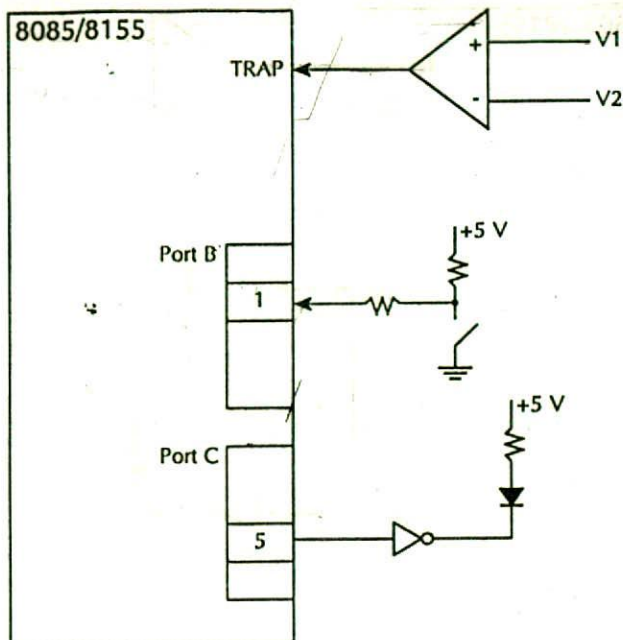
```
          JMP    3000H      ; Go to RST6.5 service routine.
          ORG    2000H
START     LXI    SP,20C0H   ; Initialize SP.
          MVI    A,40H      ; Initialize SOD. Turn LED off.
          SIM               ;
                            ;
          MVI    A,0DH      ; Set data for RST6.5.
          SIM               ; Enable RST6.5.
          EI                ; Enable interrupt.
LOOP      NOP
          JMP    LOOP       ; Wait for RST6.5 request.
          ORG    3000H      ; RST6.5 service routine entry.
          RIM               ; Get the data from SID>
          XRI    80H        ; Invert the input data in bit
                            ;   7 of A.
          ORI    40H        ; Enable SOD.
          SIM               ; Write to SOD.
          JMP    START
```

ii) *Solution*

   c)
```
          ORG    0030H      ; RST6 vector address entry
          JMP    3000H      ; Go to RST6 service routine
          ORG    2000H
CSR       EQU    00H        ; Define CSR address.
PORTB     EQU    02H        ; Define Port B address.
PORTC     EQU    03H        ; Define Port C address.
BEGIN     LXI    SP,20C0H   ; Initialize SP.
          MVI    A,0CH      ; Set Port B input and Port C
                            ;   output.
          OUT    CSR        ; Write to CSR.
```
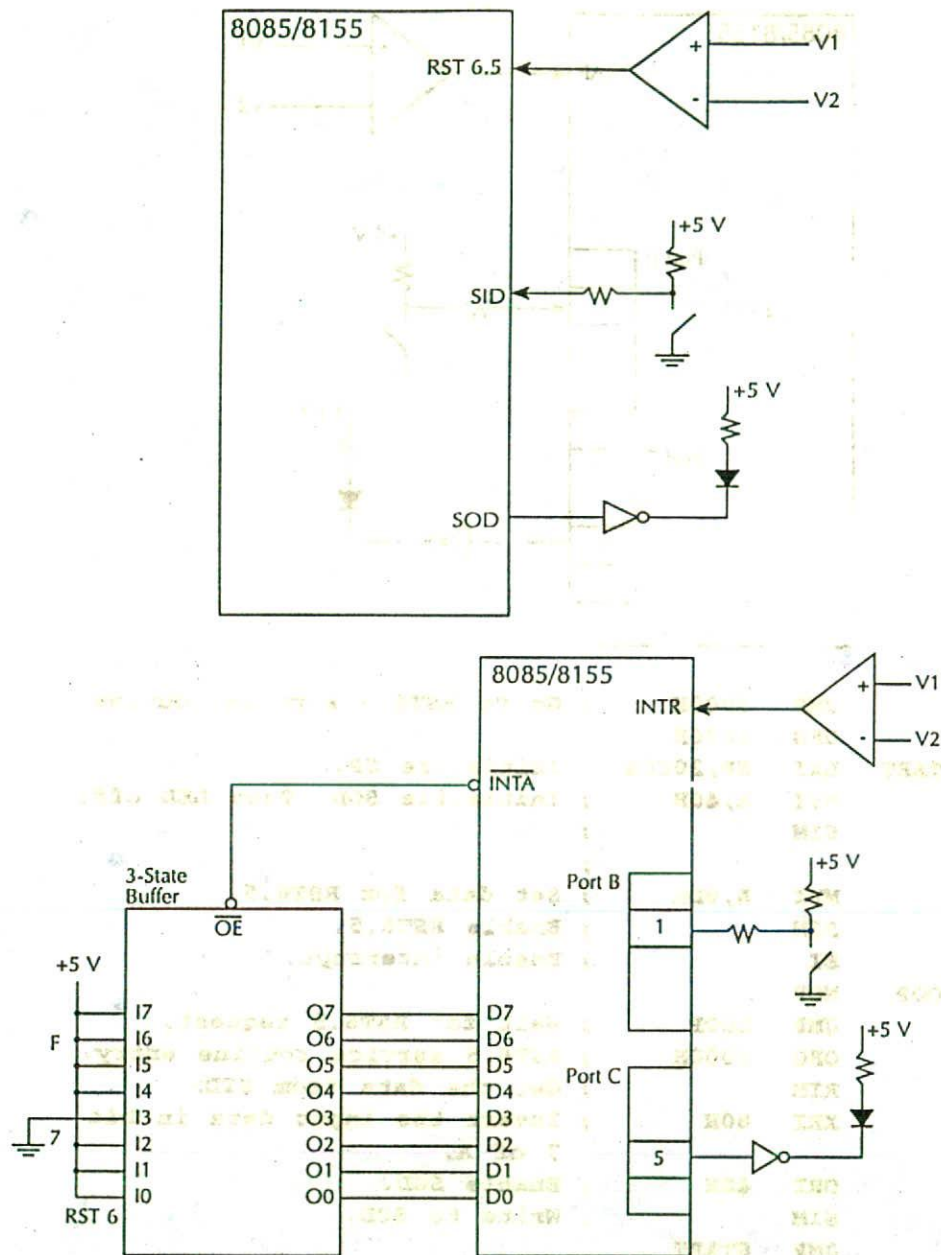
```
        OUT   PORTC      ; Initialize bit 5 of port C. Turn
                         ; LED off.
        EI               ; Enable interrupt.
LOOP    NOP              ; Wait for INTR request.
        JMP   LOOP
        ORG   3000H      ; RST6 service routine entry.
        IN    PORTB      ; Get the switch input from Port
                         ; B.
        XRI   02H        ; Invert the input data.
        RAL              ; Move the data in bit 1 of A to
                         ; bit 5.
```

```
RAL                    ;
RAL                    ;
RAL                    ;
OUT     PORTC          ; Write to Port C.
JMP     BEGIN
```

2.10 8085-Based System Design

In order to illustrate the concepts associated with 8085-based system design, a microcomputer with 2K EPROM (2716), 256 byte RAM, and 3 ports (8155) is designed. A hardware schematic is included. Also, an 8085 assembly language program is provided to multiply 4-bit unsigned numbers entered via DIP switches connected to port A. The 8-bit product is displayed on two seven-segment displays interfaced via port B. Repeated addition will be used for multiplication. Figure 2.27 shows a schematic of the hardware design. Full decoding using the 74LS138 decoder is utilized. Texas Instruments TIL 311's displays with on-chip decoder are used. The memory and I/O maps are given in the following:

1. Memory map

   2716

A15	A14	A13	A12	A11	A10	A9	A8	AD7	AD6	AD5	AD4	AD3	AD2	AD1	AD0
0	0	0	0	0	{- - - - - - - - - - - - - - - - all zeros to ones - - - - - - - - - - - -}										

   Result 0000H-07FFH

   8155

A15	A14	A13	A12	A11	A10	A9	A8	AD7	AD6	AD5	AD4	AD3	AD2	AD1	AD0
0	0	0	0	1	0	0	0	{- - - - - - - - - - - all zeros to ones - - - - - - - -}							

   Result 0800H-08FFH

2. I/O map using standard I/O

   Ports

A15 AD7	A14 AD6	A13 AD5	A12 AD4	A11 AD3	A10 AD2	A9 AD1	A8 AD0		
0	0	0	0	1	0	0	0	CSR	08H
0	0	0	0	1	0	0	1	Port A	09H
0	0	0	0	1	0	1	0	Port B	0AH
0	0	0	0	1	0	1	1	Port C	0BH

A listing of the 8085 assembly language program for performing 4 bit × 4 bit multiplication is provided below:

```
CSR EQU    08H
PORT A  EQU    09H
PORT B  EQU    0AH
REPEAT  MVI    A,02H  ; Configure Port A as input
        OUT    CSR    ; and Port B as output
        MVI    L,00H  ; Initialize product to zero
        IN     Port A ; Input multiplier and
                        multiplicand
```
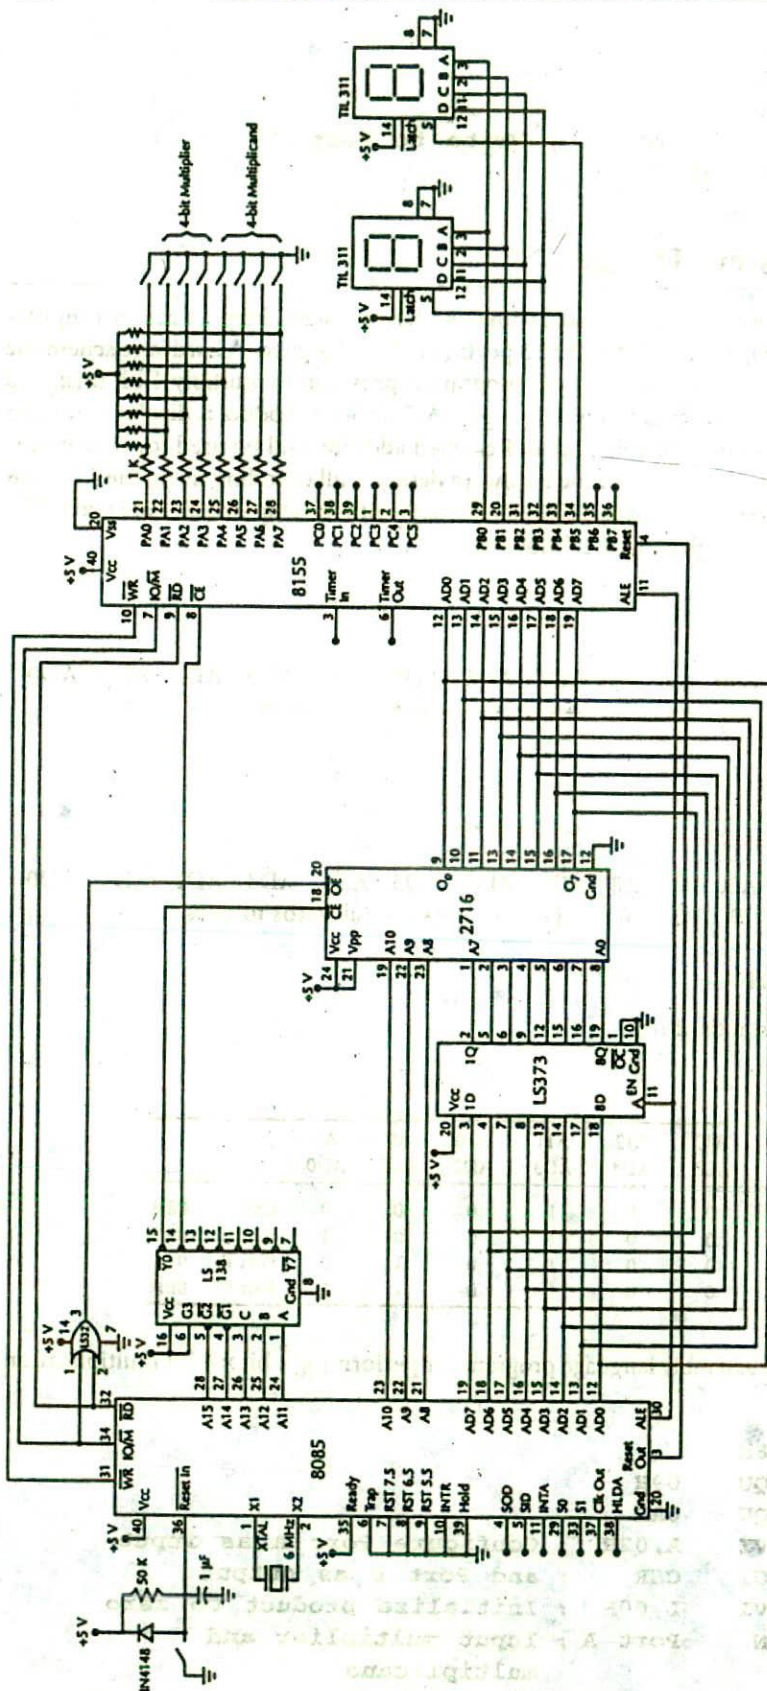
FIGURE 2.27 8085-based system design.

```
                MOV   B,A      ; Save multiplicand and
                               ; multiplier in B
                ANI   OFH      ; Mask multiplicand and retain
                               ; multiplier
                MOV   C,A      ; Save 4-bit multiplier in C
                MOV   A,B      ; Move multiplicand and
                               ; multiplier to A
                RAR            ; Move 4-bit multiplicand
                RAR            ; (upper nibble of A)
                RAR            ; into LOW nibble
                RAR            ; of
                               ; accumulator
                ANI   OFH      ; Mask high nibble and retain
                MOV   D,A      ; multiplicand
    START       ADD   D        ; Perform repeated addition
                DCR   C        ; Decrement multiplier value
                MOV   L,A      ; Save sum
                JNZ   START    ; if Z = 0, repeat addition
                               ; else product in L
                MOV   A,L      ; Move product to A
                ANI   OFH      ; Retain product LOW nibble and
                MOV   H,A      ; Save in H
                MOV   A,L      ; Move product to A
                RAR            ; Move
                RAR            ; high nibble
                RAR            ; of product
                RAR            ; to LOW nibble
                               ; of A
                ANI   2FH      ; Retain high nibble of
                               ; product and enable latch
                OUT   PORTB    ; of high hex display and
                               ; disable low latch
                               ; Display high nibble of product
                MOV   A,H      ; Move LOW product nibble to A
                ANI   1FH      ; Enable LOW latch and disable
                               ; high latch
                OUT   PORTB    ; Display LOW nibble of product
                JMP   REPEAT   ; On low display and continue
```

The above program can be assembled. The 2716 can then be programmed by using an EPROM programmer with the machine code of the preceding program starting at location 000H. Then upon activation of the switch at the 8085 reset input, the dip switch low and high nibbles will be multiplied and the result of the multiplication will be displayed on the two TIL311s. By changing DIP switch inputs at port A, new results can be displayed.

Questions and Problems

2.1   Compare the main differences of 8085AH, 8085AH-2, and 8085AH-1 with the 8085.

2.2   What is the primary purpose of the 8085 H-L pair with respect to external memory? List two of its main functions.

2.3 Identify the addressing modes of the following instructions:
   i)   `MOV A, M`
   ii) `RAR`
   iii) `STAX D`
   iv) `LDA START`

2.4 Assume register pair BC contains $7F02_{16}$. For the following 8085 assembly language program, determine the carry, zero, parity, and sign flags after execution of the MVI A,05H instruction:

```
ORG    4000H
LXI    SP, 2050H
PUSH   B
MVI    A,02H
ADI    03H
POP    PSW
MVI    A, 05H
HLT
```

2.5 What function is performed by each of the following 8085 instructions:
   i)   `XOR A`
   ii) `MOV D,D`
   iii) `DAD H`

2.6 i) What is the function of the 8085 ALE pin?
     ii) If a crystal of 4 MHz is connected to the 8085 X1X2 pins, what is the 8085 internal clock period?

2.7 Write an 8085 assembly language program to divide an 8-bit unsigned number in accumulator by 4. Neglect remainder. Use minimum number of instructions. Store result in A.

2.8 Using the simplest possible algorithm, write an 8085 assembly language program with a minimum number of instructions to divide a 16-bit unsigned number in DE by 16. Neglect remainder. Store result in DE.

2.9 Assume the contents of DE are 2050H. Write an 8085 assembly language program to unconditionally branch to 2050H. Do not use any conditional or unconditional Jump or CALL or return instructions. Use three instructions maximum.

2.10 Write an 8085 assembly language program using a minimum number of instructions to add the 16-bit numbers in BC, DE, and HL. Store the 16-bit result in DE.

2.11 Write an 8085 assembly language program to add two 24-bit numbers located in memory address 2000_{16} through 2005_{16} with the most significant byte of the first number in 2000_{16} and the most significant byte of the second number in 2003_{16}. Store the 24-bit result in three consecutive bytes starting at address 2000_{16}.

2.12 Write a subroutine in 8085 assembly language to divide a 16-bit unsigned number in BC by 32_{10}. Neglect the remainder. Use the simplest possible algorithm. Also, write the main program in 8085 assembly language to perform all the initializations. The main program will then call the subroutine, store the 16-bit quotient in DE, and stop.

2.13 Write an 8085 assembly language program to shift the contents of the DE register twice to left without using any ROTATE instructions. After shifting, if the contents of DE are nonzero, store $FFFF_{16}$ in the HL pair. On the other hand, if the contents of DE pair are zero, then store 0000_{16} in the HL pair.

2.14 Write an 8085 assembly language program to check the parity of an 8-bit number in "A" without using any instructions involving the parity flag. Store EE in location 3000 if the parity is even; otherwise, store DD in location 3000_{16}.

2.15 Write an 8085 assembly language program to move a block of data of length 100_{10} from the source block starting at 2000_{16} to the destination block starting at 3000_{16}.

2.16 Write a subroutine in 8085 assembly language to divide an 8-bit unsigned number Xi by 2. Also, write the main program in 8085 assembly language which will call the subroutine to compute

$$\sum_{i=1}^{3} \frac{Xi}{2}$$

Store the result in location 5000_{16}. Use a minimum number of instructions.

2.17 Assume an 8085/8355-based microcomputer. Suppose that four switches are connected to bits 0 through 3 of port A, an LED to bit 4 of port A, and another LED to bit 2 of port B. If the number of low switches is even, turn the port A LED ON and port B LED OFF. If the number of low switches is odd, turn port A LED OFF and port B LED ON. Write an 8085 assembly language program to accomplish the above.

2.18 Draw a simplified diagram using an 8085, one 8156, and one 8355 to include the following memory map and I/O ports:
   i)   *8355*   *8156*
        2000_{16} thru $27FF_{16}$ 4000_{16} thru $40FF_{16}$
   ii) Using memory-mapped I/O, configure ports A and B of the 8355 at addresses 8000_{16} and 8001_{16}. Show only connections for the lines which are pertinent.

2.19 It is desired to interface a pump to a 8085/8355/8156-based microcomputer. The pump will be started by the microcomputer at the trailing edge of a start pulse via the SOD pin. When the pump runs, a HIGH status signal from the pump called "Pump Running" will be used to interrupt the microcomputer via its INTR interrupt. In response, the microcomputer will turn an LED on connected to bit 7 of the 8355 port A. Assume DDRA as the data direction register for this port.
   i)   Draw a simplified schematic for accomplishing the above.
   ii) Write the main and interrupt service routines using the 8085 assembly language programs.

2.20 Will the circuit shown work? If so, determine the memory and I/O maps in hex. If not, justify briefly, modify the circuit and then determine the memory and I/O maps. Use only the pins and signals as shown. Also, assume all don't cares to be zeros.
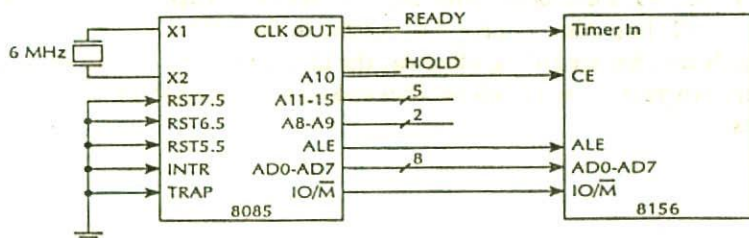
2.21
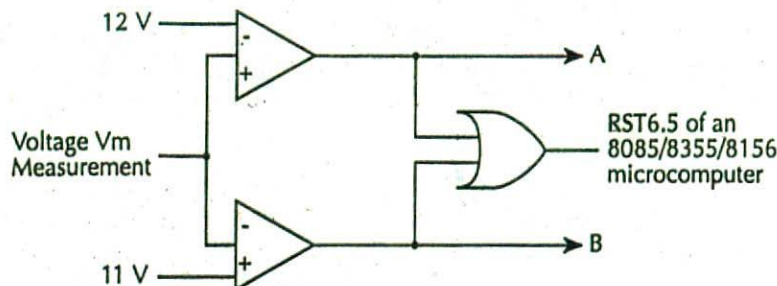
In the above, if Vm > 12 V, turn an LED ON connected at bit 3 of port A. On the other hand, if Vm < 11 V, turn the LED OFF. Use ports, registers, and memory locations as needed.

    i)   Draw a hardware block diagram showing the microcomputer and the connections of the above diagram to its ports.

    ii)   Write main program and service routines in 8085 assembly language.

2.22   Interface the following 8156 RAM chip to obtain the following memory map: 0400H thru 04FFH. Show only the connections for the pins which are shown. Assume all unused address lines to be zero.
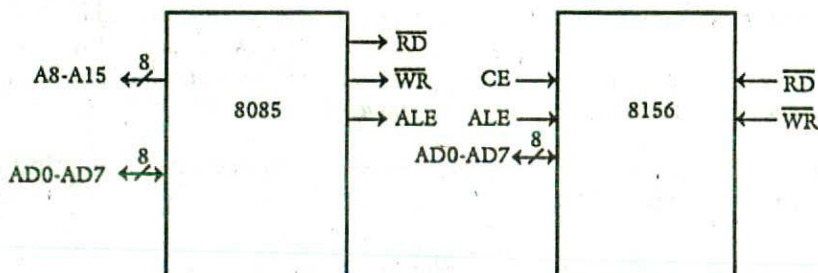
2.23
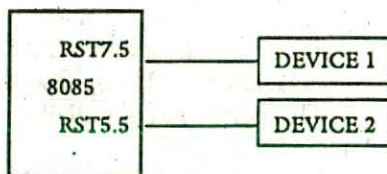
If device 2 is presently being serviced by the 8085 and the device 1 interrupt occurs, explain briefly what the user needs to do in the service routine of device 2 in order that device 1 will be serviced before device 2.

2.24   Assume an 8085/8355 microcomputer. Suppose that two switches are connected to bit 1 of Port A and the SID line. Also, an LED is connected to the SOD line. It is desired to turn the LED ON if both switches are HIGH; otherwise the LED is to be turned off. Write an 8085 assembly language program to input both switches and turn the LED ON or OFF based on the above conditions.
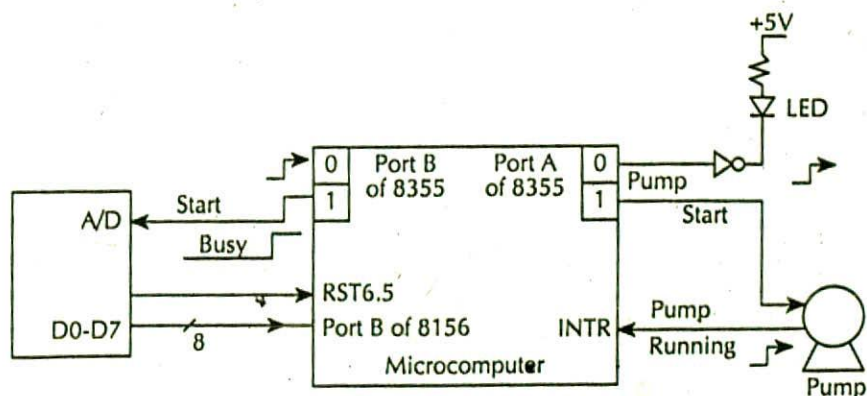
2.25 It is desired to interface a pump and an A/D converter to an 8085/8355/8156-based microcomputer as follows: The pump can be started by a HIGH at the pump start signal. If the pump starts, the pump-running signal goes to HIGH; otherwise the pump running stays LOW. The A/D converter can be started by a HIGH and the conversion is completed when the Busy signal is HIGH. Start the pump and A/D converter at nearly the same time. If the pump runs, the LED is to be turned on. On the other hand, if the Busy signal goes to HIGH, the A/D converter's output is to be read.

i)   Draw a simplified hardware schematic to accomplish the above.
ii) Write main program and service routines in 8085 assembly language. Include the interrupt priority concept in the programs.

2.26 Assume an 8085/8355-based microcomputer. Write an 8085 assembly language program to turn an LED (LED 1) ON connected to bit 3 of port A if a 16-bit number in DE register pair is negative after shifting once to the left; if the number is positive after shifting, turn the LED (LED 1) OFF. Also, after shifting, if there is a sign change of the 16-bit number, turn another LED (LED 2) ON connected to bit 1 of port A; otherwise, turn the LED (LED 2) OFF.

2.27 Design an 8085/2716/8155-based microcomputer to include the following:
i)   4K EPROM, 512 bytes static RAM, four 8-bit and two 6-bit ports. Use standard I/O and linear decoding.
ii) Repeat i) except use fully decoding using a 3 × 8 decoder.
iii) Repeat i) except use memory-mapped I/O.
iv) Repeat i) except use memory-mapped I/O and full decoding.

Draw a neat schematic of each design and also determine memory and I/O maps in each case.

INTEL 8086

This chapter describes the internal architecture, addressing modes, instruction set, and I/O techniques associated with the 8086 microprocessor. Interfacing capabilities to typical memory and I/O chips such as the 2716, 6116, and 8255 are included.

A design technique is presented showing interconnection of the 8086 to 2716 EPROM, 6116 RAM, and 8255 I/O chips. The memory and I/O maps are then determined.

3.1 Introduction

The 8086 is Intel's first 16-bit microprocessor.

The 8086 is designed using the HMOS technology and contains approximately 29,000 transistors. The 80C86A is the low power version of the 8086 designed using HCMOS technology. The 8086 is packaged in a 40-pin CERDIP or plastic package and requires a single 5V power supply. The 8086 can be operated at three different clock speeds. The standard 8086 runs at 5 MHz internal clock frequency, whereas the 8086-1 and 8086-2 run at internal clock frequencies of 10 and 8 MHz, respectively. An external clock generator/driver chip such as the Intel 8284 is needed to generate the 8086 clock input signal. The 8284 divides the external crystal input internally by three. This means that for a 5-MHz 8086 internal clock, the 8284's X1 and X2 pins must be connected to a 15-MHz crystal. The 8284 will then generate a 5 MHz clock at its CLK pin which should be connected to the 8086 CLK input.

The 8086 has a 20-bit address and, hence, it can directly address up to one megabyte (2^{20}) of memory. The 8086 uses a segmented memory. An interesting feature of the 8086 is that it prefetches up to six instruction bytes from memory and queues them in order to speed up instruction execution.

There are some advantages of working with the segmented memory. First of all, after initializing the 16-bit segment registers, the 8086 has to deal with only 16-bit effective addresses. That is, the 8086 has to manipulate and store 16-bit address components. Secondly, because of memory segmentation, the 8086 can be effectively used in time-shared systems. For example, in a time-shared system, several users share a microprocessor. The microprocessor works with one user's program for say, 10 milliseconds. After spending 10 milliseconds with each of the other users, the microprocessor returns to execute the first user's program. Each time the microprocessor switches from one user's program to the next, it must execute a new section of code and new sections of data. Segmentation makes it easy to switch from one user program (and data) to another.

The memory of an 8086-based microcomputer is organized as bytes. Each byte can be uniquely addressed with 20-bit addresses of 00000_{16}, 00001_{16}, 00002_{16} ..., $FFFFF_{16}$. An 8086 16-

bit word consists of any two consecutive bytes; the low-addressed byte is the low byte of the word and the high-addressed byte contains the high byte as follows:

Low byte of the word	High byte of the word
07_{16}	26_{16}
Address 00520_{16}	Address 00521_{16}

The 16-bit word stored at the even address 00520_{16} is 2607_{16}.

Next consider a word stored at an odd address as follows:

Low byte of the word	High byte of the word
05_{16}	$3F_{16}$
Address 01257_{16}	Address 01258_{16}

The 16-bit word stored at the odd address 01257_{16} is $3F05_{16}$. Note that for word addresses, the programmer uses the low-order address (odd or even) to specify the whole 16-bit word.

The 8086 always accesses a 16-bit word to or from memory. The 8086 can read a 16-bit word in one operation if the first byte of the word is at an even address. On the other hand, the 8086 must perform two memory accesses to two consecutive memory even addresses, if the first byte of the word is at an odd address. In this case, the 8086 discards the unwanted bytes of each. For example, consider MOV BX, [ADDR]. Note that the X or H (or L) following the 8086 register name in an instruction indicates whether the transfer is 16-bit (for X) or 8-bit (for H or L). The instruction, MOV BX, [ADDR] moves the contents of a memory location addressed by ADDR into the 8086 16-bit register BX. Now, if ADDR along with the data segment register provides a 20-bit even address such as 30024_{16}, then this MOV instruction loads the low (BL) and high (BH) bytes of the 8086 16-bit register BX with the contents of memory locations 30024_{16} and 30025_{16}, respectively, in a single access. Now, if ADDR is an odd address such as 40005_{16}, then the MOV BX, [ADDR] instruction loads BL and BH with the contents of memory locations 40005_{16} and 40006_{16}, respectively, in two accesses. Note that the 8086 accesses locations 40004_{16} and 40005_{16} in the first operation but discards the contents of 40004_{16}, and in the second operation accesses 40006_{16} and 40007_{16} but ignores the contents of 40007_{16}.

Next, consider a byte move such as MOV BH, [ADDR]. If ADDR is an even address such as 50002_{16}, then this MOV instruction accesses both 50002_{16} and 50003_{16}, but loads BH with the contents of 50002_{16} and ignores the contents of 50003_{16}. However, if ADDR is an odd address such as 50003_{16}, then this MOV loads BH with the contents of 50003_{16} and ignores the contents of 50002_{16}.

The 8086 family consists of two types of 16-bit microprocessors — the 8086 and 8088. The 8088 has an 8-bit external data path to memory and I/O, while the 8086 has a 16-bit external data path. This means that the 8088 will have to do two read operations to read a 16-bit word into memory. In most other respects, the processors are identical. Note that the 8088 accesses memory in bytes. No alterations are needed to run software written for one microprocessor on the other. Because of similarities, only the 8086 will be considered here. The 8088 was used in designing the original IBM Personal computer.

An 8086 can be configured as a small uniprocessor system (minimum mode if the $\overline{\text{MN}/\text{MX}}$ pin is tied to HIGH) or as a multiprocessor system (maximum mode when MN/MX pin is tied to LOW). In a given system, the MN/MX pin is permanently tied to either HIGH or LOW. Some of the 8086 pins have dual functions depending on the selection of the MN/MX pin level. In the minimum mode (MN/MX pin high), these pins transfer control signals

directly to memory and input/output devices. In the maximum mode (MN/MX pin low), these same pins have different functions which facilitate multiprocessor systems. In the maximum mode, the control functions normally present in minimum mode are performed by a support chip, the 8288 bus controller.

Due to technological advances, Intel introduced the high performance 80186 and 80188 which are enhanced versions of the 8086 and 8088, respectively. The 8-MHz 80186/80188 provides two times greater throughput than the standard 5-MHz 8086/8088. Both have integrated several new peripheral functional units such as a DMA controller, a 16-bit timer unit, and an interrupt controller unit into a single chip. Just like the 8086 and 8088, the 80186 has a 16-bit data bus and the 80188 has an 8-bit data bus; otherwise, the architecture and instruction set of the 80186 and 80188 are identical. The 80186/80188 has an on-chip clock generator so that only an external crystal is required to generate the clock. The 80186/80188 can operate at 6 MHz, 8 MHz, and other frequencies. Like the 8085, the crystal frequency is divided by 2 internally. In other words, external crystals of 12 or 16 MHz must be connected to generate the 6- or 8-MHz internal clock frequency. The 80186/80188 is fabricated in a 68-pin package. Both processors have on-chip priority interrupt controller circuits to provide five interrupt pins. Like the 8086/8088, the 80186/80188 can directly address one megabyte of memory. The 80186/80188 is provided with 10 new instructions beyond the 8086/8088 instruction set. Examples of these instructions include INS and OUTS for inputting and outputting string byte or string word. The 80286, on the other hand, has added memory protection and management capabilities to the basic 8086 architecture. An 8-MHz 80286 provides up to six times greater throughput than the 5-MHz 8086. The 80286 is fabricated in a 68-pin package. The 80286 can be operated at 6, 8, 10, 12.5, 16.67, or 20 MHz clock frequency. The 80286 is typically used in a multiuser or multitasking system. The 80286 was used as the CPU of the IBM PC/AT Personal computer. Intel's 32-bit microprocessor family includes 80386, 80486 and Pentium microprocessors which will be covered later in this book.

3.2 8086 Architecture

Figure 3.1 shows a block diagram of the 8086 internal architecture. As shown in the figure, the 8086 microprocessor is internally divided into two separate functional units. These are the Bus Interface Unit (BIU) and the Execution Unit (EU). The BIU fetches instructions, reads data from memory and ports, and writes data to memory and I/O ports. The EU executes instructions that have already been fetched by the BIU. The BIU and EU function independently. The BIU interfaces the 8086 to the outside world. The BIU contains segment registers, instruction pointer, instruction queue, and address generation/bus control circuitry to provide functions such as fetching and queuing of instructions, and bus control.

The BIU's instruction queue is a First-In-First-Out (FIFO) group of registers in which up to six bytes of instruction code are prefetched from memory ahead of time. This is done in order to speed up program execution by overlapping instruction fetch with execution. This mechanism is known as pipelining.

The bus control logic of the BIU generates all the bus control signals such as read and write signals for memory and I/O. The 8086 contains the on-chip logical address to physical address mapping hardware. The programmer works with the logical address which includes the 16-bit contents of a segment register and a 16-bit displacement or offset value. The 8086 on-chip mapping hardware translates this logical address to 20-bit physical address which it then generates on its twenty addressing pins.

The BIU has four 16-bit segment registers. These are the Code Segment (CS), the Data Segment (DS), the Stack Segment (SS), and the Extra Segment (ES). The 8086's one megabyte memory is divided into segments of up to 64K bytes each. The 8086 can directly address four segments (256K byte within the 1 Mbyte memory) at a particular time. Programs obtain access to code and data in the segments by changing the segment register contents to point to the
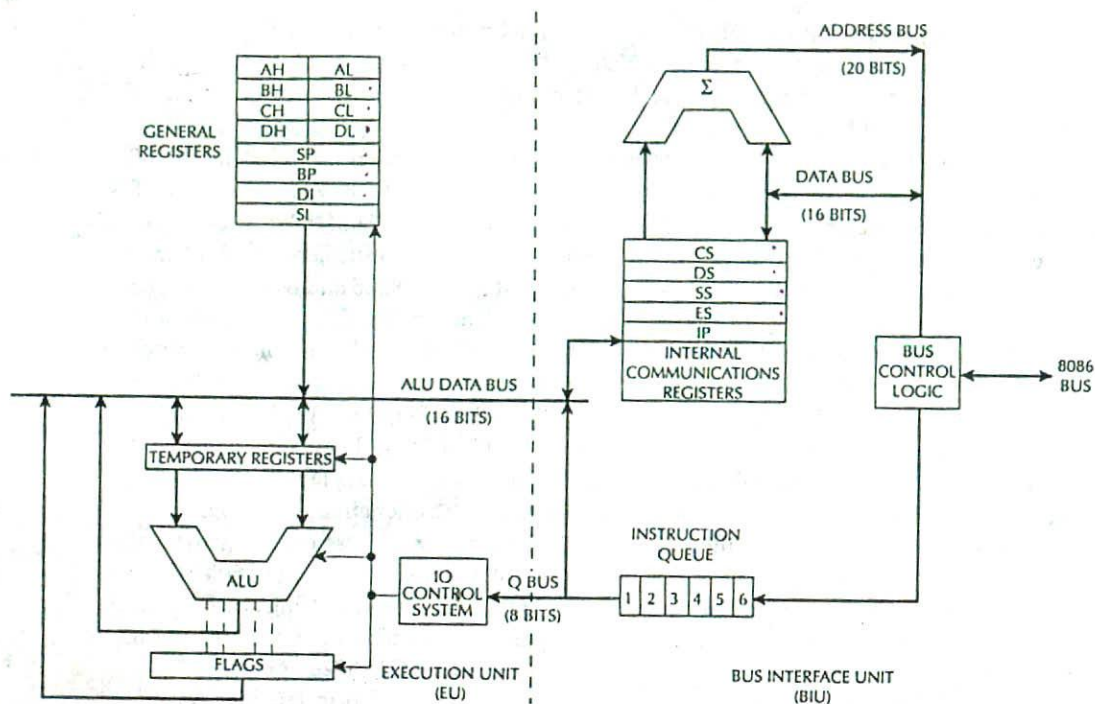
FIGURE 3.1   Internal architecture of the 8086.

desired segments. All program instructions must be located in main memory pointed to by the 16-bit CS register with a 16-bit offset in the segment contained in the 16-bit instruction pointer (IP). The BIU computes the 20-bit physical address internally using the programmer-provided logical address (16-bit contents of CS and IP) by logically shifting the contents of CS four bits to left and then adding the 16-bit contents of IP. For example, if $[CS] = 456A_{16}$ and $[IP] = 1620_{16}$, then the 20-bit physical address is generated by the BIU as follows:

$$\text{Four times logically shifted } [CS] \text{ to left} = 456A0_{16}$$

$$+ [IP] \text{ as offset} = 1620_{16}$$

$$20\text{-bit physical address} = 46CC0_{16}$$

The SS register points to the current stack. The 20-bit physical stack address is calculated from SS and SP for stack instructions such as PUSH and POP. The programmer can use the BP register instead of SP for accessing the stack using the based addressing mode. In this case, the 20-bit physical stack address is calculated from BP and SS.

The DS register points to the current data segment; operands for most instructions are fetched from this segment. A 16-bit offset (Effective Address, EA) along with the 16-bit contents of DS are used for computing the 20-bit physical address.

The ES register points to the extra segment in which data (in excess of 64K pointed to by DS) is stored. String instructions use ES and DI to determine the 20-bit physical address for the destination, and DS and SI for the source address.

The segments can be contiguous, partially overlapped, fully overlapped, or disjoint. An example of how five segments (segment 0 through segment 4) may be stored in physical memory are shown below:
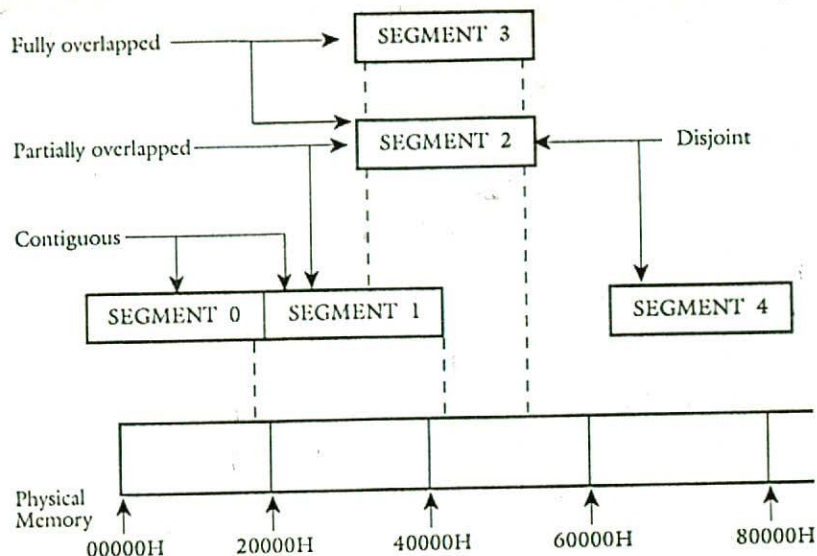
In the above, SEGMENTs 0 and 1 are contiguous (adjacent), SEGMENTs 1 and 2 are partially overlapped, SEGMENTs 2 and 3 are fully overlapped, and SEGMENTs 2 and 4 are disjoint. Every segment must start on 16-byte memory boundaries.

Typical examples of values of segments should then be selected based on physical addresses starting at 00000_{16}, 00010_{16}, 00020_{16}, 00030_{16} ..., $FFFF0_{16}$. A physical memory location may be mapped into (contained in) one or more logical segments. Many applications can be written to simply initialize the segment registers and then forget them. One can then work with a 64K memory as with the 8085.

The EU decodes and executes instructions. A decoder in the EU control system translates instructions. The EU has a 16-bit ALU for performing arithmetic and logic operations.

The EU has eight 16-bit general registers. These are AX, BX, CX, DX, SP, BP, SI, and DI. The 16-bit registers AX, BX, CX, and DX can each be used as two 8-bit registers (AH, AL, BH, BL, CH, CL, DH, DL). For example, the 16-bit register DX can be considered as two 8-bit registers DH (high byte of DX) and DL (low byte of DX). The general-purpose registers AX, BX, CX, and DX are named after special functions carried out by each one of them. For example, the AX is called the 16-bit accumulator while the AL is the 8-bit accumulator. The use of accumulator registers is assumed by some instructions. The Input/Output (IN or OUT) instructions always use AX or AL for inputting/outputting 16- or 8-bit data to or from an I/O port.

Multiplication and division instructions also use AX or AL. The AL register is the same as the 8085 A register.

The BX register is called the base register. This is the only general-purpose register, the contents of which can be used for addressing 8086 memory. All memory references utilizing these register contents for addressing use DS as the default segment register. The BX register is similar to 8085 HL register. In other words, 8086 BH and BL are equivalent to 8085 H and L registers, respectively.

The CX register is known as the counter register. This is because some instructions such as shift, rotate, and loop instructions use the contents of CX as a counter. For example, the instruction LOOP START will automatically decrement CX by 1 without affecting flags and will check if [CX] = 0. If it is zero, the 8086 executes the next instruction; otherwise the 8086 branches to the label START.

The data register DX is used to hold high 16-bit result (data) in 16×16 multiplication or high 16-bit dividend (data) before a $32 \div 16$ division and the 16-bit remainder after the division.
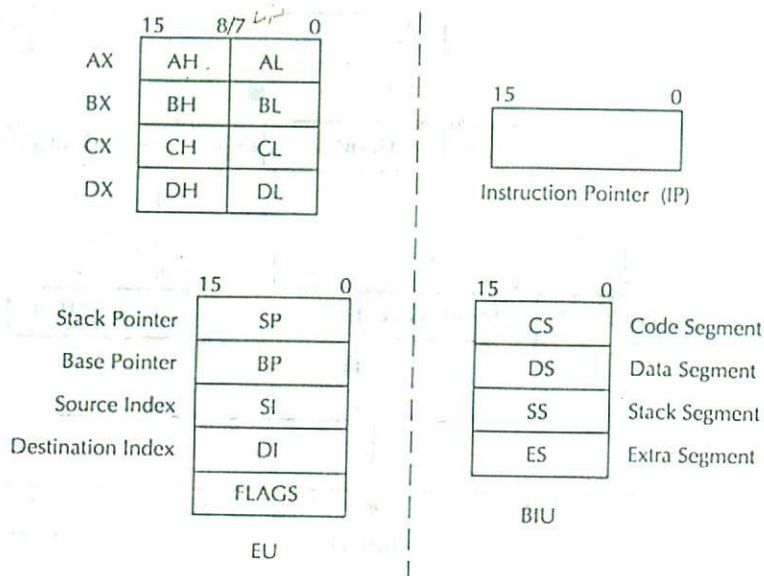
FIGURE 3.2　8086 registers.

The two pointer registers, SP (stack pointer) and BP (base pointer), are used to access data in stack segment. The SP is used as an offset from the current SS during execution of instructions that involve stack segment in external memory. The SP contents are automatically updated (incremented or decremented) due to execution of POP or PUSH instruction.

The base pointer contains an offset address in the current SS. This offset is used by the instructions utilizing the based addressing mode.

The FLAG register in the EU holds the status flags after an ALU operation. Figure 3.2 shows the 8086 registers.

The 8086 has six one-bit flags. Figure 3.3 shows the flag register. The AF (Auxiliary carry Flag) is used by BCD arithmetic instructions. AF = 1 if there is a carry from the low nibble (4-bit) into the high nibble or a borrow from the high nibble into the low nibble of the low-order 8-bit of a 16-bit number. The CF (Carry Flag) is set if there is a carry from addition or borrow from subtraction. The OF (Overflow Flag) is set if there is an arithmetic overflow, that is, if the size of the result exceeds the capacity of the destination location. An interrupt on overflow instruction is available which will generate an interrupt in this situation. The SF (Sign Flag) is set if the most significant bit of the result is one (negative) and is cleared to zero for non-negative result. The PF (Parity Flag) is set if the result has even parity; PF is zero for odd parity of the result. The ZF (Zero Flag) is set if result is zero; ZF is zero for nonzero result.

The 8086 has three control bits in the flag register which can be set or reset by the programmer: Setting DF (Direction Flag) to one causes string instructions to autodecrement the appropriate index register(s), and clearing DF to zero causes string instructions to autoincrement. Setting IF (Interrupt Flag) to one causes the 8086 to recognize external maskable interrupts; clearing IF to zero disables these interrupts. Setting TF (Trace Flag) to one places the 8086 in the single-step mode. In this mode, the 8086 generates an internal interrupt

15				12	11	10	9	8	7	6	5	4	3	2	1	0
	·	·	·		OF	DF	IF	TF	SF	ZF		AF		PF		CF
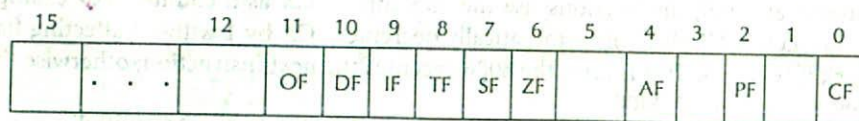
FIGURE 3.3　8086 flag register.

after execution of each instruction. The user can write a service routine at the interrupt address vector to display the desired registers and memory locations. The user can thus debug a program.

3.3 8086 Addressing Modes

The 8086 has 8 basic addressing modes (register, immediate, and memory). The basic and other 8086 addressing modes can be classified into five groups:

1. Addressing modes for accessing immediate and register data (register and immediate modes)
2. Addressing modes for accessing data in memory (memory modes)
3. Addressing modes for accessing I/O ports (I/O modes)
4. Relative addressing mode
5. Implied addressing mode

The assembler directives for the Microsoft 8086 asembler written for IBM PC will be used to illustrate the above 8086 addressing modes.

3.3.1 Addressing Modes for Accessing Immediate and Register Data (Register and Immediate Modes)

3.3.1.a Register Addressing Mode

This mode specifies the source operand, destination operand, or both to be contained in an 8086 register. An example is MOV DX, CX which moves the 16-bit contents of CX into DX. Note that in the above both source and destination operands are in register mode. Another example is MOV CL, DL which moves 8-bit contents of DL into CL. MOV BX, CH is an illegal instruction; the register sizes must be the same.

3.3.1.b Immediate Addressing Mode

In immediate mode, 8- or 16-bit data can be specified as part of the instruction. For example, MOV CL, 03H moves the 8-bit data 03H into CL. On the other hand, MOV DX, 0502H moves the 16-bit data 0502H into DX. Note that in both of the above MOV instructions, the source operand is in immediate mode and the destination operand is in register mode.

A constant such as "VALUE" can be defined by the assembler EQUATE directive such as VALUE EQU 35H. An 8086 instruction with immediate mode such as MOV BH, VALUE can then be used to load 35H into BH. Note that even though the immediate mode specifies data with the instruction, these immediate data must be part of the program located in the code segment. That is, the memory must be addressed by the 8086 CS and IP registers. This is because these data are considered part of the instruction.

3.3.2 Addressing Modes for Accessing Data in Memory (Memory Modes)

The 8086 must use a segment register whenever it accesses the memory. Also, every memory addressing instruction uses an Intel-defined standard default segment register which is DS in most cases. However, a segment override prefix can be placed before most of the memory addressing instructions whose default segment register is to be overridden. For example, INC BYTE PTR [START] will increment the 8-bit content of memory location in DS with offset START by one. However, segment DS can be overridden by ES as follows: INC ES: BYTE PTR [START]; segments cannot be overridden for some cases such as stack reference instructions (such as PUSH and POP). There are six modes in this category. These are

a. Direct addressing mode
b. Register indirect addressing mode
c. Based addressing mode
d. Indexed addressing mode
e. Based indexed addressing mode
f. String addressing mode

3.3.2.a Direct Addressing Mode

In this mode, the 16-bit effective address (EA) is directly included with the instruction. As an example, consider MOV CX; DS:START. This instruction moves the 16-bit contents of a 20-bit physical memory location computed from START and DS into CX.

Note that in the above instruction, the source is in direct addressing mode. If the 16-bit value assigned to the offset START by the programmer using an assembler directive such as EQU is 0040_{16} and [DS] = 3050_{16}, then the BIU generates the 20-bit physical address 30540_{16} on the 8086 address pins and then initiates a memory read cycle to read the 16-bit data from memory location starting at 30540_{16} location. The memory logic places the 16-bit contents of locations 30540_{16} and 30541_{16} on the 8086 data pins. The BIU transfers these data to the EU; the EU then moves these data to CX; [30540_{16}] to CL and [30541_{16}] to CH.

3.3.2.b Register Indirect Addressing Mode

In this mode, the EA is specified in either a pointer register or an index register. The pointer register can be either base register BX or base pointer register BP and index register can be either the Source Index (SI) register or the Destination Index (DI) register. The 20-bit physical address is computed using DS and EA. For example, consider MOV [DI], BX. The destination operand of the above instruction is in register indirect mode, while the source operand is in register mode. The instruction moves the 16-bit content of BX into a memory location offset by the value of EA specified in DI from the current contents in DS*16. Now, if [DS] = 5004_{16}, [DI] = 0020_{16}, and [BX] = 2456_{16}, then after MOV [DI], BX, contents of BL (56_{16}) and BH (24_{16}) are moved to memory locations 50060_{16} and 50061_{16}, respectively.

Using this mode, one instruction can operate on many different memory locations if the value in the base or index register is updated.

3.3.2.c Based Addressing Mode

In this mode, EA is obtained by adding a displacement (signed 8-bit or unsigned 16-bit) value to the contents of BX or BP. The segment registers used are DS and SS. When memory is accessed using BX, the 20-bit physical address is computed from BX and DS. On the other hand, when the user stack is accessed using BP, the 20-bit physical address is computed from BP and SS. This allows the programmer to access the stack without changing the SP contents. Note that SP is called the system stack pointer since some 8086 stack instructions such as PUSH and POP automatically use SP as the stack pointer while BP is called the user stack pointer. As an example of this mode, consider MOV AL, START [BX]. Note that some assemblers use MOV AL, [BX + START] rather than MOV AL, START [BX]. The source operand in the above instruction is in based mode. EA is obtained by adding the value of START and [BX]. The 20-bit physical address is produced from DS and EA. The 8-bit content of this memory location is moved to AL. The displacement START can be either unsigned 16-bit or signed 8-bit. However, a byte is saved for the machine code representation of the instruction if 8-bit displacement is used. The 8086 sign-extends the 8-bit displacement and then adds it to [BX] in the above MOV instruction for determining EA. On the other hand, the 8086 adds an unsigned 16-bit displacement directly with [BX] for determining EA.

Based addressing provides a convenient way to address a structure which may be stored at different places in memory:
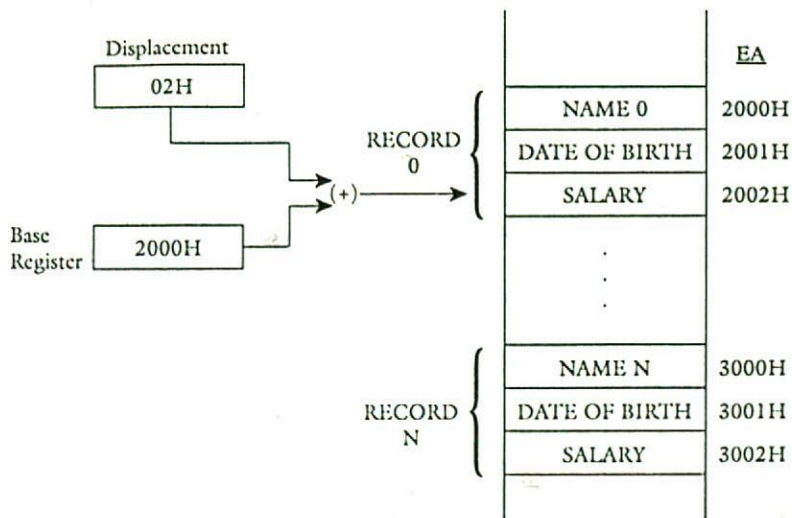
For example, the element salary in record 0 of the employee NAME 0 can be loaded into an 8086 internal register such as AL using the instruction MOV AL, ALPHA [BX], where ALPHA is the 8-bit displacement 02H and BX contains the starting address of the record 0. Now, in order to access the salary of RECORD N, the programmer simply changes the contents of the base register to 3000H.

If BP is specified as a base register in an instruction, the 8086 automatically obtains the operand from the current SS (unless a segment override prefix is present). This makes based addressing with BP a very convenient way to access stack data. BP can be used as a stack pointer in SS to access local variables. Consider the following instruction sequence:

```
PUSH BP             ;  Save BP
MOV  BP, SP         ;  Establish BP
PUSH DX             ;  Save
PUSH AX             ;  registers
SUB  SP, 4          ;  Allocate 2 words of
  :                 ;  stack for accessing stack
MOV  [BP - 6], AX   ;  Arbitrary instructions for
MOV  [BP - 8], BX   ;  accessing stack data using BP
ADD  SP, 4          ;  Deallocate storage
POP  AX             ;  Restore
POP  DX             ;  all registers
POP  BP             ;  that were pushed before
```

This instruction sequence is arbitrarily chosen to illustrate the use of BP for accessing the stack.

Figure 3.4 shows the 8086 stack during various stages. Figure 3.4A shows the stack before execution of the instruction sequence.

The instruction sequence from PUSH BP to SUB SP, 4 pushes BP, DX, and AX and then subtracts 4 from SP, and this allocates 2 words of the stack. The stack at this point is shown in Figure 3.4B. Note that in 8086, SP is decremented by 2 for PUSH and incremented by 2 for POP. The [BP] is not affected by PUSH or POP. The instruction sequence MOV [BP – 6], AX saves AX in the stack location addressed by [BP – 6] in SS. The instruction MOV [BP – 8], BX writes the [BX] into the stack location [BP – 8] in SS. These instructions are arbitrarily chosen to illustrate how BP can be used to access the stack. These two local variables can be accessed by the subroutine, using BP. The instruction ADD SP, 4 releases two words of the allocated
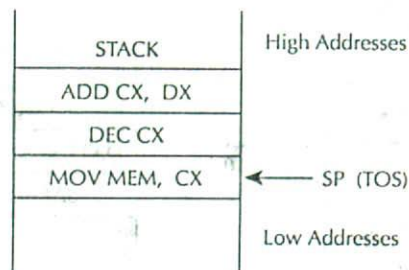
A) Stack before executing the instruction sequence.

(Assumed)
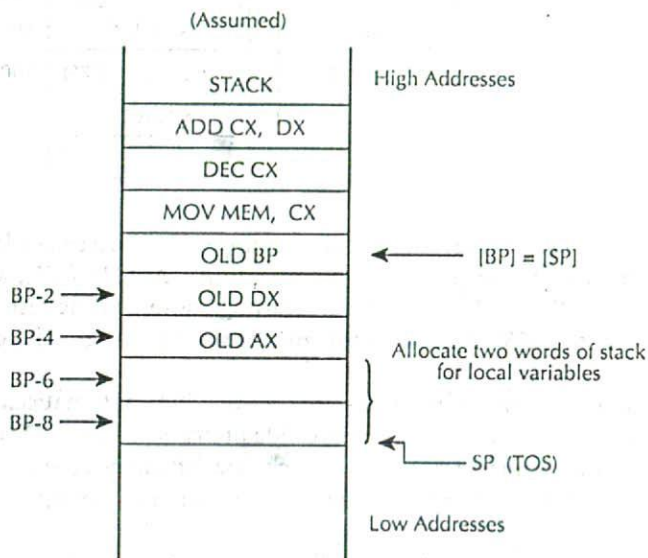
B) Stack after execution of the instruction sequence
from PUSH BP to SUB SP, 4

FIGURE 3.4   Accessing stack using BP.

stack. The stack at this point is shown in Figure 3.4C. The last three POP instructions restore the contents of AX, DX, and BP to their original values and return the stack as it was before the instruction sequence was executed. This is shown in Figure 3.4D.

3.3.2.d Indexed Addressing Mode

In this mode, the effective address is calculated by adding the unsigned 16-bit or sign-extended 8-bit displacement and the contents of SI or DI.

As an example, MOV BH, ARRAY[SI] moves the contents of the 20-bit address computed from the displacement ARRAY, SI and DS into BH. The 8-bit displacement is provided by the programmer using an assembler directive such as EQU. For 16-bit displacement, the EU adds this to SI to determine EA. On the other hand, for 8-bit displacement the EU sign-extends it to 16 bits and then adds to SI for determining EA.

The Indexed addressing mode can be used to access a single table. The displacement can be the starting address of the table. The content of SI or DI can then be used as an index from the starting address to access a particular element in the table.

STACK
ADD CX, DX
DEC CX
MOV MEM, CX
OLD BP
OLD DX
OLD AX

←———— SP (TOS)

Low Addresses
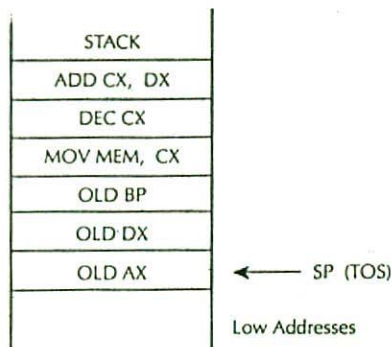
C) Stack after execution of ADD SP, 4

STACK
ADD CX, DX
DEC CX
MOV MEM, CX

High Addresses

←———— SP (TOS)

Low Addresses
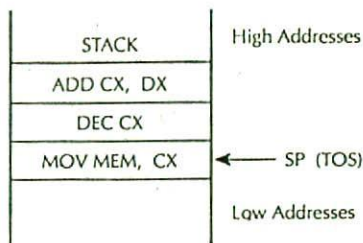
D) Stack before execution of POP BP.

FIGURE 3.4   (continued).

3.3.2.e Based Indexed Addressing Mode

In this mode, the EA is computed by adding a base register (BX or BP), an index register (SI or DI), and a displacement (unsigned 16-bit or sign-extended 8-bit). As an example, consider MOV ALPHA [SI] [BX], CL if [BX] = 0200H, value of ALPHA = 08H, [SI] = 1000H, and [DS] = 3000 H, then 8-bit content of CL is moved to 20-bit physical address 31208_{16}.

Based indexed addressing mode provides a convenient way for a subroutine to address an array allocated on a stack. Register BP can be loaded with the offset in segment SS (top of the stack after the subroutine has saved registers and allocated local storage). The displacement can be the value which is the difference between the top of the stack and the beginning of the array. An index register can then be used to access individual array elements as shown in Figure 3.5.

In the following, [BP] = top of the stack = 2005H; displacement = difference between the top of the stack and start of the array = 04H; [SI or DI] = N = 16-bit number (0, 2, 4, 6 in the example). As an example, the instruction MOV DX, 4 [SI] [BP] with [SI] = 6 will read the array (3) which is the content of 200FH in SS into DX. Since in the based indexed mode, the contents of two registers such as BX and SI can be varied, two-dimensional arrays such as matrices can also be accessed.

3.3.2.f String Addressing Mode

This mode uses index registers. The string instructions automatically assume SI points to the first byte or word of the source operand and DI points to the first byte or word of the destination operand. The contents of SI and DI are automatically incremented (by clearing DF to 0 by CLD instruction) or decremented (by setting DF to 1 by STD instruction) to point to the next byte or word. The segment register for the source is DS and may be overridden.
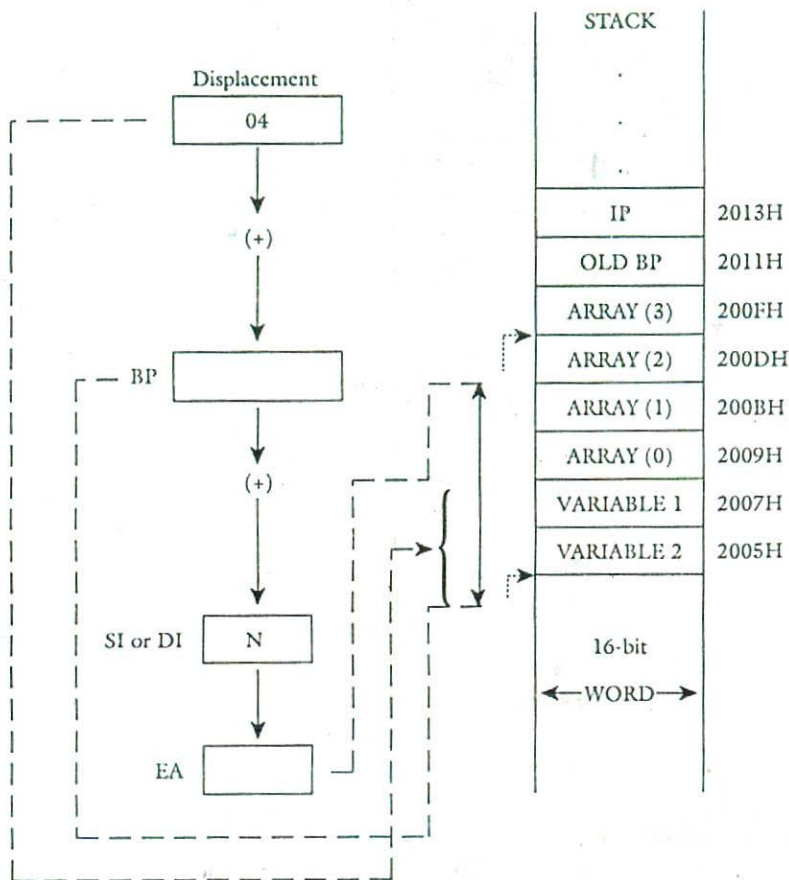
FIGURE 3.5   An example of Based Indexed Addressing.

The segment register for the destination must be ES and cannot be overridden. As an example, consider MOVS BYTE. If $[DF] = 0$, $[DS] = 2000_{16}$, $[SI] = 0500_{16}$, $[ES] = 4000_{16}$, $[DI] = 0300_{16}$, $[20500]_{16} = 38_{16}$, and $[40300_{16}] = 45_{16}$, then after execution of the MOVS BYTE, $[40300_{16}] = 38_{16}$, $[SI] = 0501_{16}$, and $[DI] = 0301_{16}$. The contents of other registers and memory locations are unchanged. Note that SI and DI can be used in either the source or destination operand of a two-operand instruction, except for string instructions in which SI points to the source (may be overridden) and DI must point to the destination.

3.3.3 Addressing Modes for Accessing I/O Ports (I/O Modes)

Standard I/O uses port addressing modes. For memory-mapped I/O, memory addressing modes are used. There are two types of port addressing modes: direct and indirect.

In direct port mode, the port number is an 8-bit immediate operand. This allows fixed access to ports numbered 0 to 255. For example, OUT 05H, AL outputs [AL] to 8-bit port 05H. In indirect port mode, the port number is taken from DX allowing 64K 8-bit ports or 32K 16-bit ports. For example, if $[DX] = 5040_{16}$, then IN AL, DX inputs the 8-bit content of port 5040_{16} into AL. On the other hand, IN AX, DX inputs the 8-bit contents of ports 5040_{16} and 5041_{16} into AL and AH, respectively. Note that 8-bit and 16-bit I/O transfers must take place via AL and AX, respectively.

3.3.4 Relative Addressing Mode

Instructions using this mode specify the operand as a signed 8-bit displacement relative to PC. An example is JNC START. This instruction means that if carry = 0, then PC is loaded with current PC contents (next instruction address) plus the 8-bit signed value of START; otherwise the next instruction is executed. Relative mode with signed 8-bit displacement provides a range of -128_{10} to $+127_{10}$ (0 being positive). If branching beyond this range is necessary, one must use the unconditional 8086 JUMP instruction which uses direct mode.

3.3.5 Implied Addressing Mode

Instructions using this mode have no operands. An example is CLC which clears the carry flag to zero.

3.4 8086 Instruction Set

The 8086 instruction set includes equivalents of the 8085 instructions plus many new ones. The new instructions contain operations such as signed and unsigned multiplication and division, bit manipulation instructions, string instructions, and interrupt instructions.

The 8086 has approximately 117 different instructions with about 300 op codes. The 8086 instruction set contains no operand, single operand, and two operand instructions. Except for string instructions which involve array operations, the 8086 instructions do not permit memory-to-memory operations. Table 3.1 lists a summary of 8086 instructions in alphabetical order. Tables supplied in Appendix E provide a detailed description of the 8086 instructions including the instruction execution times.

Examples of some of the 8086 instructions are given in the following.

1. *Data transfer instructions*

   - MOV CX, DX copies the 16-bit content of DX into CX. MOV AX, 0205H moves immediate data 0205H into 16-bit register AX. MOV CH, [BX] moves the 8-bit content of memory location addressed by BX and segment register DS into CH. If [BX] = 0050H, [DS] = 2000H, [20050H] = 08H, then after MOV CH, [BX], the content of CH will be 08H.

     MOV START [BP], CX moves the 16-bit (CL to first location and then CH) content of CX into two memory locations addressed by the sum of the displacement START and BP, and segment register SS. For example, if [CX] = 5009H, [BP] = 0030H, [SS] = 3000H, START = 06H, then after MOV START [BP], CX physical memory location [30036H] = 09H and [30037H] = 50H. Note that the segment register SS can be overridden by CS using MOV CS: START [BP], CX.

   - PUSH START [BX] pushes the 16-bit contents of two memory locations starting at the 20-bit physical address computed from START, BX, and DS after decrementing SP by 2.

   - POP ES pops the top stack word into ES and then increments SP by 2.

   - XCHG START [BX], AX exchanges the 16-bit word in AX with the contents of two consecutive memory locations starting at 20-bit physical address computed from START, BX, and DS. [AL] is exchanged with the content of the first location and [AH] is exchanged with the content of the next location.

   - XLAT can be used to convert a code such as ASCII into another code such as EBCDIC. This instruction is equivalent to MOV AL, [AL][BX].

Table 3.1   Summary of 8086 Instructions

Instructions	Interpretation	Comments
AAA	ASCII adjust [AL] after addition	This instruction has implied addressing mode; this instruction is used to adjust the content of AL after addition of two ASCII characters
AAD	ASCII adjust for division	This instruction has implied addressing mode; converts two unpacked BCD digits in AX into equivalent binary numbers in AL; AAD must be used before dividing two unpacked BCD digits by an unpacked BCD byte
AAM	ASCII adjust after multiplication	This instruction has implied addressing mode; after multiplying two unpacked BCD numbers, adjust the product in AX to become an unpacked BCD result; ZF, SF, and PF are affected
AAS	ASCII adjust [AL] after subtraction	This instruction has implied addressing mode used to adjust [AL] after subtraction of two ASCII characters
ADC mem/reg 1, mem/reg 2	[mem/reg 1] ← [mem/reg 1] + [mem/reg 2] + CY	Memory or register can be 8- or 16-bit; all flags are affected; no segment registers are allowed; no memory-to-memory ADC is permitted
ADC mem, data	[mem] ← [mem] + data + CY	Data can be 8- or 16-bit; mem uses DS as the segment register; all flags are affected
ADC reg, data	[reg] ← [reg] + data + CY	Data can be 8- or 16-bit; register cannot be segment register; all flags are affected
ADD mem/reg 2, mem/reg 1	[mem/reg 1] ← [mem/reg 2] + [mem/reg 1]	Add two 8- or 16-bit data; no memory-to-memory ADD is permitted; all flags are affected; mem uses DS as the segment register; reg 1 or reg 2 cannot be segment register
ADD mem, data	[mem] ← [mem] + data	Mem uses DS as the segment register; data can be 8- or 16-bit; all flags are affected
ADD reg, data	[reg] ← [reg] + data	Data can be 8- or 16-bit; no segment registers are allowed; all flags are affected
AND mem/reg 1, mem/reg 2	[mem/reg 1] ← [mem/reg 1] ∧ [mem/reg 2]	This instruction logically ANDs 8- or 16-bit data in [mem/reg 1] with 8- or 16-bit data in [mem/reg 2]; all flags are affected; OF and CF are cleared to zero; no segment registers are allowed; no memory-to-memory operation is allowed; mem uses DS as the segment register
AND mem, data	[mem] ← [mem] ∧ data	Data can be 8- or 16-bit; mem uses DS as the segment register; all flags are affected with OF and CF always cleared to zero
AND reg, data	[reg] ← [reg] ∧ data	Data can be 8- or 16-bit; reg cannot be segment register; all flags are affected with OF and CF cleared to zero
CALL PROC (NEAR)	Call a subroutine in the same segment with signed 16-bit displacement (to CALL a subroutine in ±32K)	NEAR in the statement BEGIN PROC NEAR indicates that the subroutine 'BEGIN' is in the same segment and BEGIN is 16-bit signed; CALL BEGIN instruction decrements SP by 2 and then pushes IP onto the stack and then adds the signed 16-bit value of BEGIN to IP and CS is unchanged; thus, a subroutine is called in the same segment (intrasegment direct)
CALL reg 16	CALL a subroutine in the same segment addressed by the contents of a 16-bit general register	The 8086 decrements SP by 2 and then pushes IP onto the stack, then specified 16-bit register contents (such as BX, SI, and DI) provide the new value for IP; CS is unchanged (intrasegment indirect)

Table 3.1   Summary of 8086 Instructions (*continued*)

Instructions	Interpretation	Comments
CALL mem 16	CALL a subroutine addressed by the content of a memory location pointed to by 8086 16-bit register such as BX, SI, and DI	The 8086 decrements SP by 2 and pushes IP onto the stack; the 8086 then loads the contents of a memory location addressed by the content of a 16-bit register such as BX, SI, and DI into IP; [CS] is unchanged (intrasegment indirect)
CALL PROC (FAR)	CALL a subroutine in another segment	FAR in the statement BEGIN PROC FAR indicates that the subroutine 'BEGIN' is in another segment and the value of BEGIN is 32 bit wide
		The 8086 decrements SP by 2 and pushes CS onto the stack and moves the low 16-bit value of the specified 32-bit number such as 'BEGIN' in CALL BEGIN into CS; SP is again decremented by 2; IP is pushed onto the stack; IP is then loaded with high 16-bit value of BEGIN; thus, this instruction CALLS a subroutine in another code segment (intersegment direct)
CALL DWORDPTR [reg 16]	CALL a subroutine in another segment	This instruction decrements SP by 2, and pushes CS onto the stack; CS is then loaded with the contents of memory locations addressed by [reg 16 + 2] and [reg 16 + 3] in DS; the SP is again decremented by 2; IP is pushed onto the stack; IP is then loaded with the contents of memory locations addressed by [reg 16] and [reg 16 + 1] in DS; typical 8086 registers used for reg 16 are BX, SI, and DI (intersegment indirect)
CBW	Convert a byte to a word	Extend the sign bit (bit 7) of AL register into AH
CLC	CF ← 0	Clear carry to zero
CLD	DF ← 0	Clear direction flag to zero
CLI	IF ← 0	Clear interrupt enable flag to zero to disable maskable interrupts
CMC	CF ← CF'	One's complement carry
CMP mem/reg 1, mem/reg 2	[mem/reg 1] − [mem/reg 2], flags are affected	mem/reg can be 8- or 16-bit; no memory-to-memory comparison allowed; result of subtraction is not provided; all flags are affected
CMP mem/reg, data	[mem/reg] − data, flags are affected	Subtracts 8- or 16-bit data from [mem or reg] and affects flags; no result is provided
CMPS BYTE or CMPSB	FOR BYTE [[SI]] − [[DI]], flags are affected [SI] ← [SI] ± 1 [DI] ← [DI] ± 1	8- or 16-bit data addressed by [DI] in ES is subtracted from 8- or 16-bit data addressed by SI in DS and flags are affected without providing any result; if DF = 0, then SI and DI are incremented by one for byte and two for word; if DF = 1, then SI and DI are decremented by one for byte and two for word; the segment register ES in destination cannot be overridden
CMPS WORD or CPSW	FOR WORD [[SI]] − [[DI]], flags are affected [SI] ← [SI] + 2 [DI] ← [DI] + 2	
CWD	Convert a word to 32 bits	Extend the sign bit of AX (bit 15) into DX
DAA	Decimal adjust [AL] after addition	This instruction uses implied addressing mode; this instruction converts [AL] into BCD; DAA should be used after BCD addition
DAS	Decimal adjust [AL] after subtraction	This instruction uses implied addressing mode; converts [AL] into BCD; DAS should be used after BCD subtraction
DEC reg 16	[reg 16] ← [reg 16] − 1	This is a one-byte instruction; used to decrement a 16-bit register except segment register; does not affect the carry flag
DEC mem/reg 8	[mem] ← [mem] − 1 or [reg 8] ← [reg 8] − 1	Used to decrement a byte or a word in memory or an 8-bit register content; segment register cannot be decremented by this instruction; does not affect carry flag

Table 3.1   Summary of 8086 Instructions (*continued*)

Instructions	Interpretation	Comments
DIV mem/reg	16/8 bit divide: $\dfrac{[AX]}{[mem\ 8\ /\ reg\ 8]}$ [AH] ← Remainder [AL] ← Quotient 32/16 bit divide: $\dfrac{[DX]\ [AX]}{[mem\ 16\ /\ reg\ 16]}$ [DX] ← Remainder [AX] ← Quotient	Mem/reg is 8-bit for 16-bit by 8-bit divide and 16-bit for 32-bit by 16-bit divide; this is an unsigned division; no flags are affected; division by zero automatically generates an internal interrupt
ESC external OP code, source	ESCAPE to external processes	This instruction is used to pass instructions to a coprocessor such as the 8087 floating point coprocessor which simultaneously monitors the system bus with the 8086; the coprocessor OP codes are 6-bit wide; the coprocessor treats normal 8086 instructions as NOP's; the 8086 fetches all instructions from memory; when the 8086 encounters an ESC instruction, it usually treats it as NOP; the coprocessor decodes this instruction and carries out the operation using the 6-bit OP code independent of the 8086; for ESC OP code, memory, the 8086 accesses data in memory for the coprocessor; for ESC data, register, the coprocessor operates on 8086 registers; the 8086 treats this as an NOP
HLT	HALT	Halt
IDIV mem/reg	Same as DIV mem/reg	Signed division. No flags are affected.
IMUL mem/reg	For 8 × 8 [AX] ← [AL] * [mem 8/reg 8] FOR 16 × 16 [DX] [AX] ← [AX]* [mem 16/reg 16]	Mem/reg can be 8- or 16-bit; only CF and OF are affected; signed multiplication
IN AL, DX	[AL] ← [PORT DX]	Input AL with the 8-bit content of a port addressed by DX; this is a one-byte instruction
IN AX, DX	[AX] ← [PORT DX]	Input AX with the 16-bit content of a port addressed by DX and DX + 1; this is a one-byte instruction
IN AL, PORT	[AL] ← [PORT]	Input AL with the 8-bit content of a port addressed by the second byte of the instruction
IN AX, PORT	[AX] ← [PORT]	Input AX with the 16-bit content of a port addressed by the 8-bit address in the second byte of the instruction
INC reg 16	[reg 16] ← [reg 16] + 1	This is a one-byte instruction; used to increment a 16-bit register except the segment register; does not affect the carry flag
INC mem/reg 8	[mem] ← [mem] + 1 or [reg 8] ← [reg 8] + 1	This is a two-byte instruction; can be used to increment a byte or word in memory or an 8-bit register content; segment registers cannot be incremented by this instruction; does not affect the carry flag
INT n (n can be zero thru 255)	[SP] ← [SP] – 2 [[SP]] ← Flags IF ← 0 TF ← 0 [SP] ← [SP] – 2 [[SP]] ← [CS] [CS] ← 4n + 2 [SP] ← [SP] –2	Software interrupts can be used as supervisor calls; that is, request for service from an operating system; a different interrupt type can be used for each type of service that the operating system could supply for an application or program; software interrupt instructions can also be used for checking interrupt service routines written for hardware-initiated interrupts

Table 3.1 Summary of 8086 Instructions (*continued*)

Instructions	Interpretation	Comments
INTO	$[[SP]] \leftarrow [IP]$ $[IP] \leftarrow 4n$ Interrupt on Overflow	Generates an internal interrupt if OF = 1; executes INT 4; can be used after an arithmetic operation to activate a service routine if OF = 1; when INTO is executed and if OF = 1, operations similar to INT n take place
IRET	Interrupt Return	POPS IP, CS and Flags from stack; IRET is used as return instruction at the end of a service routine for both hardware and software interrupts
JA/JNBE disp 8	Jump if above/jump if not below or equal	Jump if above/jump if not below or equal with 8-bit signed displacement; that is, the displacement can be from -128_{10} to $+127_{10}$, zero being positive; JA and JNBE are the mnemonic which represent the same instruction; Jump if both CF and ZF are zero; used for unsigned comparison
JAE/JNB/JNC disp 8	Jump if above or equal/jump if not below/jump if no carry	Same as JA/JNBE except that the 8086 Jumps if CF = 0; used for unsigned comparison
JB/JC/JNAE disp 8	Jump if below/jump if carry/jump if not above or equal	Same as JA/JNBE except that the jump is taken CF = 1, used for unsigned comparison
JBE/JNA disp 8	Jump if below or equal/jump if not above	Same as JA/JNBE except that the jump is taken if CF = 1 or ZF = 0; used for unsigned comparison
JCXZ disp 8	Jump if CX = 0	Jump if CX = 0; this instruction is useful at the beginning of a loop to bypass the loop if CX = 0
JE/JZ disp 8	Jump if equal/jump if zero	Same as JA/JNBE except that the jump is taken if ZF = 1; used for both signed and unsigned comparison
JG/JNLE disp 8	Jump if greater/jump if not less or equal	Same as JA/JNBE except that the jump is taken if $((SF \oplus OF)$ or ZF) = 0; used for signed comparison
JGE/JNL disp 8	Jump if greater or equal/ jump if not less	Same as JA/JNBE except that the jump is taken if $(SF \oplus OF)$ = 0; used for signed comparison
JL/JNGE disp 8	Jump if less/Jump if not greater nor equal	Same as JA/JNBE except that the jump is taken if $(SF \oplus OF)$ = 1; used for signed comparison
JLE/JNG disp 8	Jump if less or equal/ jump if not greater	Same as JA/JNBE except that the jump is taken if $((SF \oplus OF)$ or ZF) = 1; used for signed comparison
JMP Label	Unconditional Jump with a signed 8-bit (SHORT) or signed 16-bit (NEAR) displacement in the same segment	The label START can be signed 8-bit (called SHORT jump) or signed 16-bit (called NEAR jump) displacement; the assembler usually determines the displacement value; if the assembler finds the displacement value to be signed 8-bit (−128 to +127, 0 being positive), then the assembler uses two bytes for the instruction: one byte for the OP code followed by a byte for the displacement; the assembler sign extends the 8-bit displacement and then adds it to IP; [CS] is unchanged; on the other hand, if the assembler finds the displacement to be signed 16-bit (±32 K), then the assembler uses three bytes for the instruction: one byte for the OP code followed by 2 bytes for the displacement; the assembler adds the signed 16-bit displacement to IP; [CS] is unchanged; therefore, this JMP provides a jump in the same segment (intrasegment direct jump)
JMP reg 16	$[IP] \leftarrow [reg\ 16]$ [CS] is unchanged	Jump to an address specified by the contents of a 16-bit register such as BX, SI, and DI in the same code segment; in the example JMP BX, [BX] is loaded into IP and [CS] is unchanged (intrasegment memory indirect jump)
JMP mem 16	$[IP] \leftarrow [mem]$ [CS] is unchanged	Jump to an address specified by the contents of a 16-bit memory location addressed by 16-bit register

Table 3.1   Summary of 8086 Instructions (*continued*)

Instructions	Interpretation	Comments
		such as BX, SI, and DI; in the example, JMP [BX] copies the content of a memory location addressed by BX in DS into IP; CS is unchanged (intrasegment memory indirect jump)
JMP Label (FAR)	Unconditionally jump to another segment	This is a 5-byte instruction: the first byte is the OP code followed by four bytes of 32-bit immediate data; bytes 2 and 3 are loaded into IP; bytes 4 and 5 are loaded into CS to JUMP unconditionally to another segment (intersegment direct)
JMP DWORDPTR [reg 16]	Unconditionally jump to another segment	This instruction loads the contents of memory locations addressed by [reg 16] and [reg 16 + 1] in DS into IP; it then loads the contents of memory locations addressed by [reg 16 + 2] and [reg 16 + 3] in DS into CS; typical 8086 registers used for reg 16 are BX, SI, and DI (intersegment indirect)
JNE/JNZ disp 8	Jump if not equal/jump if not zero	Same as JA/JNBE except that the jump is taken if ZF = 0; used for both signed and unsigned comparison
JNO disp 8	Jump if not overflow	Same as JA/JNBE except that the jump is taken if OF = 0
JNP/JPO disp 8	Jump if no parity/jump if parity odd	Same as JA/JNBE except that the jump is taken if PF = 0
JNS disp 8	Jump if not sign	Same as JA/JNBE except that the jump is taken if SF = 0
JO disp 8	Jump if overflow	Same as JA/JNBE except that the jump is taken if OF = 1
JP/JPE disp 8	Jump if parity/jump if parity even	Same as JA/JNBE except that the jump is taken if PF = 1
JS disp 8	Jump if sign	Same as JA/JNBE except that the jump is taken if SF = 1
LAHF	[AH] ← Flag low-byte	This instruction has implied addressing mode; it loads AH with the low byte of the flag register; no flags are affected
LDS reg, mem	[reg] ← [mem] [DS] ← [mem + 2]	Load a 16-bit register (AX, BX, CX, DX, SP, BP, SI, DI) with the content of specified memory and load DS with the content of the location that follows; no flags are affected; DS is used as the segment register for mem
LEA reg, mem	[reg] ← [offset portion of address]	LEA (load effective address) loads the value of the source operand rather than its content to register (such as SI, DI, BX) which are allowed to contain offset for accessing memory; no flags are affected
LES reg, mem	[reg] ← [mem] [ES] ← [mem + 2]	DS is used as the segment register for mem; in the example LES DX, [BX], DX is loaded with 16-bit value from a memory location addressed by 20-bit physical address computed from DS and BX; the 16-bit content of the next memory is loaded into ES; no flags are affected
LOCK	LOCK bus during next instruction	Lock is a one-byte prefix that causes the 8086 (configured in maximum mode) to assert its bus LOCK signal while following instruction is executed; this signal is used in multiprocessing; the LOCK pin of the 8086 can be used to LOCK other processors off the system bus during execution of an instruction; in this way, the 8086 can be assured of uninterrupted access to common system resources such as shared RAM

Table 3.1   Summary of 8086 Instructions (*continued*)

Instructions	Interpretation	Comments
LODS BYTE or LODSB	FOR BYTE [AL] ← [[SI]] [SI] ← [SI] ± 1	Load 8-bit data into AL or 16-bit data into AX from a memory location addressed by SI in segment DS; if DF = 0, then SI is incremented by 1 for byte or
LODS WORD or LODSW	FOR WORD [AX] ← [[SI]] [SI] ← [SI] ± 2	incremented by 2 for word after the load; if DF = 1, then SI is decremented by 1 for byte or decremented by 2 for word; LODS affects no flags
LOOP disp 8	Loop if CX not equal to zero	Decrement CX by one, without affecting flags and loop with signed 8-bit displacement (from −128 to +127, zero being positive) if CX is not equal to zero
LOOPE/LOOPZ disp 8	Loop while equal/loop while zero	Decrement CX by one without affecting flags and loop with signed 8-bit displacement if CX is equal to zero, and if ZF = 1 which results from execution of the previous instruction
LOOPNE/LOOPNZ disp 8	Loop while not equal/loop while not zero	Decrement CX by one without affecting flags and loop with signed 8-bit displacement if CX is not equal to zero and ZF = 0 which results from execution of previous instruction
MOV mem/reg 2, mem/reg 1	[mem/reg 2] ← [mem/reg 1]	mem uses DS as the segment register; no memory-to-memory operation allowed; that is, MOV mem, mem is not permitted; segment register cannot be specified as reg; no flags are affected
MOV mem, data	[mem] ← data	mem uses DS as the segment register; 8- or 16-bit data specifies whether memory location is 8- or 16-bit; no flags are affected .
MOV reg, data	[reg] ← data	Segment register cannot be specified as reg; data can be 8- or 16-bit; no flags are affected
MOV segreg, mem/reg	[segreg] ← [mem/reg]	mem uses DS as segment register; used for initializing CS, DS, ES, and SS; no flags are affected
MOV mem/reg, segreg	[mem/reg] ← [segreg]	mem uses DS as segment register; no flags are affected .
MOVS BYTE or MOVSB	FOR BYTE [[DI]] ← [[SI]] [SI] ← [SI] ± 1	Move 8-bit or 16-bit data from the memory location addressed by SI in segment DS location addressed by DI in ES; segment DS can be overridden by a
MOVS WORD or MOVSW	FOR WORD [[DI]] ← [[SI]] [SI] ← [SI] ± 2	prefix but destination segment must be ES and cannot be overridden; if DF = 0, then SI is incremented by one for byte or incremented by two for word; if DF = 1, then SI is decremented by one for byte or by two for word
MUL mem/reg	FOR 8 × 8 [AX] ← [AL]* [mem/reg] FOR 16 × 16 [DX] [AX] ← [AX]* [mem/reg]	mem/reg can be 8- or 16-bit; only CF and OF are affected; unsigned multiplication
NEG mem/reg	[mem/reg] ← [mem/reg]' + 1	mem/reg can be 8- or 16-bit; performs two's complement subtraction of the specified operand from zero, that is, two's complement of a number is formed; all flags are affected except CF = 0 if [mem/reg] is zero; otherwise CF = 1
NOP	No Operation	8086 does nothing
NOT reg	[reg] ← [reg]'	mem and reg can be 8- or 16-bit; segment registers are not allowed; no flags are affected; ones complement reg
NOT mem	[mem] ← [mem]'	mem uses DS as the segment register; no flags are affected; ones complement mem
OR Mem/reg 1, Mem/reg 2	[mem/reg 1] ← [mem/reg 1] ∨ [mem/reg 2]	No memory-to-memory operation is allowed; [mem] or [reg 1] or [reg 2] can be 8- or 16-bit; all flags are affected with OF and CF cleared to zero; no

Table 3.1   Summary of 8086 Instructions (*continued*)
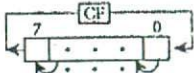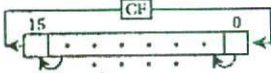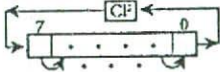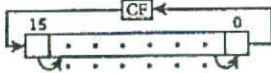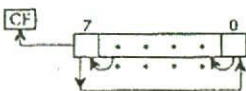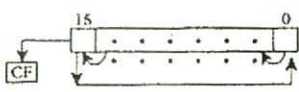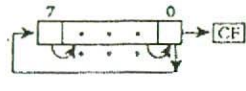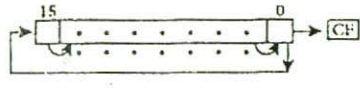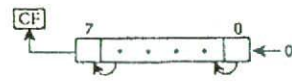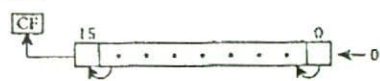
Instructions	Interpretation	Comments
		segment registers are allowed; mem uses DS as segment register
OR mem, data	[mem] ← [mem] ∨data	mem and data can be 8- or 16-bit; mem uses DS as segment register; all flags are affected with CF and OF cleared to zero
OR reg, data	[reg] ← [reg] ∨ data	reg and data can be 8- or 16-bit; no segment registers are allowed; all flags are affected with CF and OF cleared to zero
OUT DX, AL	[PORT] ← [AL] DX	Output the 8-bit contents of AL into an I/O Port addressed by the 16-bit content of DX; this is a one-byte instruction
OUT DX, AX	[PORT] ← [AX] DX	Output the 16-bit contents of AX into an I/O Port addressed by the 16-bit content of DX; this is a one-byte instruction
OUT PORT, AL	[PORT] ← [AL]	Output the 8-bit contents of AL into the Port specified in the second byte of the instruction
OUT PORT, AX	[PORT] ← [AX]	Output the 16-bit contents of AX into the Port specified in the second byte of the instruction
POP mem	[mem] ← [[SP]] [SP] ← [SP] + 2	mem uses DS as the segment register; no flags are affected
POP reg	[reg] ← [[SP]] [SP] ← [SP] + 2	Cannot be used to POP segment registers or flag register
POP segreg	[segreg] ← [[SP]] [SP] ← [SP] + 2	POP CS is illegal
POPF	[Flags] ← [[SP]] [SP] ← [SP] + 2	This instruction pops the top two stack bytes in the 16-bit flag register
PUSH mem	[SP] ← [SP] − 2 [[SP]] ← [mem]	mem uses DS as segment register; no flags are affected; pushes 16-bit memory contents
PUSH reg	[SP] ← [SP] − 2 [[SP]] ← [reg]	reg must be a 16-bit register; cannot be used to PUSH segment register or Flag register
PUSH segreg	[SP] ← [SP] − 2 [[SP]] ← [segreg]	PUSH CS is illegal
PUSHF	[SP] ← [SP] − 2 [[SP]] ← [Flags]	This instruction pushes the 16-bit Flag register onto the stack
RCL mem/reg, 1	ROTATE through carry left once byte or word in mem/reg	FOR BYTE FOR WORD
RCL mem/reg, CL	ROTATE through carry left byte or word in mem/reg by [CL]	Operation same as RCL mem/reg, 1 except the number of rotates is specified in CL for rotates up to 255; zero or negative rotates are illegal
RCR mem/reg, 1	ROTATE through carry right once byte or word in mem/reg	FOR BYTE FOR WORD
RCR mem/reg, CL	ROTATE through carry right byte or word in mem/reg by [CL]	Operation same as RCR mem/reg, 1 except the number of rotates is specified in CL for rotates up to 255; zero or negative rotates are illegal

Table 3.1 Summary of 8086 Instructions (*continued*)

Instructions	Interpretation	Comments
RET	.POPS IP for intrasegment CALLS .POPS IP and CS for intersegment CALLS	The assembler generates an intrasegment return if the programmer has defined the subroutine as NEAR; for intrasegment return, the following operations take place: [IP] ← [[SP]], [SP] ← [SP] + 2; on the other hand, the assembler generates an intersegment return if the subroutine has been defined as FAR; in this case, the following operations take place: [IP] ← [[SP]], [SP] ← [SP] + 2, [CS] ← [[SP]], [SP] ← [SP] + 2; an optional 16-bit displacement 'START' can be specified with the intersegment return such as RET START; in this case, the 16-bit displacement is added to the SP value; this feature may be used to discard parameter pushed onto the stack before the execution of the CALL instruction
ROL mem/reg, 1	ROTATE left once byte or word in mem/reg	FOR BYTE FOR WORD
ROL mem/reg, CL	ROTATE left byte or word by the content of CL	[CL] contains rotate count up to 255; zero and negative shifts are illegal; CL is used to rotate count when the rotate is greater than once; mem uses DS as the segment register
ROR mem/reg, 1	ROTATE right once byte or word in mem/reg	FOR BYTE FOR WORD
ROR mem/reg, CL	ROTATE right byte or word in mem/reg by [CL]	Operation same as ROR mem/reg, 1; [CL] specifies the number of rotates for up to 255; zero and negative rotates are illegal; mem uses DS as the segment register
SAHF	{Flags, low-byte} ← [AH]	This instruction has the implied addressing mode; the content of the AH register is stored into the low-byte of the flag register; no flags are affected
SAL mem/reg, 1	Shift arithmetic left once byte or word in mem or reg	FOR BYTE FOR WORD Mem uses DS as the segment register; reg cannot be segment registers; OF and CF are affected; if sign bit is changed during or after shifting, the OF is set to one

Table 3.1   Summary of 8086 Instructions (*continued*)
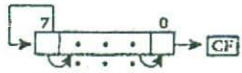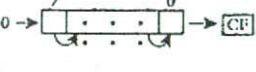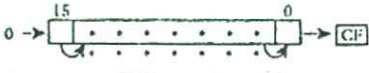
Instructions	Interpretation	Comments
SAL mem/reg, CL	Shift arithmetic left byte or word by shift count on CL	Operation same as SAL mem/reg, 1; CL contains shift count for up to 255; zero and negative shifts are illegal; [CL] is used as shift count when shift is greater than one; OF and SF are affected; if sign bit of [mem] is changed during or after shifting, the OF is set to one; mem uses DS as segment register
SAR mem/reg, 1	SHIFT arithmetic right once byte or word in mem/reg	FOR BYTE FOR WORD
SAR mem/reg, CL	SHIFT arithmetic right byte or word in mem/reg by [CL]	Operation same as SAR mem/reg, 1; however, shift count is specified in CL for shifts up to 255; zero and negative shifts are illegal
SBB mem/reg 1, mem/reg 2	[mem/reg 1] ← [mem/reg 1] – [mem/reg 2] – CY	Same as SUB mem/reg 1, mem/reg 2 except this is a subtraction with borrow
SBB mem, data	[mem] ← [mem] – data – CY	Same as SUB mem, data except this is a subtraction with borrow
SBB reg, data	[reg] ← [reg] – data – CY	Same as SUB reg, data except this is a subtraction with borrow
SBB A, data	[A] ← [A] – data – CY	Same as SUB A, data except this is a subtraction with borrow
SCAS BYTE or SCASB	FOR BYTE [AL] – [[DI]], flags are affected, [DI] [DI] ± 1	8- or 16-bit data addressed by [DI] in ES is subtracted from 8- or 16-bit data in AL or AX and flags are affected without affecting [AL] or [AX] or string data; ES cannot be overridden; if DF = 0, then DI is incremented by one for byte and two for word; if DF = 1, then DI is decremented by one for byte or decremented by two for word
SCAS WORD or SCASW	FOR WORD [AX] – [[DI]], flags are affected, [DI] ← [DI] ± 2	
SHL mem/reg, 1	SHIFT logical left once byte or word in mem/reg	Same as SAL mem/reg, 1
SHL mem/reg, CL	SHIFT logical left byte or word in mem/reg by the shift count in CL	Same as SAL mem/reg, CL except overflow is cleared to zero
SHR mem/reg, 1	SHIFT right logical once byte or word in mem/reg	FOR BYTE FOR WORD
SHR mem/reg, CL	SHIFT right logical byte or word in mem/reg by [CL]	Operation same as SHR mem/reg, 1; however, shift count is specified in CL for shifts up to 255; zero and negative shifts are illegal
STC	CF ← 1	Set carry to one
STD	DF ← 1	Set direction flag to one
STI	IF ← 1	Set interrupt enable flag to one to enable maskable interrupts
STOS BYTE or STOSB	FOR BYTE [[DI]] ← [AL] [DI] ← [DI] ± 1	Store 8-bit data from AL or 16-bit data from AX into a memory location addressed by DI in segment ES; segment register ES cannot be overridden; if DF = 0, then DI is incremented by one for byte or incremented by two for word after the store
STOS WORD or STOSW	FOR WORD [[DI]] ← [AX] [DI] ← [DI] ± 2	

Table 3.1 · Summary of 8086 Instructions (*continued*)

Instructions	Interpretation	Comments
SUB mem/reg 1, mem/reg 2	[mem/reg 1] ← [mem/reg 1] − [mem/reg 2]	No memory-to-memory SUB permitted; all flags are affected; mem uses DS as the segment register
SUB mem, data	[mem] ← [mem] − data	Data can be 8- or 16-bit; mem uses DS as the segment register; all flags are affected
SUB reg, data	[reg] ← [reg] − data	Data can be 8- or 16-bit; all flags are affected
TEST mem/reg 1, mem/reg 2	[mem/reg 1] ∧ [mem/reg 2], no result; flags are affected	No memory-to-memory TEST is allowed; no result is provided; all flags are affected with CF and OF cleared to zero; [mem], [reg 1] or [reg 2] can be 8- or 16-bit; no segment registers are allowed; mem uses DS as the segment register
TEST mem, data	[mem] ∧ data, no result; flags are affected	Mem and data can be 8- or 16-bit; no result is provided; all flags are affected with CF and OF cleared to zero; mem uses DS as the segment register
TEST reg, data	[reg] ∧ data, no result; flags are affected	Reg and data can be 8- or 16-bit; no result is provided; all flags are affected with CF and OF cleared to zero; reg cannot be segment register;
WAIT	8086 enters wait state	Causes CPU to enter wait state if the 8086 TEST pin is high; while in wait state, the 8086 continues to check TEST pin for low; if TEST pin goes back to zero, the 8086 executes the next instruction; this feature can be used to synchronize the operation of 8086 to an event in external hardware
XCHG mem, reg	[mem] ↔ [reg]	reg and mem can be both 8- or 16-bit; mem uses DS as the sement register; reg cannot be segment register; no flags are affected
XCHG reg, reg	[reg] ↔ [reg]	reg can be 8- or 16-bit; reg cannot be segment register; no flags are affected
XLAT	[AL] ← [AL] + [BX]	This instruction is useful for translating characters from one code such as ASCII to another such as EBCDIC; this is a no-operand instruction and is called an instruction with implied addressing mode; the instruction loads AL with the contents of a 20-bit physical address computed from DS, BX, and AL; this instruction can be used to read the elements in a table where BX can be loaded with a 16-bit value to point to the starting address (offset from DS) and AL can be loaded with the element number (0 being the first element number); no flags are affected; the XLAT instruction is equivalent to MOV AL, [AL] [BX]
XOR mem/reg 1, mem/reg 2	[mem/reg 1] ← [mem/reg 1] ⊕ [mem/reg 2]	No memory-to-memory operation is allowed; [mem] or [reg 1] or [reg 2] can be 8- or 16-bit; all flags are affected with CF and OF cleared to zero; mem uses DS as the segment register
XOR mem, data	[reg] ← [mem] ⊕ data	Data and mem can be 8- or 16-bit; mem uses DS as the segment register; mem cannot be segment register; all flags are affected with CF and OF cleared to zero
XOR reg, data	[reg] ← [reg] ⊕ data	Same as XOR mem, data

Suppose that an 8086-based microcomputer is interfaced to an ASCII keyboard and an IBM printer (EBCDIC code). In such a system, any number entered into the microcomputer will be in ASCII code which must be converted to EBCDIC code before outputting to the printer. If the number entered is 4, then the ASCII code for 4 (34H) must be translated to the EBCDIC code for 4 (F4H). A look-up table containing EBCDIC code for all the decimal numbers (0 to 9) can be stored in memory at the 16-bit starting address in the data segment of memory; with code for 0 stored at 3030H, code for 1 at 3031H, and so on. Now, if AL is loaded with the ASCII code 34H and BX is loaded with 3000H, then by using the XLAT instruction, content of memory location 3034H in the data segment containing the EBCDIC code for 4 (F4H) from the look-up table is read into AL, thus replacing the ASCII code for 4 with the EBCDIC code for 4.

2. *Examples of input/output instructions*

- Consider fixed port addressing in which the 8-bit port address is directly specified as part of the instruction. IN AL, 38H inputs 8-bit data from port 38H into AL. IN AX, 38H inputs 16-bit data from ports 38H and 39H into AX. OUT 38H, AL outputs the contents of AL to port 38H. OUT 38H, AX, on the other hand, outputs the 16-bit contents of AX to ports 38H and 39H.

- For the variable port addressing, the port address is 16-bit and is contained in the DX register. Consider ports addressed by 16-bit address contained in DX. Assume $[DX] = 3124_{16}$ in all the following examples:

  IN AL, DX inputs 8-bit data from 8-bit port 3124_{16} into AL.

  IN AX, DX inputs 16-bit data from ports 3124_{16} and 3125_{16} into AL and AH respectively.

  OUT DX, AL outputs 8-bit data from AL into port 3124_{16}.

  OUT DX, AX outputs 16-bit data from AL and AH into ports 3124_{16} and 3125_{16} respectively.

Variable port addressing allows up to 65,536 ports with addresses from 0000H to FFFFH. The port addresses in the variable port addressing can be calculated dynamically in a program. For example, assume that an 8086-based microcomputer is connected to three printers via three separate ports. Now, in order to output to each one of the printers, separate programs are required if fixed port addressing is used. However, with variable port addressing one can write a general subroutine to output to the printers and then supply the address of the port for a particular printer for which data output is desired to register DX in the subroutine.

3. *Examples of address initialization instructions*

- LEA reg, mem loads an offset (mem) directly into the specified register. This instruction is useful when address computation is required. LEA BX, 5000H and MOV BX, 5000H accomplish the same task. That is, both of these instructions load 5000H into BX. On the other hand, LEA DI, [SI] [BX] loads the 16-bit value computed from BX and SI into DI whereas MOV DI, [SI] [BX] loads the 16 bit contents of a memory location computed from SI and BX into DI. LEA can be used to load these addresses. For example, LEA can be used to load the address of the table used by the XLAT instructions.

- LDS reg, mem can be used to initialize SI and DS to point to the start of the source string before using one of the string instructions. For example, LDS SI, [BX] loads the 16-bit contents of memory offset by [BX] in the data segment into SI and the 16-bit contents of memory offset by [BX + 2] in the data segment into DS.

- LES reg, mem can be used to point to the start of the destination string before using one of the string instructions. For example, LES DI, [BX] loads the 16-bit contents of

memory offset by [BX] in the data segment to DI and then initializes ES with the 16-bit contents of memory offset by [BX + 2] in DS.

4. *Examples of flag register instructions*

  • PUSHF pushes the 16-bit flag register onto the stack.
  • LAHF loads AH with the condition codes from the low byte of the flag register.

5. *Explanation of arithmetic instructions*

  • Numerical data received by an 8086-based microcomputer from a terminal is usually in ASCII code. The ASCII codes for numbers 0 to 9 are 30H through 39H. Two 8-bit data can be entered into an 8086-based microcomputer via a terminal. The ASCII codes for these data (with 3 as the upper middle for each type) can be added. AAA instruction can then be used to provide the correct unpacked BCD. Suppose that ASCII codes for 2 (32_{16}) and 5 (35_{16}) are entered into an 8086-based microcomputer via a terminal. These ASCII codes can be added and then the result can be adjusted to provide the correct unpacked BCD using AAA instructions as follows:

```
ADD CL, DL   ;   [CL] = 32₁₆ = ACSII for 2
             ;   [DL] = 35₁₆ = ASCII for 5
             ;   Result [CL] = 67₁₆
MOV AL, CL   ;   Move ASCII result
             ;   into AL since AAA
             ;   adjust only [AL]
AAA          ;   [AL] = 07, unpacked
             ;   BCD for 7
```

Note that in order to send the unpacked BCD result 07_{16} back to the terminal, [AL] = 07 can be ORed with 30H to provide 37H, the ASCII code for 7.

  • DAA is used to adjust the result of adding two packed BCD numbers in AL to provide a valid BCD number. If after the addition, the low 4-bit of the result in AL is greater than 9 (or if AF = 1), then the DAA adds 6 to the low 4 bits of AL. Then, if the high 4 bits of the result in AL is greater than 9 (or if CF = 1), then DAA adds 60H to AL. As an example, consider adding two packed BCD digits 55 with 18 as follows:

```
ADD AL, DL   ;   [AL] = 55 BCD
             ;   [DL] = 18 BCD
             ;   Result = [AL] = 6DH
DAA          ;   Since low nibble
             ;   D = 1101₂ > 9, add i.e. 1101₂ + 0110₂ →
```

$$1 \qquad \underline{0011}_2$$

$$\uparrow \qquad \uparrow$$

carry 3BCD

  • The ASCII codes for two 8-bit numbers in an 8086-based microcomputer can be subtracted. Assume

```
[AL] = 35H = ASCII for 5
[DL] = 37H = ASCII for 7
```

The following instruction sequence provides the correct subtraction result:

```
SUB AL, DL  ;  [AL] = 1111 1110₂ = subtraction result
            ;  in 2's complement
            ;  CF = 1
AAS         ;  [AL] = BCD02
            ;  CF = 1 means
            ;  borrow to be
            ;  used in multibyte BCD subtraction
```

AAS adjusts [AL] and leaves zeros in the upper nibble. To output to the ASCII terminal from the microcomputer, BCD data can be ORed with 30H to produce the correct ASCII code.

- DAS can be used to adjust the result of subtraction in AL of two packed BCD numbers to provide the correct packed BCD. If low 4-bit in AL is greater than 9 (or if AF = 1), then DAS subtracts 6 from the low 4-bit of AL. Then, if the upper 4-bit of the result in AL is greater than 9 (or if CF = 1), DAS subtracts 60 from AL. While performing these subtractions, any borrows from LOW and HIGH nibbles are ignored. For example, consider subtracting BCD 55 in DL from BCD 94 in AL.

```
SUB AL, DL  ;  [AL] = 3FH  →  Low nibble = 1111
DAS         ;  CF = 0                    -6 =  1010
                                          1    1001
            ;  [AL] = 39 BCD              ↑
                                        ignore
```

- IMUL mem/reg provides signed 8 × 8 or signed 16 × 16 multiplication. As an example, if [CL] = FDH = -3_{10}, [AL] = FEH = -2_{10}, then after IMUL CL, register AX contains 0006H.

- Consider 16 × 16 unsigned multiplication, MUL WORDPTR [BX]. If [BX] = 0050H, [DS] = 3000H, [30050H] = 0002H, [AX] = 0006H, then after MUL WORDPTR [BX], [DX] = 0000H, [AX] = 000CH.

- Consider DIV BL. If [AX] = 0009H, [BL] = 02H, then after DIV BL,

$$[AH] = \text{Remainder} = 01H$$
$$[AL] = \text{Quotient} = 04H$$

- Consider IDIV WORDPTR [BX]. If [BX] = 0020H, [DS] = 2000H, [20020H] = 0004H, [DX] [AX] = 00000011H, then after IDIV. WORDPTR [BX],

$$[DX] = \text{Remainder} = 0001H$$
$$[AX] = \text{Quotient} = 0004H$$

- AAD converts two unpacked BCD digits in AH and AL to an equivalent binary number in AL. AAD must be used before dividing two unpacked BCD digits in AX by an unpacked BCD byte. For example, consider dividing [AX] = unpacked BCD 0508 (58 decimal) by [DH] = 07H. [AX] must first be converted to binary by using AAD. The register AX will then contain 003AH = 58 decimal. After DIV DH, [AL] = quotient = 08 unpacked BCD, [AH] = remainder = 02 unpacked BCD.

- Consider CBW. This instruction extends the sign from the AL register to AH register. For example, if [AL] = E2H, then after CBW, AH will contain FFH since the most significant bit of E2H is one. Note that sign extension is useful when one wants to perform an arithmetic operation on two signed numbers of different sizes. For example, the 8-bit signed number 02H can be subtracted from 16-bit signed number 2005H as follows:

$$2005H = 0010\ 00000\ 00000101$$

2's complement of 02H = $\boxed{1111\ 11111}$ 11111110

sign extend ⎯⎯⎯⎯⎯⎯⎯⎯⎯⎯⎯⎯⎯⎯

ignore carry ⎯⎯⎯⎯⎯→ $\overline{1\ 0010\ 0000\ 0000\ 0011}$

$$\begin{array}{cccc} 2 & 0 & 0 & 3H \end{array}$$
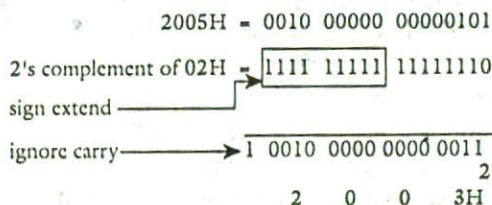
Another example of sign extension is that in order to multiply a signed 8-bit number by a signed 16-bit number, one must first sign-extend the signed 8-bit into a signed 16-bit number and then the instruction IMUL can be used for 16×16 signed multiplication.

- AAM adjusts the product of two unpacked BCD digits in AX. If [AL] = BCD3 = 00000011_2 and [CH] = BCD8 = $0000\ 1000_2$, then after MUL CH, [AX] = $0000000000011000_2 = 0018H$, and after using AAM, [AX] = $00000010\ 00000100_2 =$ unpacked BCD 24. The following instruction sequence accomplishes this:

```
MUL  CH
AAM
```

Note that the 8086 does not allow multiplication of two ASCII codes. Therefore, before multiplying two ASCII bytes received from a terminal, one must mask the upper 4-bits of each one of these bytes and then multiply them as two unpacked BCD digits and then use AAM for adjustment. In order to convert the unpacked BCD product back to ASCII, for sending back to the terminal, one must OR the product with 3030H.

6. *Examples of logical, shift and rotate instructions*

- TEST BL, 3 logically ANDs the contents of BL with 00000011_2 but does not store the result in BL. All flags are affected.
- All shift and rotate instructions include two operands. The destination operand specifies the register or memory to be shifted or rotated while the source operand specifies the number of times the register or memory contents are to be shifted. For example, SHL DX, 1 logically shifts the 16-bit contents of DX once to the left. On the other hand, SHL DX,CL, with CL = 5, logically shifts the 16-bit contents of DX five times to the left. For all 8086 shift and rotate instructions, a shift count of one must be represented in the source operand by '1' while shift counts from 2_{10} to 255_{10} must be represented in the source operand by the contents of CL.

7. *Example of string instructions*

- LODS can be represented in four forms:

```
For Byte   LODS BYTE
              or
           LODSB
```

```
For Word    LODS WORD
              or
            LODSW
```

If [SI] = 0020H, [DS] = 3000H, [30020H] = 05H, DF = 0, then after LODS BYTE or LODSB, [AL] = 05H, [SI] = 0021H.

- If [DS] = 2000H, [ES] = 3000H, [SI] = 0020H, [DI] = 0050H, DF = 0, [20020H] = 0205H, [30050H] 4071H, then after execution of MOVSW or MOVS WORD, memory location 30050H will contain 0205H. Since DF = 0, register SI will contain 0022H and register DI will contain 0052H.

- REP, a one-byte prefix, can be used with MOVS to cause the instruction MOVS to continue executing until CX = 0. Each time the instructions such as REP MOVSB or REP MOVSW are executed, CX is automatically decremented by 1, and if [CX] ≠ 0, MOVS is reexecuted and then CX is decremented by 1 until CX = 0; the next instruction is then executed. REP MOVSB or REP MOVSW can be used to move string bytes until the string length (loaded into CX before the instructions REP MOVSB or REP MOVSW) is decremented for zero.

- REPE/REPZ or REPNE/REPNZ prefix can be used with CMPS or SCAS to cause one of these instructions to continue executing until ZF = 0 (for REPNE/REPNZ prefix) or CX = 0. Note that REPE and REPZ are two mnemonics for the same prefix byte. Similarly, REPNE and REPNZ also provide the same purpose.

- If CMPS is prefixed with REPE or REPZ, the operation is interpreted as "compare while not end-of-string (CX not zero) and strings are equal (ZF = 1)". If CMPS is preceded by REPNE or REPNZ, the operation is interpreted as "compare while not end-of-string (CX not zero) and strings not equal (ZF = 0)". Thus, repeated CMPS can be used to find matching or differing string elements.

- If SCAS is prefixed with REPE or REPZ, the operation is interpreted as "scan while not end-of-string (CX not 0) and string-element = scan-value (ZF = 1)". This form may be used to scan for departure from a given value. If SCAS is prefixed with REPNE or REPNZ, the operation is interpreted as "scan while not end-of-string (CX not 0) and string-element is not equal to scan-value (ZF = 0)". This form may be used to locate a value in a string.

- Repeated string instructions are interruptible; the processor recognizes the interrupt before processing the next string element. Upon return from the interrupt, the repeated operation is resumed from the point of interruption. When multiple prefixes (such as LOCK and segment override) are specified in addition to any of the repeat prefixes, program execution does not resume properly upon return from interrupt.

  The processor remembers only one prefix in effect at the time of the interrupt, the prefix that immediately precedes the string instructions. Upon return from interrupt, program execution resumes at this point but any additional prefixes specified are not recognized. If multiple prefix must be used with a string instruction, maskable interrupts should be disabled for the duration of the repeated execution. However, this will not prevent a nonmaskable interrupt from being recognized.

- Note that the segment register for destination for all string instructions is always ES and cannot be overridden. However, DS can be overridden for the source using a prefix. For example, ES: MOVSB instruction uses the segment register as ES for both source and destination.

8. *Examples of unconditional transfers*

- There are two types of Jump instructions. These are intersegment Jumps (both IP and CS change; Jump to a different code segment) and intrasegment Jumps (IP changes and CS is fixed; Jump in the same code segment). For example, JMP FAR BEGIN (or some 8086 assemblers use JMP FAR PTR BEGIN) unconditionally branches to a label BEGIN in a different code segment. JMP START jumpts to a label START in the same code segment.

- CALL instructions can be intersegment and intrasegment. For example, CALL DWORD PTR [BX] pushes CS and IP onto the stack and loads IP and CS with the contents of four consecutive locations pointed to by BX. CALL BX, on the other hand, pushes IP onto the stack, the new value of IP is loaded from BX and CS in unchanged.

9. *Examples of conditional transfers*

All 8086 conditional branch instructions use 8-bit signed displacement. That is, the displacement covers a branch range of -128_{10} to $+127_{10}$ with 0 being positive. In order to branch out of this range, the 8086 unconditional jump instructions (having direct mode) must be used.

Conditional jumps are typically used with compare instructions to find the relationship (equal to, greater than, or less than) between two numbers. The use of conditional instructions depends on whether the numbers to be compared are signed or unsigned. The 8-bit number $1111\ 1110_2$, when considered as signed, has a value of -2_{10}; the same number has a value of $+254_{10}$ when considered as unsigned. This number, when considered signed, will be smaller than zero, and when unsigned will be greater than zero. Some new terms are used to differentiate between signed and unsigned conditional transfers. For the unsigned numbers, the terms used are "below and above", while for signed numbers, the terms "less than" and "greater than" are used. The number $0111\ 1110_2$, when considered signed is greater than $0000\ 0000_2$, while the same number $0111\ 1110_2$ when considered unsigned is above $0000\ 0000_2$. The conditional transfer instructions for equality of two numbers are the same for both signed and unsigned numbers. This is because when two numbers are compared for equality irrespective of whether they are signed or unsigned, they will provide a zero result (ZF = 1) if equal or a nonzero result (ZF = 0) if not equal. Therefore, the same instructions apply for both signed and unsigned numbers for "equal to" or "not equal to" conditions, and the various signed and unsigned conditional branch instructions for determining the relationship between two numbers are as follows:

Signed		Unsigned	
Name	Alternate name	Name	Alternate name
JE disp8 (JUMP if equal)	JZ disp8 (JUMP if result zero)	JE disp8 (JUMP if equal)	JZ disp8 (JUMP if zero)
JNE disp8 (JUMP if not equal)	JNZ disp8 (JUMP if not zero)	JNE disp8 (JUMP if not equal)	JNZ disp8 (JUMP if not zero)
JG disp8 (JUMP if greater)	JNLE disp8 (JUMP if not less or or equal)	JA disp8 (JUMP if above)	JNBE disp8 (JUMP if not below or equal)
JGE disp8 (JUMP if greater or equal)	JNL disp8 (JUMP if not less)	JAE disp8 (JUMP if above or equal)	JNB disp8 (JUMP if not below)
JL disp8 (JUMP if less than)	JNGE disp8 (JUMP if not greater or equal)	JB disp8 (JUMP if below)	JNAE disp8 (JUMP if not above or equal)
JLE disp8 (JUMP if less or equal)	JNG disp8 (JUMP if not greater)	JBE disp8 (JUMP if below or equal)	JNA disp8 (JUMP if not above)

There are also conditional transfer instructions that are concerned with the setting of status flags rather than relationship between two numbers. The table below lists these instructions:

JC disp8	JUMP if carry, i.e., CF = 1
JNC disp8	JUMP if no carry, i.e., CF = 0
JP disp8	JUMP if parity, i.e., PF = 1
JNP disp8	JUMP if no parity, i.e., PF = 0
JO disp8	JUMP if overflow, i.e., OF = 1
JNO disp8	JUMP if no overflow, i.e., OF = 0
JS disp8	JUMP if sign, i.e., SF = 1
JNS disp8	JUMP if no sign, i.e., SF = 0
JZ disp8	JUMP if result zero, i.e., Z = 1
JNZ disp8	JUMP if result not zero, i.e., Z = 0

10. *Examples of LOOP instructions*

- LOOP BEGIN automatically decrements CX by 1 without affecting the flags and jumps to the label BEGIN if CX ≠ 0; goes to the next instruction if CX = 0 after autodecrement.
- LOOPNE NEXT automatically decrements CX by 1 and jumps to the label NEXT if CX ≠ 0 and ZF = 0. However, if CX = 0 after autodecrement or ZF = 1, execution will go on to the next instruction. Note that LOOPNE can be read as 'Loop while not equal and CX not 0.'

11. *Examples of interrupt instructions*

- INT n is a software interrupt instruction. This instruction pushes Flags, CS, and IP onto the stack and loads CS from the memory location 4n and IP from the memory location 4n+2. An interrupt service routine can be written at this location.
- IRET pops IP, CS, and flags from the stack.

12. *Examples of processor or control instructions*

- STD sets direction flag to one.
- CLI clears the IF flag in the status register to zero to disable maskable interrupts.

3.5 8086 Assembler-Dependent Instructions

Some 8086 instructions do not define whether an 8-bit or 16-bit operation is to be executed. Instructions with one of the 8086 registers as an operand typically define the operation as 8-bit or 16-bit based on the register size. An example is MOV CL, [BX] which moves an 8-bit number with the offset defined by [BX] in the data segment into register CL; MOV CX, [BX], on the other hand, moves the number from offsets [BX] and [BX + 1] in the data segment into CX.

The string instructions may define this in two ways. Typical examples are MOVSB or MOVS BYTE for 8-bit and MOVSW or MOVS WORD for 16-bit. Memory offsets can also be specified by including BYTEPTR for 8-bit and WORDPTR for 16-bit with the instruction. Typical examples are INC BYTEPTR [BX] and INC WORDPTR [BX].

3.6 ASM-86 Assembler Directives

The ASM-86 is the assembler written by Intel for the 8086 microprocessor. Other 8086 assemblers include the Microsoft 8086 assembler written for the IBM personal computer and the Hewllet Packard 8086 assembler written for the HP 64000 microcomputer development system. These assemblers allow the programmer to assign the values of CS, DS, SS, and ES.

These assemblers can be used to declare a variable's type as byte (8-bit), word (16-bit), or double word (4 bytes or 2 words) as follows:

```
START DB 0  ;  START is declared
            ;  as a byte offset
            ;  and its content is initialized to zero.

BEGIN DW 0  ;  BEGIN is declared
            ;  as a word offset
            ;  and its content is initialized to zero.

NAME DD 0   ;  NAME is declared
            ;  as a double word
            ;  (4 bytes) offset
            ;  and initialized to zero.
```

The EQU directive can be used to assign a name to constants.

In the following, typical ASM-86 assembler directives such as SEGMENT, ENDS, ASSUME, and DUP are discussed.

3.6.1 SEGMENT and ENDS Directives

A section of a program or a data array can be defined by the SEGMENT and ENDS directives as follows:

```
          ORG     5000H
JOHN    SEGMENT
X1        DB      0
X2        DB      0
X3        DB      0
JOHN    ENDS
```

The segment name is JOHN. The assembler will assign a numeric value 5000H to JOHN. The programmer must use the 8086 instructions to load JOHN into DS as follows:

```
MOV BX, JOHN
MOV DS, BX
```

Note that the segment registers (except CS) must be loaded via a 16-bit register (AX, BX, CX, or DX).

3.6.2 Assume Directive

As mentioned before, the 8086, at any time, can directly address four physical segments which include a code segment, a data segment, a stack segment, and an extra segment. An 8086 program may contain a number of logical segments containing code, data, and stack. The ASSUME pseudoinstruction assigns a logical segment to a physical segment at any given time. That is, the ASSUME directive tells the assembler what addresses will be in the segment registers at execution time.

For example, the statement ASSUME CS: PROGRAM, DS: DATA, SS: STACK directs the assembler to use the logical code segment PROGRAM as CS containing the instructions, the logical data segment DATA as DS containing data, and the logical STACK segment STACK as

SS containing the stack. But, it is the responsibility of the programmer to set these segment registers in the program.

3.6.3 DUP Directive

The DUP directive can be used to initialize several locations to zero. For example, the statement START DW 4 DUP (0) reserves four words starting at the offset START in DS and initializes them to zero. The DUP directive can also be used to reserve several locations which need not be initialized. A question mark must be used with DUP in this case. For example, the statement BEGIN DB 100 DUP (?) reserves 100 bytes of uninitialized data space to an offset BEGIN in the data segment. Note that BEGIN should be typed in the label field, DB in the op field, and 100 DUP (?) in the operand field.

A typical example illustrating the use of these directives is given below:

```
DATA         SEGMENT
ADDR         DW 3005H
ADDR         DW 2003FH
DATA         ENDS
STACK        SEGMENT
             DW 60 DUP (0)   ;  Assign 60₁₀ words.
                             ;  of stack with zeros.
STACK_TOP    LABEL WORD      ;  Initialize Stack_Top
                             ;  to the next
STACK        ENDS            ;  location after the
                             ;  top of the stack.
CODE         SEGMENT
             ASSUME CS: CODE, DS: DATA, SS: STACK
             MOV AX,   STACK
             MOV SS,   AX
             LEA SP,   STACK_TOP
             MOV AX,   DATA
             MOV DS,   AX
             LEA SI,   ADDR
             LEA DI,   ADDRR
             -
             -              }   Main Program
             -                     Body
             -
CODE         ENDS
```

Note that LABEL is a directive used to initialize STACK_TOP to the next location after the top of the stack. The statement STACK_TOP LABEL WORD gives the name STACK_TOP to the next address after the 60 words are set aside for the stack. The WORD in this statement indicates that PUSH into and POP from the stack are done as words.

When the assembler translates an assembly language program, it computes the displacement, or offset, of each instruction code byte from the start of a logical segment that contains it. For example, in the above program the CS: CODE in the ASSUME statement directs the assembler to compute the offsets or displacements of the following instructions from the start of the logical segment CODE. This means that when the program is run, the CS will contain the 16-bit value where the logical segment CODE was located in memory. The assembler keeps track of the instruction byte displacements which are loaded into IP. The 20-bit physical address generated from CS and IP are used to fetch each instruction.

Another example to store data bytes in a data segment and to allocate stack is given in the following:

```
        DSEG        SEGMENT
        ARRAY       DB 02H, F1H, A2H    ;   store 3 bytes
                                        ;   of data in an
        DSEG        ENDS                ;   address defined
                    -                   ;   by DSEG as DS
                    -                   ;   and ARRAY as
                    -                   ;   offset
        SSEG        SEGMENT
                    DW 10 DUP (0)       ;   Allocate
                                        ;   10 word stack
  STACK_TOP LABEL WORD                  ;   Label initial
                                        ;   TOS
        SSEG        ENDS
                    -
                    -
                    -
                    -
                    MOV AX, DSEG        ;   Initialize
                    MOV DS, AX          ;   DS
                    MOV AX, SSEG        ;   Initialize
                    MOV SS, AX          ;   SS
                    MOV SP, STACK_TOP   ;   Initialize SP
                    -
                    -
                    -
```

Note that typical 8086 assemblers on the Microsoft software and HP64000 system use the ORG directive to load CS and IP. For example, CS and IP can be initialized with 0500H and 0020H as follows:

> For Microsoft 8086 assembler: ORG 05000020H
> For HP64000 8086 assembler: ORG 0500H:0020H

Example 3.1

Determine the effect of each one of the following 8086 instructions:

i)   PUSH [BX]
ii) DIV DH
iii) CWD
iv) MOVSB
v)   MOV START [BX], AL

Assume the following data prior to execution of each one of the above instructions independently. Assume all numbers in hexadecimal.

[DS]	= 3000H	[SI]	= 0400H
[ES]	= 5000H	[DI]	= 0500H
[DX]	= 0400H	DF	≐ 0
[SP]	= 5000H	[BX]	= 6000H
[SS]	= 6000H	Value of START	= 05H
[AX]	= 00A9H		

[36000H] = 02H, [36001H] = 03H
[50500H] = 05H
[30400H] = 02H, [30401H] = 03H

Solution

i) 20-bit physical memory addressed by DS and BX = 36000H. 20-bit physical location pointed to by SP and SS = 65000H. PUSH [BX] pushes [36001H] and [36000H] into stack locations 64FFFH and 64FFEH, respectively. The SP is then decremented by 2 to contain the 20-bit physical address 64FFEH (SS = 6000H, SP = 4FFEH). Therefore, [64FFFH] = 03H and [64FFEH] = 02H.

ii) Before unsigned division, [DX] = 0400H, DH contains 04H, and [AX] = 00A9H = 169_{10}. After DIV DH, [AH] = Remainder = 01H and [AL] = Quotient = 2AH = 42_{10}.

iii) CWD sign extends AX register into the DX register. Since the sign bit of [AX] = 0, after CWD, [DX AX] = 000000A9H.

iv) MOVSB moves the content of memory addressed by the source [DS] and [SI] to the destination addressed by [ES] and [DI] and then it increments SI and DI by 1 for byte move. Since DF = 0, [DS] = 3000H, [SI] = 0400H, [ES] = 5000H, [DI] = 0500H, and the content of physical memory location 30400H is moved to physical memory 50500H. Since [30400H] = 02H, the location 50500H will also contain 02H. Since DF = 0 after MOVSB, [SI] = 0401H, [DI] = 0501H.

v) Since [BX] = 6000H, [DS] = 3000H, START = 05H, and the physical memory for destination = 36005H. After MOV START [BX], AL, memory location 36005H will contain A9H.

Example 3.2

Write 8086 assembly program to clear 100_{10} consecutive bytes.

Solution

```
0000                     CODE  SEGMENT  AT  1000H       ; Program to
                                                        ; clear 100 bytes

                         ASSUME CS:CODE,  DS:DATA
0000  B8 2000            MOV AX,  2000H                 ; Initialize
0003  8E D8              MOV DS,  AX                     ; Data Segment
0005  7D 1E 0000 R       LEA BX,  ADDR                   ; Initialize BX
0009  B9 0064            MOV CX,  100                    ; Initialize loop
                                                        ; count
000C  C6 07 00   START:MOV BYTEPTR[BX],  0              ; Clear memory byte
000F  43                 INC BX                          ; Update pointer
0010  E2 FA              LOOP START                      ; Decrement CX and
                                                        ; loop
0012  F4                 HLT                             ; Halt
0013               CODE  ENDS
0000               DATA  SEGMENT  AT  2000H
0000  3000         ADDR  DW 3000H                        ; Store the initial
                                                        ; address
0002               DATA  ENDS                            ; End program
                         END
```

Microsoft (R) Macro Assembler Version 5.10 5/13/92 13:12:12
                                                           Symbols-1

Segments and Groups:

Name	Length	Align	Combine Class
CODE	0013	AT	1000
DATA	0002	AT	2000

Symbols:

Name	Type	Value	Attr
ADDR	L WORD	0000	DATA
START	L NEAR	000C	CODE
@CPU	TEXT	0101h	
@FILENAME	TEXT	ex32	
@VERSION	TEXT	510	

     28 Source Lines
     28 Total Lines
      9 Symbols

47886 + 445965 Bytes symbol space free

     0 Warning Errors
     0 Severe Errors

Example 3.3

Write 8086 assembly program to computer $\sum_{i=1}^{N} XiYi$ where Xi and Yi are signed 8-bit numbers. N = 100. Assume no overflow.

Solution

```
0000                  CODE SEGMENT AT 1000H
                      ASSUME CS:CODE, DS:DATA
0000  B8 2000         MOV AX, 2000H          ; Initialize
0003  8E D8           MOV DS, AX             ; Data Segment
0005  B9 0064         MOV CX, 100            ; Initialize loop
                                             ; count
0008  8D 1E 0000 R    LEA BX, ADDR1          ; Load ADDR1
                                             ; into BX
000C  8D 36 0002 R    LEA SI, ADDR2          ; Load ADDR2
                                             ; into SI
0010  BA 0000         MOV DX, 0000H          ; Initialize sum
                                             ; to zero
0013  8A 07    START:MOV AL, [BX]            ; Load data into
                                             ; AL
0015  F6 2C           IMUL BYTEPTR [SI]      ; Signed 8x8
                                             ; multiplication
```

```
0017   03 D0          ADD DX, AX           ;  Sum XiYi
0019   43             INC BX               ;  Update pointer
001A   46             INC SI               ;  Update pointer
001B   E2 F6          LOOP START           ;  Decrement CX
                                           ;  and loop
001E              CODE ENDS
0000              DATA SEGMENT AT 2000H
0000   0000       ADDR1 DW 0000H           ;  Location of Xi
0002   1000       ADDR2 DW 1000H           ;  Location of Yi
0004              DATA ENDS
                  END                      ;  End program
```

Microsoft (R) Macro Assembler Version 5.10   5/13/92 13:26:33

Segments and Groups:

Name	Length	Align	Combine Class
CODE	001E	AT	1000
DATA	0004	AT	2000

Symbols:

Name	Type	Value	Attr
ADDR1	L WORD	0000	DATA
ADDR2	L WORD	0002	DATA
START	L NEAR	0013	CODE
@CPU	TEXT	0101h	
@FILENAME	TEXT	ex33	
@VERSION	TEXT	510	

```
    39  Source Lines
    39  Total Lines
    10  Symbols

47848 + 445003 Bytes symbol space free

     0  Warning Errors
     0  Severe Errors
```

Example 3.4

Write 8086 assembly language program to add two words; each word contains four packed BCD digits. The first word is stored in two consecutive locations with the low byte at the offset pointed to by SI at 0500H, while the second word is stored in two consecutive locations with the low byte pointed to by BX at the offset 1000H. Store the result in memory pointed to by BX.

Solution

```
0000              CODE SEGMENT AT 1000H
                  ASSUME CS:CODE, DS:DATA
0000   BE 2000    MOV AX, 2000H            ;  Initialize
```

```
0003   8E D8                MOV DS, AX           ;  Data Segment
0005   B9 0002              MOV CX, 2            ;  Initialize loop
                                                 ;  count
0008   BE 0500              MOV SI, 0500H        ;  Initialize SI
000B   BB 1000              MOV BX, 1000H        ;  Initialize BX
000E   F8                   CLC                  ;  Clear carry
000F   8A 04          START:MOV AL, [SI]         ;  Move data
0011   12 07                ADC AL, [BX]         ;  Perform addition
0013   27                   DAA                  ;  BCD adjust
0014   88 07                MOV [BX], AL         ;  Store result
0016   46                   INC SI               ;  Update pointer
0017   43                   INC BX               ;  Update pointer
0018   E2 F5                LOOP START           ;  Decrement CX and
                                                 ;  loop
001A   F4                   HLT                  ;  Halt
001B              CODE ENDS
0000              DATA SEGMENT AT 2000H
0000              DATA ENDS
                  ENDS                           ;  End program
```

Microsoft (R) Macro Assembler Version 5.10

Segments and Groups:

Name	Length	Align	Combine Class
CODE	001B	AT	1000
DATA	0000	AT	2000

Symbols:

Name	Type	Value	Attr
START	L NEAR	000F	CODE
@CPU	TEXT	0101h	
@FILENAME	TEXT	ex34	
@VERSION	TEXT	510	

```
    36  Source Lines
    36  Total Lines
     8  Symbols

47890 + 445961 Bytes symbol space free

     0  Warning Errors
     0  Severe Errors
```

Example 3.5

Write an 8086 assembly language program to add two words; each contains two ASCII digits. The first word is stored in two consecutive locations with the low byte pointed to by SI at offset 0300H, while the second byte is stored in two consecutive locations with the low byte pointed to by DI at offset 0700H. Store the result in memory pointed to by DI.

Solution

```
0000                    CODE SEGMENT AT 1000H
                        ASSUME CS:CODE, DS:DATA
0000   B8 2000          MOV AX, 2000H          ; Initialize
0003   8E D8            MOV DS, AX             ; Data Segment
0005   B9 0002          MOV CX, 2             ; Initialize loop
                                              ; count
0008   BE 0300          MOV SI, 0300H         ; Initialize SI
000B   BF 0700          MOV DI, 0700H         ; Initialize DI
000E   F8               CLC                   ; Clear carry
000F   8A 04  START:    MOV AL, [SI]          ; Move data
0011   12 05            ADC AL, [DI]          ; Perform addition
0013   37               AAA                   ; ASCII adjust
0014   88 05            MOV [DI], AL          ; Store result
0016   46               INC SI                ; Update pointer
0017   47               INC DI                ; Update pointer
0018   E2 F5            LOOP START            ; Decrement CX and
                                              ; loop
001A                    HLT                   ; Halt
001B                CODE ENDS
0000                DATA SEGMENT AT 2000H
0000                DATA ENDS
                    END                       · End program
```

Microsoft (R) Macro Assembler Version 5.10 10/05/93 22:25:10
                                              Symbols-1

Segments and Groups:

Name	Length	Align	Combine Class
CODE	001B	AT	1000
DATA	0000	AT	2000

Symbols:

Name	Type	Value	Attr
START	L NEAR	000F	CODE
@CPU	TEXT	0101h	
@FILENAME	TEXT	ex35	
@VERSION	TEXT	510	

```
    36 Source Lines
    36 Total Lines
     8 Symbols

47890 + 445961 Bytes symbol space free

     0 Warning Errors
     0 Severe Errors
```

Example 3.6

Write an 8086 assembly language program to compare a source string of 50_{10} words pointed to by an offset of 2000H in the data segment with a destination string pointed to by an offset 3000H in another segment. The program should be halted as soon as a match is found or the end of string is reached.

Solution

```
0000                      CODE  SEGMENT AT 1000H
                              ASSUME  CS:CODE,DS:DATA,ES:DATAA
0000  B8  2000            MOV  AX,2000H          ;  initialize
0003  8E  D8              MOV  DS,AX             ;  Data segment
0005  B8  4000            MOV  AX,4000H          ;  Initialize
0008  8E  C0              MOV  ES,AX             ;  ES
000A  BE  2000            MOV  SI,2000H          ;  Initialize SI
000D  BF  3000            MOV  DI,3000H          ;  Initialize DI
0010  B9  0032            MOV  CX,50             ;  Initialize CX
0013  FC                  CLD                   ;  DF is cleared so
                                                ;  that SI and DI
                                                ;  will autoincrement
                                                ;  after compare
0014  F2/A7               REPNE CMPSW            ;  Repeat CMPSW until
                                                ;  CX=0 or until
                                                ;  compared words are
                                                ;  equal
0016  F4                  HLT                   ;  Halt
0017                  CODE  ENDS
0000                  DATA  SEGMENT AT 2000H
0000                  DATA  ENDS
0000                  DATAA SEGMENT AT 4000H
                          END                   ;  End program
```

Microsoft (R) Macro Assembler Version 5.10 10/05/92 11:45:41

Segments and Groups:

Name	Length	Align	Combine Class
CODE	0017	AT	1000
DATA	0000	AT	2000
DATAA	0000	AT	4000

Symbols:

Name	Type	Value	Attr
@CPU	TEXT	0101h	
@FILENAME	TEXT	TEST1	
@VERSION	TEXT	510	

20 Source Lines

```
      20 Total Lines
       8 Symbols

   47884 + 433567 Bytes symbol space free

       0 Warning Errors
       0 Severe Errors
```

Example 3.7

Write a subroutine in 8086 assembly language which can be called by a main program in a different code segment. The subroutine will multiply a signed 16-bit number in CX by a signed 8-bit number in AL. The main program will call this subroutine, store the result in two consecutive memory words, and stop. Assume SI and DI contain pointers to the signed 8-bit and 16-bit data, respectively.

Solution

```
0000                    CODE SEGMENT AT 1000H
                            ASSUME CS:CODE,DS:DATA,SS:STACK
0000  B8 5000           MOV AX, 5000H          ;  Initialize
0003  8E D8             MOV DS, AX             ;  Data Segment
0005  B8 6000           MOV AX, 6000H          ;  Initialize
0008  8E D0             MOV SS, AX             ;  Stack Segment
000A  BC 0020           MOV SP, 0020H          ;  Initialize SP
000D  BB 2000           MOV BX, 2000H          ;  Initialize BX
0010  8A 04             MOV AL, [SI]           ;  Move 8-bit data
0012  8B 0D             MOV CX, [DI]           ;  Move 16-bit data
0014  9A 0000---R       CALL FAR PTR MULTI     ;  Call MULTI
                                               ;  subroutine
0019  89 17             MOV [BX], DX           ;  Store high word of
                                               ;  result
001B  89 47 02          MOV [BX+2], AX         ;  Store low word of
                                               ;  result
001E  F4                HLT                    ;  Halt
001F                    CODE ENDS
0000                    SUBR SEMENT AT 7000H
                            ASSUME CS:SUBR
0000                    MULTI PROC FAR         ;  Must be called from
0000  51                    PUSH CX            ;  another code
                                               ;  segment
0001  50                    PUSH AX            ;
0002  98                    CBW                ;  Sign extend AL
0003  F7 E9                 IMUL CX            ;  [DX][AX]←[AX]*[CX]
0005  58                    POP AX             ;
0006  59                    POP CX             ;
0007  CB                    RET                ;  Return
0008                    MULTI ENDP             ;  End of procedure
0004                    SUBR ENDS              ;  End subroutine
0000                    DATA SEGMENT AT 5000H
0000                    DATA ENDS
0000                    STACK SEGMENT AT 6000H
0000                    STACK ENDS
                            END                ;  End program
```

Microsoft (R) Macro Assembler Version 5.10 10/05/93 22:40:32
Symbols-1

Segments and Groups:

Name	Length	Align	Combine Class
CODE	001F	AT	1000
DATA	0000	AT	5000
STACK	0000	AT	6000
SUBR	0004	AT	7000

Symbols:

Name	Type	Value	Attr
MULTI	F PROC	0000	SUBR Length=0004
@CPU	TEXT	0101h	
@FILENAME	TEXT	EX37	
@VERSION	TEXT	510	

```
    52 Source Lines
    52 Total Lines
    10 Symbols

47852 + 443951 Bytes symbol space free

     0 Warning Errors
     0 Severe Errors
```

Example 3.8

Write an 8086 assembly language program to subtract two 64-bit numbers. Assume SI and DI contain the starting address of the numbers. Store the result in memory pointed to by [DI].

Solution

```
0000                CODE SEGMENT AT 4000H
3000                    ORG 3000H
3000   05712FA2   DATAA DD 05712FA2H          ;  DATAA LOW
3004   0248A201         DD 0248A201H           ;  DATAA HIGH
3008   561A2604   DATAB DD 561A2604H          ;  DATAB LOW
300C   72A2B270         DD 72A2B270H           ;  DATAB HIGH
3010                DATA ENDS
0000                PROG SEGMENT
                        ASSUME CS:PROG,DS:DATA
0000   B8 4000          MOV AX, 4000H          ;  Initialize
0003   8E D8            MOV DS, AX             ;  DS
0005   BA 0004          MOV DX, 4              ;  Load 4 into DX
0008   8D 36 3000 R     LEA SI, DATAA          ;  Initialize SI
000C   8D 3E 30008 R    LEA DI, DATAB          ;  Initialize DI
0010   F8               CLC                    ;  Clear carry
0011   8B 04     START: MOV AX, [SI]           ;  Load DATAA
0013   19 05            SBB [DI], AX           ;  Perform subtraction
```

```
0015   83  C6  02         ADD  SI, 2             ;  Update
0018   83  C7  02         ADD  DI, 2             ;  pointers
001B   4A                 DEC  DX
001C   75  F3             JNZ  START
001E   F4                 HLT
001F                PROG  ENDS
                          END
```

Microsoft (R) Macro Assembler Version 5.10   10/10/93 22:14:0
                                              Symbols-1

Segments and Groups:

Name	Length	Align	Combine	Class
DATA	3010	AT	4000	
PROG	001F	PARA	NONE	

Symbols:

Name	Type	Value	Attr
DATAA	L DWORD	3000	DATA
DATAB	L DWORD	3008	DATA
START	L NEAR	0011	PROG
@CPU	TEXT	0101h	
@FILENAME	TEXT	examp38	
@VERSION	TEXT	510	

```
      25  Source Lines
      25  Total Lines
      10  Symbols

47518 + 428045 Bytes symbol space free

       0  Warning Errors
       0  Severe Errors
```

Example 3.9

Write an 8086 assembly language program that will perform the following operation:

$$5 * AL - 6 * BH + (BH/8) \rightarrow CX$$

Assume that the stack pointer is already initialized and all numbers are unsigned.

Solution

```
0000               PROG  SEGMENT AT 5000H
0500                     ORG 0500H
                         ASSUME CS:PROG
0500   B3  05            MOV  BL,5             ;  Compute
0502   F6  E3            MUL  BL               ;  AX←5*AL
0504   8B  C8            MOV  CX,AX            ;  Store in CX
```

```
0506   B3  06        MOV  BL,6         ;
0508   8A  C7        MOV  AL,BH        ; Compute
050A   F6  E3        MUL  BL           ; AX←6*BH
050C   29  C8        SUB  CX,AX        ; store in CX
050E   51            PUSH CX
050F   B1  03        MOV  CL,3
0511   D2  EF        SHR  BH,CL        ; Compute BH/8
0513   8A  DF        MOV  BL,BH        ; Convert BH/8 to
0515   B7  00        MOV  BH,0         ; 16-bit unsigned
                                       ; number in BX

0517   59            POP  CX
0518   03  CB        ADD  CX,BX        ; Store final result
                                       ; in CX

051A   F4            HLT
051B             PROG ENDS
                     END
```

Microsoft (R) Macro Assembler Version 5.10   1/16/80 12:07:44
                                             Symbols-1

Segments and Groups:

Name	Length	Align	Combine Class
PROG	051B	AT	5000

Symbols:

Name	Type	Value	Attr
@CPU	TEXT	0101h	
@FILENAME		TEXT	TEST1
@VERSION		TEXT	510

```
   20  Source Lines
   20  Total Lines
    6  Symbols

47808 + 433643 Bytes symbol space free

    0  Warning Errors
    0  Severe Errors
```

Example 3.10

Write an 8086 assembly program that converts a number from Fahrenheit degrees to Celsius degrees. Each number is only one byte. The source byte is assumed to reside at offset 1000H in the data segment, and the destination at an offset of 2000H in the data segment. Discard the remainder of the result.

$$\text{use } C = \frac{(F - 32)}{9} \times 5$$

Solution

```
0000                    CODE  SEGMENT
                        ASSUME  CS:CODE,DS:DATA
0000  BE 1000           MOV SI, 1000H          ;  Initialize source
                                               ;  pointer
0003  BF 2000           MOV DI, 2000H          ;  Initialize
                                               ;  destination pointer
                                               ;
0006  B4 00             MOV AH, 0              ;  Clear AX high byte
0008  8A 04             MOV AL, [SI]           ;  Get degrees F
000A  98                CBW                    ;  Sign extend
000B  2D 0020           SUB AX, 32             ;  Subtract 32
                                               ;
000E  B9 0005           MOV CX, 5              ;  Get multiplier
0011  F7 E9             IMUL CX                ;  Multiply by 5
                                               ;
0013  B9 0009           MOV CX, 9              ;  Get divisor
0016  F7 F9             IDIV CX                ;  divide by 9 to
                                               ;  get Celsius
                                               ;
0017  88 05             MOV [DI], AL           ;  Put result in
                                               ;  destination
                                               ;
001A                    CODE  ENDS             ;  End segment
                                               ;
                        END                    ;
```

Microsoft (R) Macro Assembler Version 5.10 10/10/93 22:27:4

                                                     Symbols-1

Segments and Groups:

Name	Length	Align	Combine	Class
CODE	0019	PARA	NONE	

Symbols:

Name	Type	Value	Attr
GO	L NEAR	0000	CODE
@CPU	TEXT	0101h	
@FILENAME	TEXT	examp	310
@VERSION	TEXT	510	

    22 Source Lines
    22 Total Lines
    11 Symbols

47464 + 428099 Bytes symbol space free

     0 Warning Errors
     0 Severe Errors

3.7 System Design Using the 8086

This section covers the basic concepts associated with interfacing the 8086 to its support chips such as memory and I/O. Topics such as timing diagrams and 8086 pins and signals will also be included.

3.7.1 Pins and Signals

The 8086 pins and signals are shown in Figure 3.6. Unless otherwise indicated, all 8086 pins are TTL compatible. As mentioned before, the 8086 can operate in two modes. These are minimum mode (uniprocessor system — single 8086) and maximum mode (multiprocessor system — more than one 8086). MN/MX is an input pin used to select one of these modes. When MN/MX is HIGH, the 8086 operates in the minimum mode. In this mode, the 8086 is configured (that is, pins are defined) to support small, single processor systems using a few devices that use the system bus.

When MN/MX is LOW, the 8086 is configured (that is, pins are defined in the maximum mode) to support multiprocessor systems. In this case, the Intel 8288 bus controller is added to the 8086 to provide bus controls and compatibility with the multibus architecture. Note that in a particular application, the MN/MX must be tied to either HIGH or LOW.

The AD0-AD15 lines are a 16-bit multiplexed address/data bus. During the first clock cycle AD0-AD15 are the low order 16 bits of address. The 8086 has a total of 20 address lines. The upper four lines are multiplexed with the status signals for the 8086. These are the A16/S3, A17/S4, A18/S5, and A19/S6. During the first clock period of a bus cycle (read or write cycle), the entire 20-bit address is available on these lines. During all other clock cycles for memory and I/O operations, AD15-AD0 contain the 16-bit data, and S3, S4, S5, and S6 become status lines. S3 and S4 lines are decoded as follows:

A17/S4	A16/S3	Function
0	0	Extra segment
0	1	Stack segment
1	0	Code or no segment
1	1	Data segment

Therefore, after the first clock cycle of an instruction execution, the A17/S4 and A16/S3 pins specify which segment register generates the segment portion of the 8086 address. Thus, by decoding these lines and then using the decoder outputs as chip selects for memory chips, up to 4 megabytes (one megabyte per segment) can be provided. This provides a degree of protection by preventing erroneous write operations to one segment from overlapping into another segment and destroying information in that segment. A18/S5 and A19/S6 are used as A18 and A19, respectively, during the first clock period of an instruction execution. If an I/O instruction is executed, they stay low during the first clock period. During all other cycles, A18/S5 indicates the status of the 8086 interrupt enable flag and a low A19/S6 pin indicates that the 8086 is on the bus. During a "Hold Acknowledge" clock period, the 8086 tristates the A19/S6 pin and thus allows another bus master to take control of the system bus.

The 8086 tristates AD0-AD15 during Interrupt Acknowledge or Hold Acknowledge cycles.

BHE/S7 is used as BHE (Bus High Enable) during the first clock cycle of an instruction execution. The 8086 outputs a low on this pin during read, write, and interrupt acknowledge cycles in which data are to be transferred in a high-order byte (AD15-AD8) of the data bus. BHE can be used in conjunction with AD0 to select memory banks. A thorough discussion is provided later. During all other cycles BHE/S7 is used as S7 and the 8086 maintains the output level (BHE) of the first clock cycle on this pin.

Common Signals

Name	Function	Type
AD15-AD0	Address/Data Bus	Bidirectional, 3-State
A19/S6-A16/S3	Address/Status	Output, 3-State
$\overline{BHE}$/S7	Bus High Enable/Status	Output, 3-State
MN/$\overline{MX}$	Minimum/Maximum Mode Control	Input
$\overline{RD}$	Read Control	Output, 3-State
$\overline{TEST}$	Wait On Test Control	Input
READY	Wait State Control	Input
RESET	System REset	Input
NMI	Non-Maskable Interrupt Request	Input
INTR	Interrupt Request	Input
CLK	System Clock	Input
Vcc	+5V	Input
GND	Ground	

Mimimum Mode Signals (MN/MX = Vcc)

Name	Function	Type
HOLD	Hold Request	Input
HLDA	Hold Acknowledge	Output
$\overline{WR}$	Write Control	Output, 3-State
M/$\overline{IO}$	Memory/IO Control	Output, 3-State
DT/$\overline{R}$	Data Transmit/Receive	Output, 3-State
$\overline{DEN}$	Data Enable	Output, 3-State
ALE	Address Latch Enable	Output
$\overline{INTA}$	Interrupt Acknowledge	Output

Maximum Mode Signals (MN/MX = GND)

Name	Function	Type
RQ/GT1, 0	Request/Grant Bus Access Control	Bidirectional
LOCK	Bus Priority Lock Control	Output, 3-State
S2-S0	Bus Cycle Status	Output, 3-State
QS1, QS0	Instruction Queue Status	Output

GND □ 1	40 □ VCC	
AD14 □ 2	39 □ AD15	
AD13 □ 3	38 □ AD16/S3	
AD12 □ 4	37 □ A17/S4	
AD11 □ 5	36 □ A18/S5	
AD10 □ 6	35 □ A19/S6	
AD9 □ 7	34 □ $\overline{BHE}$/S7	
AD8 □ 8	33 □ MN/$\overline{MX}$	
AD7 □ 9	32 □ $\overline{RD}$	
AD6 □ 10	31 □ HOLD	(RQ/$\overline{GT0}$)
AD5 □ 11	30 □ HLDA	(RQ/$\overline{GT1}$)
AD4 □ 12	29 □ $\overline{WR}$	($\overline{LOCK}$)
AD3 □ 13	28 □ M/$\overline{IO}$	($\overline{S2}$)
AD2 □ 14	27 □ DT/$\overline{R}$	($\overline{S1}$)
AD1 □ 15	26 □ $\overline{DEN}$	($\overline{S0}$)
AD0 □ 16	25 □ ALE	(QS0)
NMI □ 17	24 □ $\overline{INTA}$	(QS1)
INTR □ 18	23 □ $\overline{TEST}$	
CLK □ 19	22 □ READY	
GND □ 20	21 □ RESET	

8086 CPU
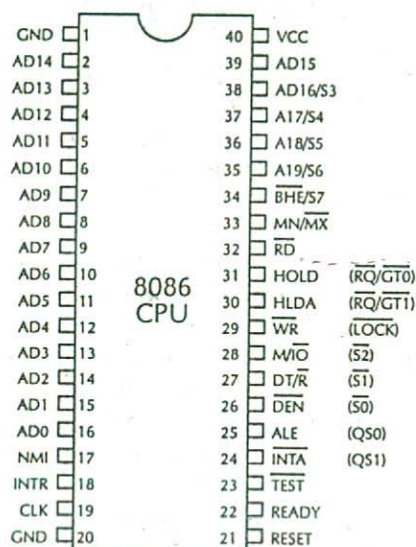
Maximum mode pin functions (e.g., $\overline{LOCK}$) are shown in parenthesis

FIGURE 3.6   8086 pins and signals.

$\overline{RD}$ is LOW whenever the 8086 is reading data from memory or an I/O location.

$\overline{TEST}$ is an input pin and is only used by the WAIT instruction. The 8086 enters a wait state after execution of the WAIT instruction until a LOW is seen on the TEST pin. This input is synchronized internally during each clock cycle on the leading edge of the CLK pin.

INTR is the maskable interrupt input. This line is not latched and, therefore, INTR must be held at a HIGH level until recognized to generate an interrupt.

NMI is the nonmaskable interrupt input activated by a leading edge.

RESET is the system reset input signal. This signal must be high for at least four clock cycles to be recognized, except after power-on which requires a 50-ms reset pulse. It causes the 8086 to initialize registers DS, SS, ES, IP, and flags to all zeros. It also initializes CS to FFFFH. Upon

removal of the RESET signal from the RESET pin, the 8086 will fetch its next instruction from 20-bit physical address FFFF0H (CS = FFFFH, IP = 0000H).

The <u>reset</u> signal to the 8086 can be generated by the 8284. The 8284 has a Schmitt Trigger input (RES) for generating reset from a low active external reset.

Since the reset vector is located at the physical address FFFF0H, there may not be enough locations available to write programs. Typical assemblers such as the Microsoft 8086 assembler use the following to jump to a different code segment to write programs:

```
ORG       0FFFF0H ;reset vector
          JMP       FAR BEGIN
          ORG       30000500H
BEGIN     —    ⎫
          —    ⎪  user
          —    ⎬  program
          —    ⎭
```

The above instruction sequence will allow the 8086 to jump to the offset BEGIN (offset 0500H) in code segment 3000H upon hardware reset where the user can write his/her programs.

To guarantee reset from power-up, the 8284 reset input must remain below 1.05 volts for 50 microseconds after Vcc has reached the minimum supply voltage of 4.5V. The $\overline{RES}$ input of the 8284 can be driven by a simple RC circuit as shown in Figure 3.7.

The values of R and C can be selected as follows:

$$Vc\,(t) = V\,(1 - exp - (t/RC))$$

where t = 50 microseconds, V = 4.5 V, Vc = 1.05V, and RC = 188 microseconds. For example, if C is chosen arbitrarily to be 0.1 μF, then R = 1.88 KΩ.
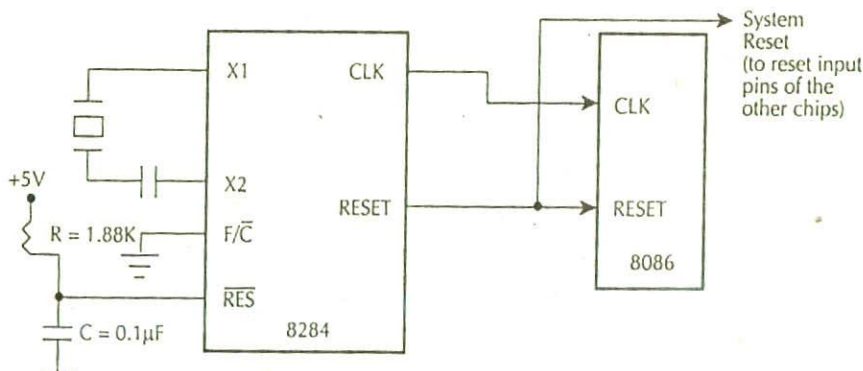
FIGURE 3.7   8086 reset and system resst.

As mentioned before, the 8086 can be configured in either minimum or maximum mode using the MN/MX input pin. In minimum mode, the 8086 itself generates all bus control signals. These signals are

- DT/$\overline{\text{R}}$ (Data Transmit/Receive). DT/$\overline{\text{R}}$ is an output signal required in minimum system that uses an 8286/8287 data bus transceiver. It is used to control direction of data flow through the transceiver.
- $\overline{\text{DEN}}$ (Data Enable) is provided as an output enable for the 8286/8287 in a minimum system which uses the transceiver.
- $\overline{\text{DEN}}$ is active LOW during each memory and I/O access and for $\overline{\text{INTA}}$ cycles.
- ALE (Address Latch Enable) is an output signal provided by the 8086 and can be used to demultiplex the AD0-AD15 into A0-A15 and D0-D15 at the falling edge of ALE. The 8086 ALE signal is similar to the 8085 ALE.
- M/IO. This 8086 output signal is similar to the 8085 IO/$\overline{\text{M}}$. It is used to distinguish a memory access (M/IO = HIGH) from an I/O access (M/IO = LOW). When the 8086 executes an I/O instruction such as IN or OUT, it outputs a LOW on this pin. On the other hand, the 8086 outputs HIGH on this pin when it executes a memory reference instruction such as MOVE AX, [SI].
- $\overline{\text{WR}}$. The 8086 outputs LOW on this pin to indicate that the processor is performing a write memory or write I/O operation, depending on the M/IO signal.
- $\overline{\text{INTA}}$. The 8086 $\overline{\text{INTA}}$ is similar to the 8085 $\overline{\text{INTA}}$. For Interrupt Acknowledge cycles (for INTR pin), the 8086 outputs LOW on this pin.
- HOLD (input), HLDA (output). These pins have the same purpose as the 8085 HOLD/HLDA pins and are used for DMA. A HIGH on the HOLD pin indicates that another master is requesting to take over the system bus. The processor receiving the HOLD request will output HLDA high as an acknowledgment. At the same time, the processor tristates the system bus. Upon receipt of LOW on the HOLD pin, the processor places LOW on the HLDA pin. HOLD is not an asynchronous input. External synchronization should be provided if the system cannot otherwise guarantee the setup time.
- CLK (input) provides the basic timing for the 8086.

The maximum clock frequencies of the 8086-1, 8086, and 8086-2 are 5 MHz, 10 MHz, and 8 MHz, respectively. Since the design of these processors incorporates dynamic cells, a minimum frequency of 2 MHz is required to retain the state of the machine. The 8086, 8086-1, and 8086-2 will be referred to as 8086 in the following. Since the 8086 does not have on-chip clock generation circuitry, an 8284 clock generator chip must be connected to the 8086 CLK pin as shown in Figure 3.8.
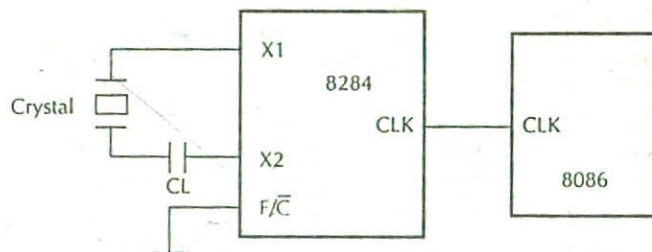
FIGURE 3.8   8284 clock generator connections to the 8086.

The crystal must have a frequency three times the 8086 internal frequency. That is, the 8284 divides the crystal clock frequency by 3. In other words, to generate a 5 MHz 8086 internal

clock, the crystal clock must be 15 MHz. To select the crystal inputs of the 8284 as the frequency source for clock generation, the F/C input must be strapped to ground. This strapping option allows either the crystal or an external frequency input as the source for clock generation. When selecting a crystal for use with the 8284, the crystal series resistance should be as low as possible. Two crystal manufacturers recommended by Intel are Crystle Corp. model CY 15A (15 MHz) and CTS Knight Inc. model CY 24 A (24 MHz).

Note that the 8284 can be used to generate the 8086 READY input signal based on inputs from slow memory and I/O devices which are not capable of transferring information at the 8086 rate.

In the maximum mode, some of the 8086 pins in the minimum mode are redefined. For example, pins HOLD, HLDA, WR, M/IO, DT/R, DEN, ALE, and INTA in the minimum mode are redefined as RQ/GT0, RQ/GT1, LOCK, S2, S1, S0, QS0, and QS1, respectively. In maximum mode, the 8288 bus controller decodes the status information from S0, S1, S2 to generate bus timing and control signals required for a bus cycle. S2, S1, S0 are 8086 outputs and are decoded as follows:

$\overline{S2}$	$\overline{S1}$	$\overline{S0}$	
0	0	0	Interrupt Acknowledge
0	0	1	Read I/O port
0	1	0	Write I/O port
0	1	1	Halt
1	0	0	Code access
1	0	1	Read memory
1	1	0	Write memory
1	1	1	Inactive

- $\overline{RQ}/\overline{GT0}$, $\overline{RQ}/\overline{GT1}$. These request/grant pins are used by other local bus masters to force the processor to release the local bus at the end of the processor's current bus cycle. Each pin is bidirectional, with $\overline{RQ}/\overline{GT0}$ having higher priority than $\overline{RQ}/\overline{GT1}$. These pins have internal pull-up resistors so that they may be left unconnected. The request/grant function of the 8086 works as follows:
  1. A pulse (one clock wide) from another local bus master ($\overline{RQ}/\overline{GT0}$ or $\overline{RQ}/\overline{GT1}$ pins) indicates a local bus request to the 8086.
  2. At the end of 8086 current bus cycle, a pulse (one clock wide) from the 8086 to the requesting master indicates that the 8086 has relinquished the system bus and tristated the outputs. Then the new bus master subsequently relinquishes control of the system bus by sending a LOW on $\overline{RQ}/\overline{GT0}$ or $\overline{RQ}/\overline{GT1}$ pins. The 8086 then regains bus control.
- LOCK . The 8086 outputs LOW on the $\overline{LOCK}$ pin to prevent other bus masters from gaining control of the system bus. The LOCK signal is activated by the 'LOCK' prefix instruction and remains active until the completion of the instruction that follows.
- QS1, QS0. The 8086 outputs to QS1 and QS0 pins to provide status to allow external tracking of the internal 8086 instruction queue as follows:

QS1	QS0	
0	0	No operation
0	1	First byte of op code from queue
1	0	Empty the queue
1	1	Subsequent byte from queue

QS0 and QS1 are valid during the clock period following any queue operation. The 8086 can be operated from a +5V to +10V power supply. There are two ground pins on the chip to distribute power for noise reduction.
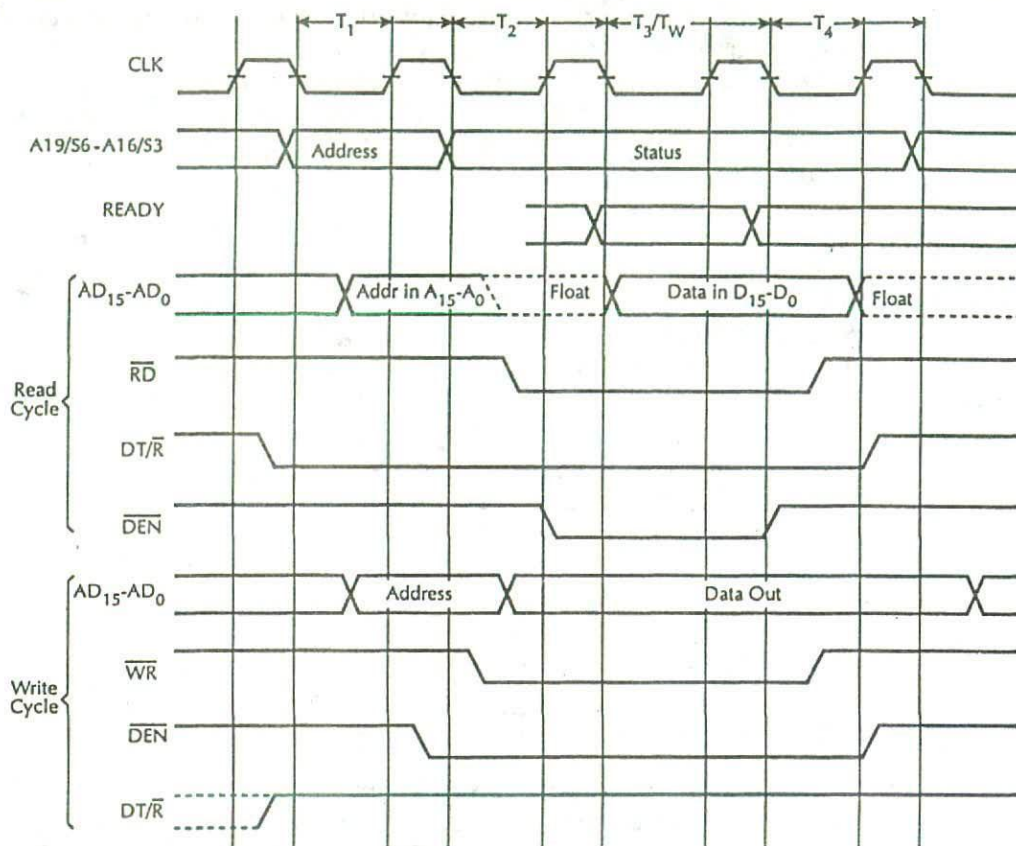
FIGURE 3.9   Basic 8086 bus cycle.

3.7.2 8086 Basic System Concepts

This section describes basic concepts associated with 8086 bus cycles, address and data bus, system data bus, and multiprocessor environment.

3.7.2.a 8086 Bus Cycle

In order to communicate with external devices via the system bus for transferring data or fetching instructions, the 8086 executes a bus cycle. The 8086 basic bus cycle timing diagram is shown in Figure 3.9. The minimum bus cycle contains four CPU clock periods called T States. The bus cycle timing diagram depicted in Figure 3.9 can be described as follows:

1. During the first T State (T1), the 8086 outputs the 20-bit address computed from a segment register and an offset on the multiplexed address/data/status bus.
2. For the second T State (T2), the 8086 removes the address from the bus and either tristates or activates the AD15-AD0 lines in preparation for reading data via AD15-AD0 lines during the T3 cycle. In case of a write bus cycle, the 8086 outputs data on AD15-AD0 lines. Also, during T2, the upper four multiplexed bus lines switch from address (A19-A16) to bus cycle status (S6, S5, S4, S3). The 8086 outputs LOW on RD (for read cycle) or WR (for write cycle) during portions of T2, T3, and T4.
3. During T3, the 8086 continues to output status information on the four A19-A16/S6-S3 lines and will either continue to output write data or input read data to or from AD15-AD0 lines. If the selected memory or I/O device is not fast enough to transfer data
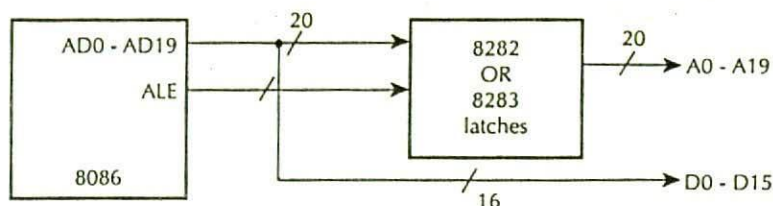
FIGURE 3.10 Separate address and data buses.

with the 8086, the memory or I/O device activates the 8086's READY input line low by the start of T3. This will force the 8086 to insert additional clock cycles (wait states TW) after T3. Bus activity during TW is the same as T3. When the selected device has had sufficient time to complete the transfer, it must activate the 8086 READY PIN HIGH. As soon as the TW clock periods end, the 8086 executes the last bus cycle, T4. The 8086 will latch data on AD15-AD0 lines during the last wait state or during T3 if no wait states are requested.

4. During T4, the 8086 disables the command lines and the selected memory and I/O devices from the bus. Thus, the bus cycle is terminated in T4. The bus cycle appears to devices in the system as an asynchronous event consisting of an address to select the device, a register or memory location within the device, a read strobe, or a write strobe along with data.

5. DEN and DT/R pins are used by the 8286/8287 transceiver in a minimum system. During the read cycle, the 8086 outputs DEN LOW during part of T2 and all of T3 cycles. This signal can be used to enable the 8286/8287 transceiver. The 8086 outputs LOW on the DT/R pin from the start of T1 and part of T4 cycles. The 8086 uses this signal to receive (read) data from the receiver during T3-T4. During a write cycle, the 8086 outputs DEN LOW during part of T1, all of T2 and T3, and part of T4 cycles. The signal can be used to enable the transceiver. The 8086 outputs HIGH on DT/R throughout the four bus cycles to transmit (write) data to the transceiver during T3-T4.

3.7.2.b 8086 Address and Data Bus Concepts

The majority of memory and I/O chips capable of interfacing to the 8086 require a stable address for the duration of the bus cycle. Therefore, the address on the 8086 multiplexed address/data bus during T1 should be latched. The latched address is then used to select the desired I/O or memory location. Note that the 8086 has a 16-bit multiplexed address and data bus, while the 8085's 8-bit data lines and LOW address byte are multiplexed. Hence, the multiplexed bus components of the 8085 family are not applicable to the 8086. To demultiplex the bus, the 8086 provides an ALE (Address Latch Enable) signal to capture the address in either the 8282 (noninverting) or 8283 (inverting) 8-bit bistable latches. These latches propagate the address through to the outputs while ALE is HIGH and latch the address in the falling edge of ALE. This only delays address access and chip select decoding by the propagation delay of the latch. Figure 3.10 shows how the 8086 demultiplexes the address and data buses.

The programmer views the 8086 memory address space as a sequence of one million bytes in which any byte may contain an eight-bit data element and any two consecutive bytes may contain a 16-bit data element. There is no constraint on byte or word addresses (boundaries). The address space is physically implemented on a 16-bit data bus by dividing the address space into two banks of up to 512K bytes as shown in Figure 3.11. These banks can be selected by BHE (Bus High Enable) and A0 as follows:
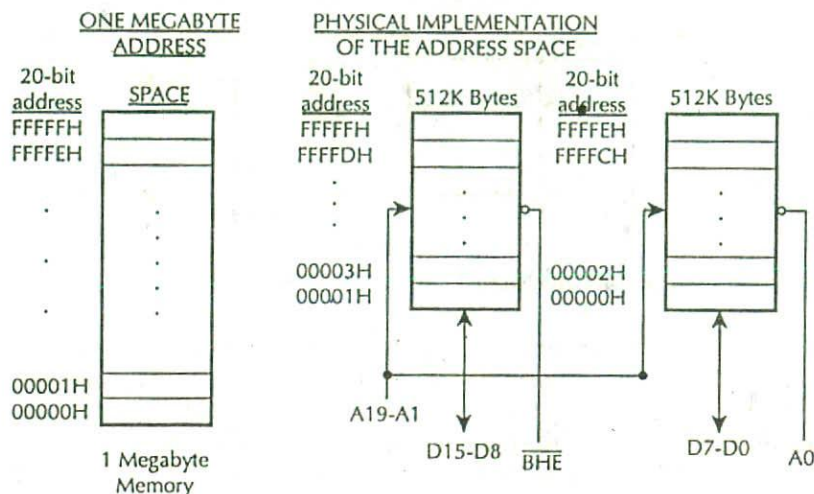
FIGURE 3.11   8086 memory.

BHE	A0	Byte transferred
0	0	Both bytes
0	1	Byte to/from odd address via D_{15}–D_8
1	0	Byte to/from even address via D_7–D_0
1	1	None

One bank is connected to D7-D0 and contains all even addressed bytes (A0 = 0). The other bank is connected to D15-D8 and contains odd-addressed bytes (A0 = 1). A particular byte in each bank is addressed by A19-A1. The even-addressed bank is enabled by LOW A0 and data bytes transferred over D7-D0 lines. The 8086 outputs HIGH on BHE (Bus High Enable) and thus disables the odd-addressed bank. The 8086 outputs LOW on BHE to select the odd-addressed bank and HIGH on A0 to disable the even-addressed bank. This directs the data transfer to the appropriate half of the data bus. Activation of A0 and BHE is performed by the 8086 depending on odd or even addresses and is transparent to the programmer. As an example, consider execution of the instruction MOV DH, [BX]. Suppose the 20-bit address computed by BX and DS is even. The 8086 outputs LOW on A0 and HIGH on BHE. This will select the even-addressed bank. The content of the selected memory is placed on the D7-D0 lines by the memory chip. The 8086 reads this data via D7-D0 and automatically places it in DH. Next, consider accessing a 16-bit word by the 8086 with low byte at an even address as shown in Figure 3.12.

For example, suppose that the 8086 executes the instruction MOV [BX], CX. Assume [BX] = 0004H, [DS] = 2000H. The 20-bit physical address for the word is 20004H. The 8086 outputs
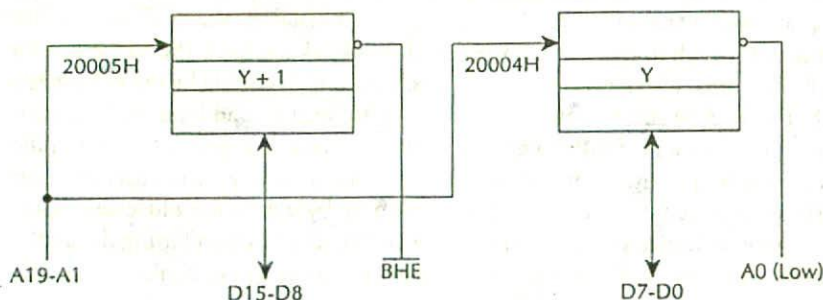
FIGURE 3.12   Even-addressed word transfer.

LOW on both A0 and BHE , enabling both banks simultaneously. The 8086 outputs [CL] to D7-D0 lines and [CH] to D15-D8 lines with WR LOW and M/IO HIGH. The enabled memory banks obtain the 16-bit data and write [CL] to location 20004H and [CH] to location 20005H. Next, consider accessing an odd-addressed 16-bit word by the 8086. For example, suppose the 20-bit physical address computed by the 8086 is 20005H. The 8086 accomplishes this transfer in two bus cycles (A19-A1 = 20005H in the first cycle, and 20006H in the second cycle). In the first bus cycle, the 8086 outputs HIGH on A0, LOW on BHE , and thus enables the odd-addressed bank and disables the even-addressed bank. The 8086 also outputs LOW on RD and HIGH on M/IO pins. In this bus cycle, the odd memory bank places [20005H] on D15-D8 lines. The 8086 reads this data into CL. In the second bus cycle, the 8086 outputs LOW on A0, HIGH on BHE , and thus enables the even-addressed bank and disables the odd-addressed bank. The 8086 also outputs LOW on RD and HIGH on M/IO pins. The selected even-addressed memory bank places [20006H] on D7-D0 lines. The 8086 reads this data into CH.

During a byte read, the 8086 floats the entire D15-D0 lines during portions of T2 cycle even though data are expected on the upper or lower half of the data bus. As will be shown later, this action simplifies the chip select decoding requirements for ROMs and EPROMs. During a byte write, the 8086 will drive the entire 16-bit data bus. The information on the half of the data bus not transferring data is indeterminate. These concepts also apply to I/O transfers.

If memory of I/O devices are directly connected to the multiplexed bus, the designer must guarantee that the devices do not corrupt the address on the bus during T1. To avoid this, the memory or I/O devices should have an output enable controlled by the 8086 read signal.

3.7.3 Interfacing with Memories

Figure 3.13 shows a general block diagram of an 8086 memory array. In Figure 3.13, the 16-bit word memory is partitioned into high and low 8-bit banks on the upper and lower halves of the data bus selected by BHE and A0.

3.7.3.a ROM and EPROM

ROMs and EPROMs are the simplest memory chips to interface to the 8086. Since ROMs and EPROMs are read-only devices, A0 and BHE are not required to be part of the chip enable/ select decoding (chip enable is similar to chip select except chip enable also provides whether
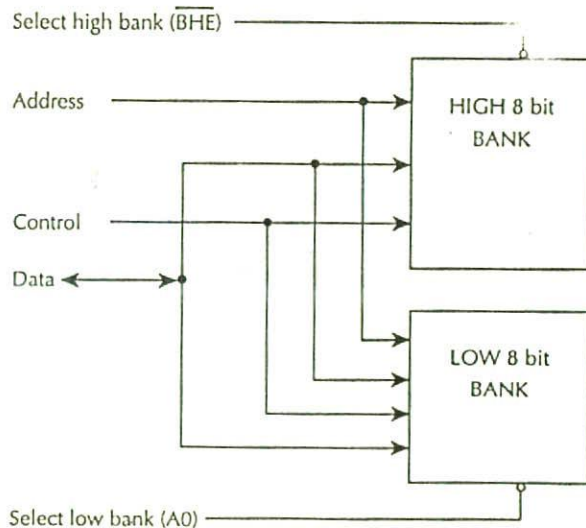
FIGURE 3.13   8086 memory array.

the chip is in active or standby power mode). The 8086 address lines must be connected to the address pins of the ROM/EPROM chips starting with A1 and higher to all the address lines of the ROM/EPROM chips. The 8086 unused address lines can be used as chip enable/select decoding. To interface the ROMs/RAMs directly to the 8086 multiplexed bus, they must have output enable signals.

Byte accesses are obtained by reading the full 16-bit word onto the bus with the 8086 discarding the unwanted byte and accepting the desired byte. If RD, WR, and M/IO are not decoded to generate separate memory and I/O commands for memory and I/O chips and the I/O space overlaps with the memory space of ROM/EPROM, then M/IO must be a condition of chip select decode.

3.7.3.b Static RAMs

Since static RAMs are read/write memories, both A0 and $\overline{\text{BHE}}$ must be included in the chip select/chip enable decoding of the devices and write timing must be considered in the compatibility analysis.

For each static RAM (containing odd or even addresses), the memory data lines must be connected to either the upper half AD15-AD8 for static RAM with odd addresses or lower half AD7-AD0 of the 8086 data lines for static RAM with even addresses.

3.7.3.c Dynamic RAMs

Dynamic RAMs store information as charges in capacitors. Since capacitors can hold charges for a few milliseconds, refresh circuitry is necessary in dynamic RAMs for retaining these charges. Therefore, dynamic RAMs are complex devices to design into a system. To relieve the designer of most of these complicated interfacing tasks, Intel provides the 8202 dynamic RAM controller as part of the 8086 family of peripheral devices. The 8202 can be interfaced with the 8086 to build a dynamic memory system. A thorough discussion on this topic can be found in the Intel manuals.

3.7.4 8086 Programmed I/O

The 8086 can be interfaced to 8- and 16-bit I/O devices using either standard or memory-mapped I/O. The standard I/O uses the instructions IN and OUT and is capable of providing 64K bytes of I/O ports. Using standard I/O, the 8086 can transfer 8- or 16-bit data to or from a peripheral device. The 64K byte I/O locations can then be configured as 64K 8-bit ports or 32K 16-bit ports. All I/O transfer between the 8086 and the peripheral devices take place via AL for 8-bit ports (AH is not involved) and AX for 16-bit ports. The I/O port addressing can be done either directly or indirectly as follows:

DIRECT
- IN AL, PORTA or IN AX, PORTB inputs 8-bit contents of port A into AL or 16-bit contents of port B into AX, respectively.

  Port A and port B are assumed as 8- and 16-bit ports, respectively.
- OUT PORTA, AL or OUT PORT B, AX outputs 8-bit contents of AL into port A or 16-bit contents of AX into port B, respectively.

INDIRECT
- IN AX, DX or IN AL, DX inputs 16-bit data addressed by DX into AX or 8-bit data addressed by DX into AL, respectively.
- OUT DX, AX or OUT DX, AL outputs 16-bit contents of AX into the port addressed by DX or 8-bit contents of AL into the port addressed by DX, respectively. In indirect addressing, register DX is used to hold the port address.

Data transfer using the memory-mapped I/O is accomplished by using memory-oriented instructions such as MOV reg 8 or reg 16, [BX] and MOV [BX], reg 8 or reg 16 for inputting and outputting 8- or 16-bit data from or to an 8-bit register or a 16-bit register addressed by the 20-bit memory-mapped port location computed from DS and BX.

Note that the indirect I/O transfer method is desirable for service routines that handle more than one device such as multiple printers by allowing the desired device (a specific printer) to be passed to the procedure as a parameter.

Devices with I/O ports can be connected to either the upper or lower half of the data bus. If the I/O port chip is connected to the 8086 lower half of the data lines (AD0-AD7), the port addresses will be even (A0 = 0). On the other hand, the port addresses will be odd ($\overline{\text{BHE}}$ = 0) if the I/O port chip is connected to the upper half of the 8086 data lines (AD8-AD15).

3.8 8086-Based Microcomputer

In this section, an 8086 will be interfaced in the minimum mode to provide 2K × 16 EPROM, 2K × 16 static RAM, and six 8-bit I/O ports. 2716 EPROM, 6116 static RAM, and 8255 I/O chips will be used for this purpose. Memory and I/O maps will also be determined. Figure 3.14 shows a hardware schematic for accomplishing this.

Three 74LS373 latches are used. The 2716 is a 2K × 8 ultraviolet EPROM with eleven address pins A0-A10 and eight data pins O0-7. Two 2716s are used. The 8086 A1-A11pins are connected to the A0-A10 pins of these chips. The 2716 even EPROM's O0-7 pins are connected to the 8086 D0-D7 pins. This is because the 8086 reads data via the D0-D7 pins for even addresses. On the other hand, the O0-7 pins of the odd 2716 are connected to the 8086 D8-D15 pins. The 8086 reads data via D8-D15 pins for odd addresses.

The 6116 is a 2KX8 SRAM. Two 6116s are used.

Table 3.2 shows memory and I/O maps of the 8086-based microcomputer of Figure 3.14. Note that the reset vector FFFF0H (CS = FFFFH, IP = 0000H) is included in the 2716.

For I/O ports, two 8255 chips are used. The 8255 is a general-purpose programmable I/O chip. The 8255 has three 8-bit I/O ports: ports A, B, and C. Ports A and B are latched 8-bit ports for both input and output. Port C is also an 8-bit port with latched output but the inputs are not latched. Port C can be used in two ways. It can either be used as a simple I/O port or as a control port for data transfer using handshaking via ports A and B.

The 8086 can configure the three ports by outputting appropriate data to the 8-bit control register. The ports can be decoded by two 8255 input pins A0 and A1 as follows:

A1	A0	
0	0	Port A
0	1	Port B
1	0	Port C
1	1	Control register

The structure of the control register is given below:
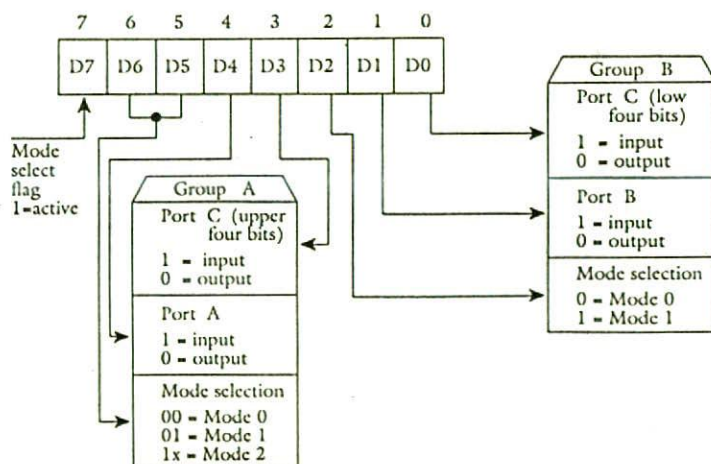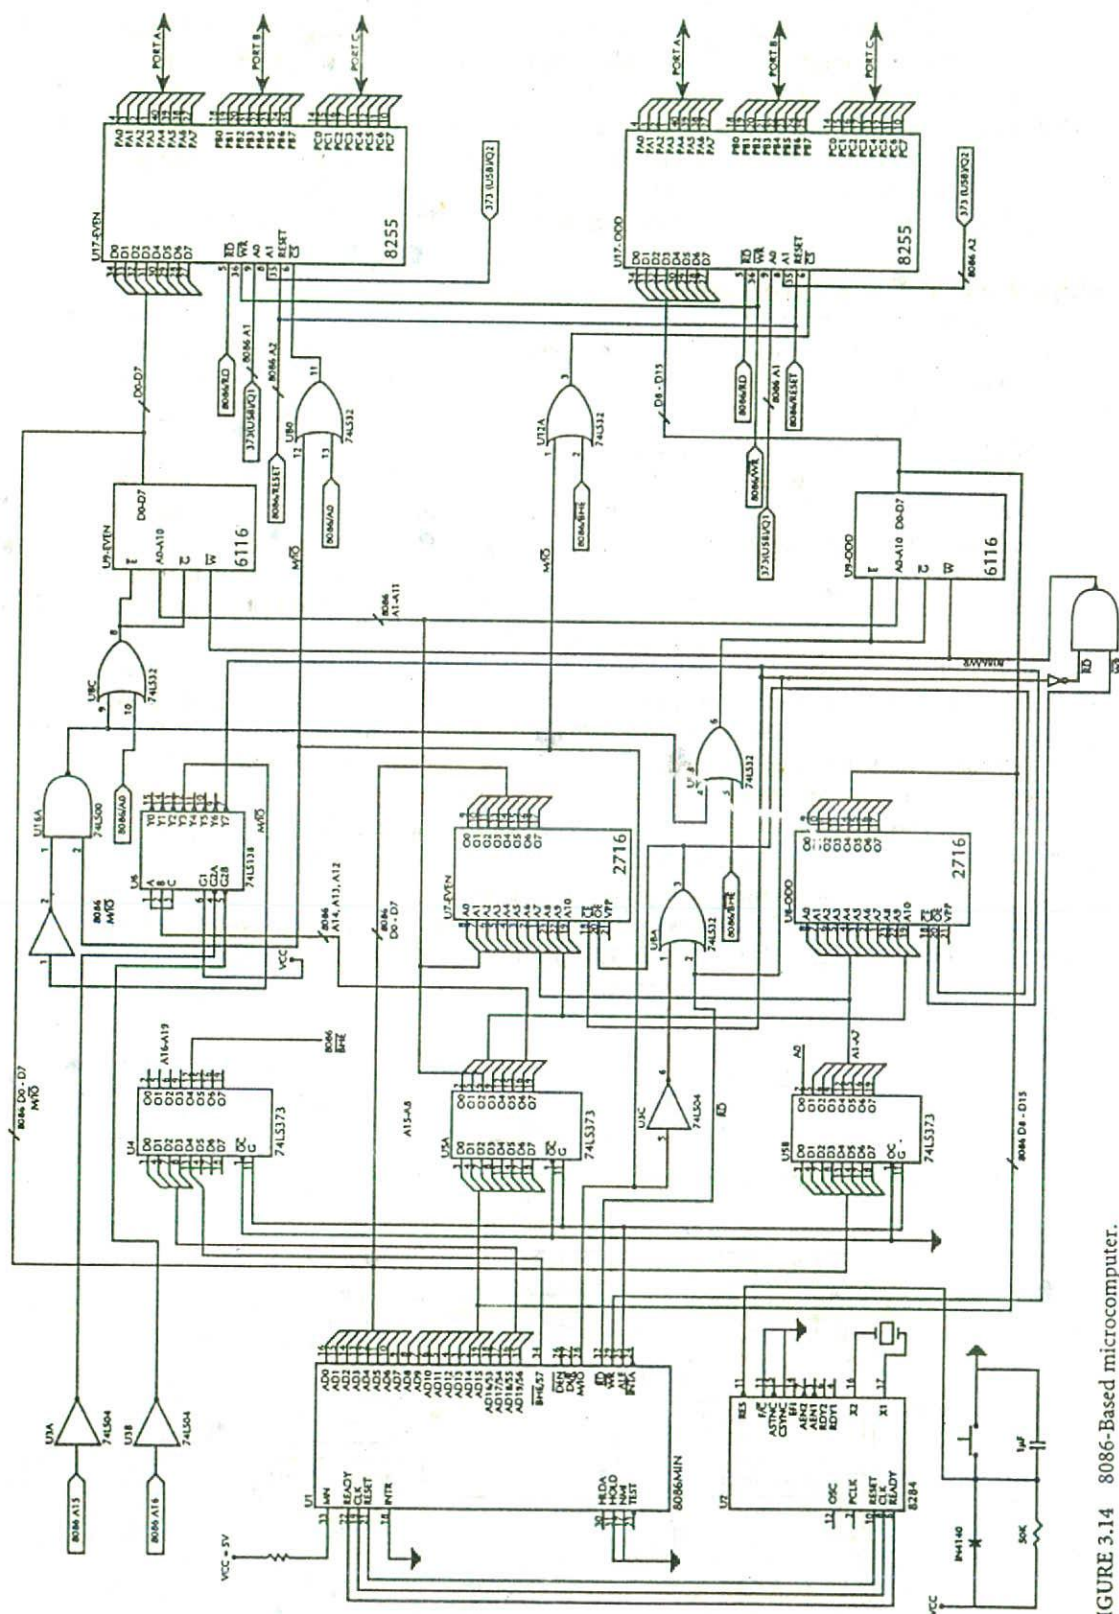
FIGURE 3.14   8086-Based microcomputer.

TABLE 3.2   Memory and I/O Maps

(a) Physical Memory Map

2716 (U7) EVEN # EPROM

A19	A18	A17	A16	A15	A14	A13	A12	A11	A10	A9	A8	A7	A6	A5	A4	A3	A2	A1	A0
1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0
										TO									
1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0
don't cares asumed high			to enable the 74LS138			output line 7 of the 74LS138													
FF000H, FF002H, … FFFFEH																			

2716 (U8) ODD # EPROM

A19	A18	A17	A16	A15	A14	A13	A12	A11	A10	A9	A8	A7	A6	A5	A4	A3	A2	A1	A0
1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	1
										TO									
1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
don't cares asumed high			to enable the 74LS138			output line 7 of the 74LS138													
FF001H, FF003H, … FFFFFH																			

6116 (U9) EVEN # RAM

A19	A18	A17	A16	A15	A14	A13	A12	A11	A10	A9	A8	A7	A6	A5	A4	A3	A2	A1	A0
1	1	1	1	1	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0
										TO									
1	1	1	1	1	0	1	1	1	1	1	1	1	1	1	1	1	1	1	0
don't cares asumed high			to enable the 74LS138			output line 3 of the 74LS138													
FB000H, FB002H, … FBFFEH																			

6116 (U10) ODD # RAM

A19	A18	A17	A16	A15	A14	A13	A12	A11	A10	A9	A8	A7	A6	A5	A4	A3	A2	A1	A0
1	1	1	1	1	0	1	1	0	0	0	0	0	0	0	0	0	0	0	1
										TO									
1	1	1	1	1	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1
don't cares asumed high			to enable the 74LS138			output line 3 of the 74LS138													
FB001H, FB003H, … FBFFFH																			

TABLE 3.2   Memory and I/O Maps (*continued*)

(b) Logical Memory Map

Chip	Segment Value	Offset
Even 2716	FF00H	0000H, 0002H, 0004H, ..., 0FFEH
Odd 2716	FF00H	0001H, 0003H, 0005H, ..., 0FFFH
Even 6116	FB00H	0000H, 0002H, 0004H, ..., 0FFEH
Odd 6116	FB00H	0001H, 0003H, 0005H, ..., 0FFFH

I/O Map	
Even 8255	Port A = F8H
	Port B = FAH
	Port C = FCH
	CSR = FEH
Odd 8255	Port A = F9H
	Port B = FBH
	Port C = FDH
	CSR = FFH

The bit 7 (D7) of the control register must be one to send the above definitions for bits 0 through 6 (D0-D6).

In this format, bits D0-D6 are divided into two groups: groups A and B. Group A configures all 8 bits of port A and upper 4 bits of port C, while Group B defines all 8 bits of port B and lower 4 bits of Port C. All bits in a port can be configured as a parallel input port by writing a 1 at the appropriate bit in the control register by the 8086 OUT instructions, and a 0 to a particular bit position will configure the appropriate port as a parallel output port. Group A has three modes of operation. These are modes 0, 1, and 2. Group B has two modes: modes 0 and 1. Mode 0 for both groups A and B provides simple I/O operation for each of the three ports. No handshaking is required. Mode 1 for both groups A and B is the strobed I/O mode used for transferring I/O data to or from a specified port in conjuction with strobes or handshaking signals. Ports A and B use the lines on port C to generate or accept these handshaking signals.

The mode 2 of group A is the strobed bidirectional bus I/O and may be used for communicating with a peripheral device on a single 8-bit data bus for both transmitting and receiving data (bidirectional bus I/O). Handshaking signals are required. Interrupt generation and enable/disable functions are also available.

When D7 = 0, the bit set/reset control word format is used for the control register as follows:
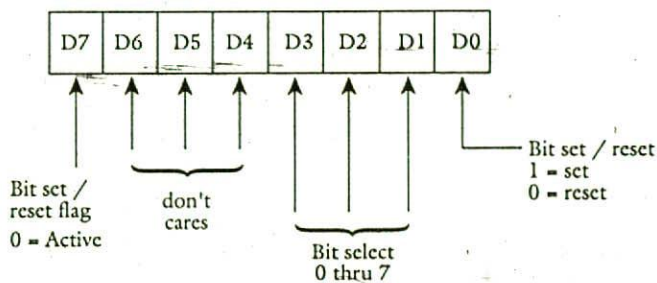
This format is used to set or reset the output on a pin of port C or when enabling the interrupt output signals for handshake data transfer. For example, the following 8-bits

$$
\begin{array}{cccc}
0 & X\ X\ X & 1\ 1\ 0 & 0 \\
\uparrow & \text{don't} & \text{bit 6} & \uparrow \\
\text{bit set/} & \text{care} & & \text{clear} \\
\text{reset mode} & & &
\end{array}
$$

will clear bit 6 of port C to zero. Note that the control word format can be output to the 8255 control register by using the 8086 OUT instruction.

Next, only mode 0 of the 8255 will only be considered in the following discussion to illustrate 8086's programmed I/O capability.

Now, let us define the control word format for mode 0 more precisely by means of numerical example. Consider that the control word format is 1 000 0 010$_2$. With this data in the control register, all 8 bits of port A are configured as outputs, 8 bits of port C are also configured as outputs, but all 8 bits of port B are defined as inputs. On the other hand, outputting 1 0 011 011$_2$ into the control register will configure all 3 8-bit ports (ports A, B, and C) as inputs.

Let us decode the I/O port addresses. One 8255 will contain the odd-addressed ports since it is enabled by $\overline{\text{BHE}}$, while the other 8255 will include the even-addressed ports since it is enabled by A0.

Since the 8086 A1 and A2 pins are utilized in addressing the ports, bits A3-A7 are don't cares and are assumed to be ones here. Note that A0 = 1 for odd-addressed ports, while A0 = 0 for even-addressed ports.

For 8255-1 (Odd-Addressed Ports)

Port name	A7	A6	A5	A4	A3	A2	A1	A0	
Port A	1	1	1	1	1	0	0	1	= F9H
Port B	1	1	1	1	1	0	1	1	= FBH
Port C	1	1	1	1	1	1	0	1	= FDH
Control Register	1	1	1	1	1	1	1	1	= FFH

For 8255-2 (Even-Addressed Ports)

Port name	A7	A6	A5	A4	A3	A2	A1	A0	
Port A	1	1	1	1	1	0	0	0	= F8H
Port B	1	1	1	1	1	0	1	0	= FAH
Port C	1	1	1	1	1	1	0	0	= FCH
Control Register	1	1	1	1	1	1	1	0	= FEH

In the above, standard I/O technique is used. The 8255s can also be interfaced to the 8086 using memory-mapped I/O. In this case the 8086 M/IO pin will not be used. The 20-bit physical addresses for the ports can be determined in a similar way by considering any unused 8086 address bits (A3-A19) as don't cares.

From the above discussion, the following points can be summarized:

1. For ROMs and EPROMs, $\overline{BHE}$ and A0 are not required to be part of chip enable/select decoding.
2. For RAMs and I/O chips, both $\overline{BHE}$ and A0 must be used in chip select logic.
3. For ROMs/EPROMs and RAMs, both odd and even addressed chips are required. However, for I/O chips, either an odd-addressed I/O chip or even-addressed I/O chip or both can be used depending on the number of ports required in an application.
4. For interfacing ROMs/EPROMs and RAMs to the 8086, the same chip select logic must be used for both even and its corresponding odd memory bank. For example, the even memory chip containing address 00000H, 00002H, 00004H, ..., must have the same chip select logic as its odd counterpart containing address 00001H, 00003H, 00005H,.... The memory map of the two memory banks will be distinguished by the value of A0 (A0 = 0 for even addresses and A0 = 1 for odd addresses).

Example 3.11

Assume 8086/8255 based microcomputer. Write an 8086 assembly language program to drive an LED connected to bit 2 of Port B based on two switch inputs at bits 6 and 7 of Port A. If both switch inputs are HIGH or LOW, turn the LED "ON"; otherwise turn the LED "OFF".

Solution

```
              TITLE     PROGRAM
1000          ORG       1000H
0000          PROG      SEGMENT
                        ASSUME  CS:PROG
=0001         PORTA     EQU  01H
```

```
=0003          PORTB    EQU 03H
=0007          CNTRL    EQU 07H
0000 B0 90              MOV AL,90H
0002 E6 07              OUT CNTRL,AL
0004 E4 01     BEGIN:   IN AL,PORTA
0006 24 C0              AND AL,0C0H
0008 7A 06              JPE LEDON
000A B0 00              MOV AL,00H
000C E6 03              OUT PORTB,AL
000E EB F4              JMP BEGIN
0010 B0 04     LEDON:   MOV AL,04H
0012 E6 03              OUT PORTB,AL
0014 EB EE              JMP BEGIN
0016           PROG     ENDS
                        END
```

Microsoft (R) Macro Assembler Version 5.10 10/23/93 21:58:39
PROGRAM   Symbols-1

Segments and Groups:

Name	Length	Align	Combine Class
PROG	0016	PARA	NONE

Symbols:

Name	Type	Value	Attr
BEGIN	L NEAR	0004	PROG
CNTRL	NUMBER	0007	
LEDON	L NEAR	0010	PROG
PORTA	NUMBER	0001	
PORTB	NUMBER	0003	
@CPU	TEXT	0101h	
@FILENAME		TEXT	test1
@VERSION		TEXT	510

```
    19 Source Lines
    19 Total Lines
    11 Symbols
```

47884 + 433567 Bytes symbol space free

```
    0 Warning Errors
    0 Severe Errors
```

Example 3.12

Assume an 8086/8255 configuration with the following port addresses:

Port A   = F9H
Port B   = FBH
Control Register = FFH

Port A has three switches connected to bits 0, 1, and 2 and port B has an LED connected to bit 2 as follows.

Write an 8086 assembly language program to turn the LED ON if port A has an odd number of HIGH switch inputs; otherwise turn the LED OFF. Do not use any instructions involving the parity flag.

Solution

```
1000                ORG    1000H
0000                CODE   SEGMENT
                           ASSUME CS:CODE
=00F9               PORTA  EQU  0F9H
=00FB               PORTB  EQU  0FBH
=00FF               CNTRL  EQU  0FFH
0000 B0 90                 MOV  AL, 90H    ; Configure Port A as
                                           ; input
0002 E6 FF                 OUT  CNTRL, AL  ; Port B as output
0004 E4 F9          BEGIN: IN   AL, PORTA  ; Input switches
0006 24 07                 AND  AL, 07H    ; Mask high five bits
0008 3C 07                 CMP  AL, 07H    ; Are all three inputs
                                           ; HIGH?
000A 74 12                 JZ   ODD        ; If so, turn LEDON
000C 3C 01                 CMP  AL, 01H    ; Is only input 0 HIGH?
000E 74 0E                 JZ   ODD        ; If so, turn LED ON
0010 3C 02                 CMP  AL, 02H    ; Is only input 1 HIGH?
0012 74 0A                 JZ   ODD        ; If so, turn LED ON
0014 3C 04                 CMP  AL, 04H    ; Is only input 2 HIGH?
0016 74 06                 JZ   ODD        ; If so, turn LED ON
0018 B0 00                 MOV  AL, 00H    ; Else turn LED
001A E6 FB                 OUT  PORTB, AL  ; OFF
001C EB E6                 JMP  BEGIN      ; REPEAT
001E B0 04          ODD:   MOV  AL, 04H    ; Turn LED
0020 E6 FB                 OUT  PORTB, AL  ; ON
0022 EB E0                 JMP  BEGIN      ; REPEAT
0024                CODE   ENDS
                           END
```

Microsoft (R) Macro Assembler Version 5.10   11/14/93 16:40:4
                                             Symbols-1

Segments and Groups:

Name	Length	Align	Combine	Class
CODE	0024	PARA	NONE	

Symbols:

Name	Type	Value	Attr
BEGIN	L NEAR	0004	CODE
CNTRL	NUMBER	00FF	
ODD	L NEAR	001E	CODE
PORTA	NUMBER	00F9	

```
PORTB  . . . . . . . . . . .     NUMBER   00FB
@CPU  . . . . . . . . . . . .    TEXT     0101h
@FILENAME  . . . . . . .                  TEXT   examp17
@VERSION  . . . . . . . .                 TEXT   510
                                                         *
```

```
     26  Source Lines
     26  Total Lines
     11  Symbols
```

```
47576 + 427923 Bytes symbol space free
```

```
      0  Warning Errors
      0  Severe Errors
```

3.9 8086 INTERRUPT SYSTEM

The 8086 interrupts can be classified into three types. These are

1. Predefined interrupts
2. User-defined software interrupts
3. User-defined hardware interrupts

The interrupt vector addresses for all the 8086 interrupts are determined from a table stored in locations 00000H through 003FEH. The starting addresses for the service routines for the interrupts are obtained by the 8086 using this table. Four bytes of the table are assigned to each interrupt: two bytes for IP and two bytes for CS. The table may contain up to 256 32-bit vectors. If fewer than 256 interrupts are defined in the system, the user is required to provide enough memory for the interrupt pointer table for obtaining the defined interrupts.

The 8086 assigns every interrupt a type code for identifying the interrupt. There are 256 type codes associated with the 256 table entries. Each entry consists of two word addresses, one for storing the IP contents and the other for storing the CS contents. Each 8086 interrupt physical address vector is 20 bits wide and is computed from the 16-bit contents of IP and CS.

For obtaining an interrupt address vector, the 8086 uses the two addresses in the pointer table where IP and CS are stored for a particular interrupt type.

For example, for the interrupt type nn (instruction INT nn), the table address for IP = 4 * nn and the table address for CS = (4 * nn) + 2. For servicing the 8086's nonmaskable interrupt (NMI pin), the 8086 assigns the type code 2 to this interrupt. The 8086 automatically executes the INT2 instruction internally to obtain the interrupt address vector as follows:

$$\text{Address for IP} = 4 * 2 = 00008H$$
$$\text{Address for CS} = (4 * 2) + 2 = 0000AH$$

The 8086 loads the values of IP and CS from the 20-bit physical addresses 00008H and 0000AH in the pointer table. The user must store the desired 16-bit values of IP and CS in these locations. Similarly, the IP and CS values for other interrupts are calculated. The 8086 interrupt pointer table layout is shown in Table 3.3.

In response to an interrupt, the 8086 pushes flags, CS, and IP onto the stack, clears TF and IF flags, and then loads IP and CS from the pointer table using the type code.

Interrupt service routines must be terminated with the IRET (Interrupt Return) instruction which pops the top three stack words into IP, CS, and flags, thus returning to the right place in the main program. The 256 interrupt type codes are assigned as follows:

TABLE 3.3   8086 Interrupt Pointer Table

Interrupt type code		20-bit Memory Address
0	IP	00000H
	CS	00002H
1	IP	00004H
	CS	00006H
2		00008H
		0000AH
	.	.
	.	.
	.	.
	.	.
	.	.
	.	.
255	IP	003FCH
	CS	003FEH

- Types 0 to 4 are for the predefined interrupts.
- Types 5 to 31 are reserved by Intel for future use.
- Types 32 to 255 are available for maskable interrupts.

3.9.1 Predefined Interrupts (0 to 4)

The predefined interrupts include DIVISION BY ZERO (type 0), SINGLE STEP (type 1) NONMASKABLE INTERRUPT pin (type 2), BREAKPOINT-INTERRUPT (type 3), and INTERRUPT ON OVERFLOW (type 4). The user must provide the desired IP and CS values in the interrupt pointer table. The user may also imitate these interrupts through hardware or software. If a predefined interrupt is not used in a system, the user may assign some other function to the associated type.

The 8086 is automatically interrupted whenever a division by zero is attempted. This interrupt is nonmaskable and is implemented by Intel as part of the execution of the divide instruction. When the TF (TRAP flag) is set by an instruction, the 8086 goes into the single step mode. The TF can be set to one as follows:

```
PUSHF                 ;  Save flags
MOV BP, SP            ;  Move [SP] to [BP]
OR [BP + 0], 0100H    ;  Set TF
POPF                  ;  Pop flags
```

Note that in the above [BP + 0] rather than [BP] is used since BP cannot be used without displacement.

Once TF is set to one, the 8086 automatically generates a TYPE 1 interrupt after execution of each instruction. The user can write a service routine at the interrupt address vector to display memory locations and/or register to debug a program. Single step is nonmaskable and cannot be enabled by STI (enable interrupt) or disabled by CLI (disable interrupt) instruction.

The nonmaskable interrupt is initiated via the 8086 NMI pin. It is edge triggered (LOW to HIGH) and must be active for two clock cycles to guarantee recognition. It is normally used

for catastrophic failures such as power failure. The 8086 obtains the interrupt vector address by automatically executing the INT2 (type 2) instruction internally.

Type 3 interrupt is used for breakpoint and is nonmaskable. The user inserts the one-byte instruction INT3 into a program by replacing an instruction. Breakpoints are useful for program debugging.

The INTERRUPT ON OVERFLOW is a type 4 interrupt. This interrupt occurs if the overflow flag (OF) is set and the INTO instruction is executed. The overflow flag is affected, for example, after execution of signed arithmetic such as IMUL (signed multiplication) instruction. The user can execute the INTO instruction after the IMUL. If there is an overflow, an error service routine written by the user at the type 4 interrupt address vector is executed.

3.9.2 User-Defined Software Interrupts

The user can generate an interrupt by executing a two-byte interrupt instruction INT nn. The INT nn instruction is not maskable by the interrupt enable flag (IF). The INT nn instruction can be used to test an interrupt service routine for external interrupts. Type codes 0 to 255 can be used. If predefined interrupt is not used in a system, the associated type code can be utilized with the INT nn instruction to generate software (internal) interrupts.

3.9.3 User-Defined Hardware (Maskable Interrupts, Type Codes 32_{10}–255_{10})

The 8086 maskable interrupts are initiated via the INTR pin. These interrupts can be enabled or disabled by STI (IF = 1) or CLI (IF = 0), respectively. If IF = 1 and INTR is active (HIGH) without occurrence of any other interrupts, the 8086, after completing the current instruction, generates INTA LOW twice, each time for about 2 cycles.

The state of the INTR pin is sampled during the last clock cycle of each instruction. In some instances, the 8086 samples the INTR pin at a later time. An example is execution of POP to a segment register. In this case, the interrupts are sampled until completion of the following instruction. This allows a 32-bit pointer to be loaded to SS and SP without the danger of an interrupt occurring between the two loads.

INTA is only generated by the 8086 in response to INTR, as shown in Figure 3.15. The interrupt acknowledge sequence includes two INTA cycles separated by two idle clock cycles. ALE is also generated by the 8086 and will load the address latches with indeterminate
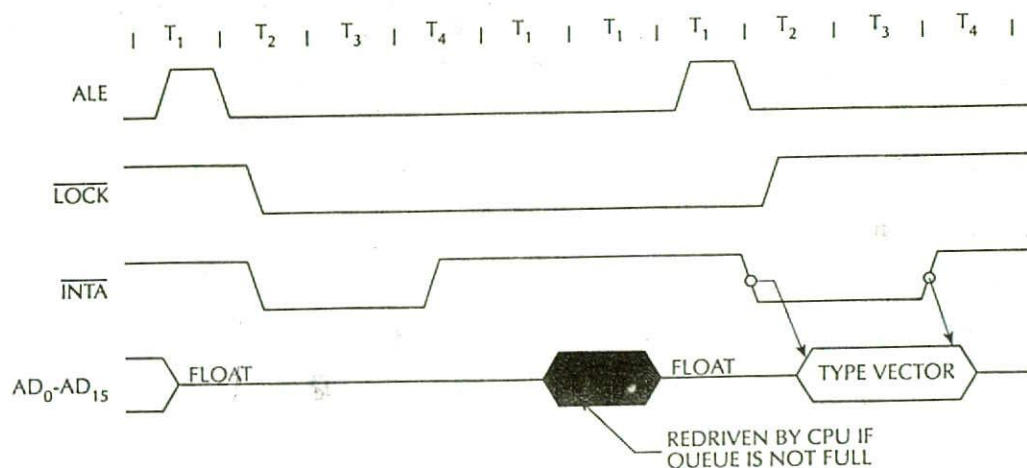

FIGURE 3.15   Interrupt acknowledge sequence.

FIGURE 3.16   Servicing the INTR in the minimum mode.

information. The ALE is useful in maximum systems with multiple 8259A priority interrupt controllers. During the INTA bus cycles, DT/R and DEN are LOW (see 8086 minimum mode bus cycle). The first INTA bus cycle indicates that an interrupt acknowledge cycle is in progress and allows the system to be ready to place the interrupt type code on the next INTA bus cycle. The 8086 does not obtain the information from the bus during the first cycle. The external hardware must place the type code on the lower half of the 16-bit data bus during the second cycle.

In the minimum mode, the M/IO is low indicating I/O operation during the INTA bus cycles. The 8086 internal LOCK signal is also low from T2 of the first bus cycle until T2 of the second bus cycle to avoid the BIU from accepting a hold request between the two INTA cycles. Figure 3.16 shows a simplified interconnection between the 8086 and 74LS244 for servicing the INTR. INTA enables 74LS244 to place the type code nn on the 8086 data bus.

In the maximum mode, the status lines S0 to S2 will enable the INTA output for each cycle via the 8288. The 8086 LOCK output will be active from T2 of the first cycle until T2 of the second to prevent the 8086 from accepting a hold request on either RQ/GT input and to prevent bus arbitration logic from releasing the bus between INTAs in multimaster systems. The LOCK output can be used in external logic to lock other devices off the system bus, thus ensuring the INTA sequence to be completed without intervention.

Once the 8086 has the interrupt-type code (via the bus for hardware interrupts, from software interrupt instructions INT nn or from the predefined interrupts), the type code is multiplied by four to obtain the corresponding interrupt vector in the interrupt vector table. The four bytes of the interrupt vector are least significant byte of the instruction pointer, most significant byte of the pointer, least significant byte of the code segment register; and most significant byte of the code segment register. During the transfer of control, the 8086 pushes the flags and current code segment register and instruction pointer into the stack. Flags TF and IF are then cleared to zero. The CS and IP values are read by the 8086 from the interrupt vector table. No segment registers are used when accessing the interrupt pointer table. S4, S3 has the value 10_2 to indicate no segment register selection.

As far as the 8086 interrupt priorities are concerned, single-step interrupt has the highest priority, followed by NMI, followed by the software interrupts (all interrupts except single step, NMI, and INTR interrupts). This means that a simultaneous NMI and single step will cause the NMI service routine to follow single step; a simultaneous software interrupt and single step will cause the software interrupt service routine to follow single step and a simultaneous NMI and software interrupt will cause the NMI service routine to be executed prior to the software interrupt service routine. An exception to this priority scheme occurs if all three nonmaskable interrupts (single step, software, and NMI) are pending. For this case, software interrupt service routine will be executed first followed by the NMI service routine, and single stepping will not be serviced. However, if software interrupt and single stepping are pending,

single stepping resumes upon execution of the instruction causing the software interrupt (the next instruction in the routine being single stepped).

The INTR is maskable and has the lowest priority. If the user does not wish to single step before INTR is serviced, the single-step routine must disable interrupts during execution of the program being single stepped, and reenable interrupts on entry to the single-step routine. To avoid single stepping before the NMI service routine, the single-step routine must check the return address on the stack for the NMI service routine address and return control to that routine without single step enabled.

A priority interrupt controller such as the 8259A can be used with the 8086 INTR to provide eight levels of hardware interrupts. The 8259A has built-in features for expansion of up to 64 levels with additional 8259As. The 8259A is programmable and can be readily used with the 8086 to obtain multiple interrupts from the single 8086 INTR pin.

Example 3.13
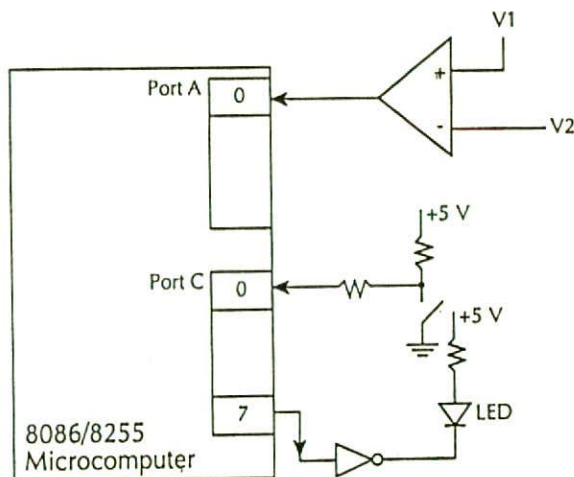
i)   In the above, the 8086/8255 microcomputer is required to perform the following:

If V1 > V2, turn the LED "ON" if the switch is open. If V1 ≤ V2, loop and test again.

Write an 8086 assembly language program to accomplish the above.

ii)   Repeat part i) using NMI and INTR interrupts. Write service routine at IP = 4000H, CS = 3000H. Assume that the stack pointer and stack segment are already initialized.

Solution

i)

```
0000            JOHN   SEGMENT AT 1000H
                       ASSUME CS: JOHN
=0001           PORTA  EQU  01H
=0005           PORTC  EQU  05H
=0007           CNTRL  EQU  07H
0000 B0 91             MOV  AL, 91H        ;  Configure
0002 E6 04             OUT  CNTRL,AL       ;  Port A and Port C
```

```
0004 E4 01    START:   IN AL, PORTA      ; Input comparator
0006 24 01             AND AL, 01H       ; Check if high
0008 74 FA             JZ START          ; Jump back if low
000A E4 03             IN AL, PORTC      ; Input switch
000C B1 02             MOV CL, 2
000E D2 D8             RCR AL, CL
0010 E6 03             OUT PORTC, AL     ; Output to LED
0012 F4                HLT
0013          JOHN     ENDS
                       END
```

Microsoft (R) Macro Assembler Version 5.10   11/17/93 22:35:4
                                             Symbols-1

Segments and Groups:

Name	Length	Align	Combine Class
JOHN	0013	AT	1000

Symbols:

Name	Type	Value	Attr
CNTRL	NUMBER	0007	
PORTA	NUMBER	0001	
PORTC	NUMBER	0005	
START	L NEAR	0004	JOHN
@CPU	TEXT	0101h	
@FILENAME	TEXT	EXP318A	
@VERSION	TEXT	510	

```
    17 Source Lines
    17 Total Lines
    10 Symbols

47576 + 427939 Bytes symbol space free

     0 Warning Errors
     0 Severe Errors
```

ii)
  Using NMI
    Main program

```
0000          JOHN     SEGMENT AT 2000H
                       ASSUME CS: JOHN
=0001         PORTA EQU 01H
=0005         PORTC EQU 05H
=0007         CNTRL EQU 07H
0000 B0 91             MOV AL, 91H       ; Configure
0002 E6 07             OUT CNTRL,AL      ; Port A and Port C
0004 9B       WAIT:    JMP WAIT
0005 F4                HLT               ; Wait for interrupt
0006          JOHN     ENDS
                       END
```

Service routine

```
0008                              ORG              00000008H
0008  40003000                    DD               40003000H
30004000                          ORG              30004000H
30004000  E4 03    START:         IN AL, PORTC    ;  Input switch
30004002  B1 02                   MOV CL, 2
30004004  D2 D8                   RCR AL, CL
30004006  E6 03                   OUT PORTC, AL   ;  Output to LED
30004008  CF                      IRET
```

Microsoft (R) Macro Assembler Version 5.10 11/17/93 22:18:20
                                                  Symbols-1

Segments and Groups:

Name	Length	Align	Combine Class
JOHN	0006	AT	2000

Symbols:

Name	Type	Value	Attr
CNTRL	NUMBER	0007	
PORTA	NUMBER	0001	
PORTC	NUMBER	0005	
@CPU	TEXT	0101h	
@FILENAME		TEXT	test
@VERSION		TEXT	510

```
   11 Source Lines
   11 Total Lines
    9 Symbols
```
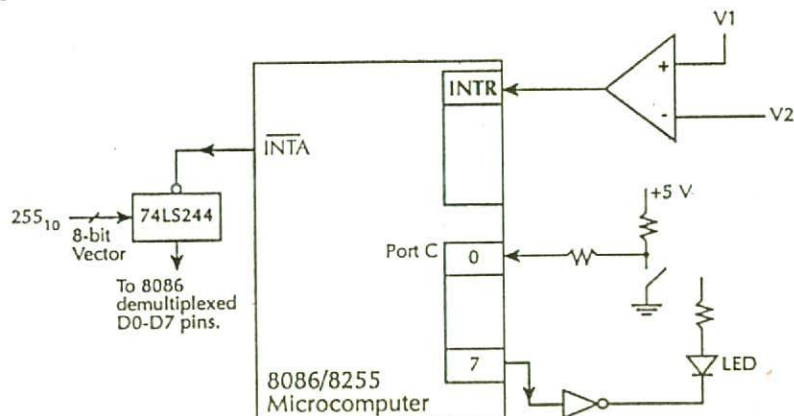
47884 + 433567 Bytes symbol space free

```
    0 Warning Errors
    0 Severe Errors
```

Using INTR (vector 256_{10})

Main program

```
0000            JOHN    SEGMENT AT 3000H
                        ASSUME CS: JOHN
=0001           PORTA   EQU 01H
=0005           PORTC   EQU 05H
=0007           CNTRL   EQU 07H
0000 FB                 STI                  ; Enable maskable
                                             ; interrupt
0001 B0 91              MOV AL, 91H          ; Configure Port A and
0003 E6 04              OUT CNTRL, AL        ; Port C
                WAIT:   JMP WAIT             ; Wait for interrupt
0005 F4                 HLT
0006            JOHN    ENDS
                        END
```

Service routine

```
03FC                            ORG    000003FCH
03FC 40003000                   DD     40003000H
30004000                        ORG    30004000H
30004000 E4 03                  IN AL, PORTC    ; Input switch
30004002 B1 02                  MOV CL, 2
30004004 D2 D8                  RCR AL, CL
30004006 E6 03                  OUT PORTC, AL   ; Output to LED
30004008 CF                     IRET
```

Microsoft (R) Macro Assembler Version 5.10   11/17/93 22:55:4
                                             Symbols-1

Segments and Groups:

Name	Length	Align	Combine Class
JOHN	0006	AT	3000

Symbols:

Name	Type	Value	Attr
CNTRL	NUMBER	0007	
PORTA	NUMBER	0001	
PORTC	NUMBER	0005	
@CPU	TEXT	0101h	
@FILENAME		TEXT	EXP318C
@VERSION		TEXT	510

```
    11 Source Lines
    11 Total Lines
     9 Symbols

47576 + 427939 Bytes symbol space free
```

```
0 Warning Errors
0 Severe Errors
```

3.10 8086 DMA

When configured in the minimum mode (MN/$\overline{\text{MX}}$ pin HIGH), the 8086 provides HOLD (DMA request) and HLDA (DMA acknowledge) signals to take over the system bus for DMA applications. The Intel DMA controller chips 8257 and 8237 can be used with the 8086. The 8257 or 8237 can request DMA transfer between the 8086 memory and I/O device by activating the 8086 HOLD pin. The 8086 will complete the current bus cycle (if there is one presently in progress) and then output HLDA, relinquishing the system bus to the DMA controller. The 8086 will not try to use the system bus until the HOLD pin is negated.

As mentioned before, the 8086 memory addresses are organized in two separate banks — one containing even-addressed bytes and the other containing odd-addressed bytes. An 8-bit DMA controller must alternately select these two banks to access logically adjacent bytes in memory.

QUESTIONS AND PROBLEMS

3.1   What is the basic difference between the 8086 and 8088 microprocessors? Name one reason why these two microprocessors are included in the i APX 86 family by Intel.

3.2   List the 8086 minimum and maximum mode signals. How are these modes selected?

3.3   What are the functions provided by 8288 bus controller in a maximum mode 8086 system?

3.4   Which bit of the 8086 FLAG register is used by the string instructions? How? Illustrate this by using the 8086 MOVSB instruction.

3.5   What is the relationship between the 8086 and 8284 input clocks? Does the 8086 have an on-chip clock circuitry? Comment.

3.6   What is the purpose of the TF bit in the FLAG register?

3.7   If [BL] = 36_{16} (ASCII code for 6) and [CL] = 33_{16} (ASCII code for 3), write an 8086 assembly program which will add the contents of BL and CL, and then provide the result in decimal. Store result in CL.

3.8   What happens to the contents of the AX register after execution of the following 8086 instruction sequence:

```
MOV AX, 0F180H
CBW
CWD
```

3.9   Determine the addressing modes for the following instructions:

```
(a)  MOV CH, 8
(b)  MOV AX, DS:START
(c)  MOV [SI], AL
```

(d) **MOV SI, BYTEPTR[BP+2][DI]**

3.10 Consider MOV BX, DS: BEGIN. How many memory accesses are required by the 8086 to execute the above instruction if BEGIN = 0401H.

3.11 Write an 8086 assembly program to implement the following Pascal segment:

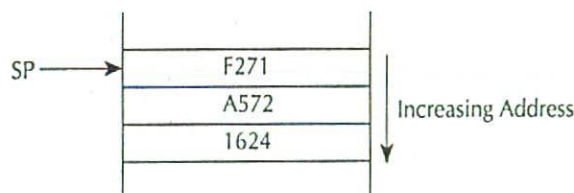```
SUM: =0; for i: = 0 to 15 do
SUM: = SUM + A(i)
```

Assume CS and DS are already initialized. A(i)s are 16-bit numbers.

**3.12**  Write and 8086 assembly program to add two 128-bit numbers stored in memory in consecutive locations.

**3.13**  What are the remainder, quotient, and registers containing them after execution of the following instruction sequence?

```
MOV DX, 0
MOV AX, -5
MOV BX, 2
IDIV BX
```

3.14 Write an 8086 assembly language program to divide A5721624H by F271H. Store the remainder and quotient onto the stack. Assume that the numbers are signed and stored in the stack as follows:

```
            ┌─────────────┐
SP ────────▶│    F271     │   │
            ├─────────────┤   │
            │    A572     │   │  Increasing Address
            ├─────────────┤   │
            │    1624     │   ▼
            ├─────────────┤
            │             │
            └─────────────┘
```

Assume that the stack segment and stack pointer are already initialized.

3.15 Write an 8086 assembly language program to compute X = Y + Z − 12FFH where X, Y, Z are 64-bit variables. The lower 32 bits of Y and Z are stored respectively at offset 5000H and 5008H followed by the upper 32 bits. Store the lower 32 bits of the 64-bit result at offset 6000H followed by the upper 32 bits.

3.16 Assume that registers AL, BX and DX CX contain a signed byte, a signed word, and a signed 32-bit number respectively. Write an 8086 assembly language program that will compute the signed 32-bit result:

$$AL + BX - DXCX \rightarrow DXCX$$

3.17 Write an 8086 assembly language program to compute X = 5 * Y + (Z/W) where offsets 5000H, 5002H, and 5004H respectively contain the 16-bit signed integers Y, Z, and W. Store the 32-bit result in memory starting at offset 5006H. Discard the remainder of Z/W.

3.18 Write an 8086 instruction sequence to clear the trap flag to zero.

3.19 Write an 8086 subroutine to compute

$$Y = \sum_{i=1}^{N} \frac{Xi^2}{N}$$

where Xis are signed 8-bit integers and N = 100. The numbers are stored in consecutive locations. Assume SI points to Xis and SP, DS, SS are already initialized.

3.20 Write an 8086 assembly program to move a block of data bytes of length 100_{10} from the source block starting at location 2000H in ES = 1000H to the destination block starting at location 3000H in the same extra segment.

3.21 Write an 8086 assembly program to logically shift a 128-bit number stored in location starting at 4000H in DS twice to the right. Store the result in memory location starting at 5000H in the data segment. Assume DS is initialized.

3.22 Connect one 2732, and one 8255 to an 8086 to obtain even 2732 locations and odd addresses for the 8255's Port A, Port B, Port C, and the Control register. Show only the connections for the pins shown in the figure below. Assume all unused address lines to be zeros. Use latches and gates as required.
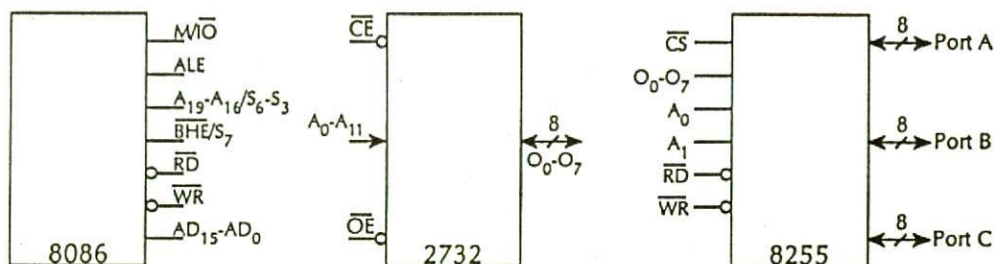
3.23 Determine the number of 2732 4K × 8 EPROMs and 6116 2K × 8 static RAMs to provide a 4K ×16 EPROM and a 2K ×16 word RAM. Draw a neat schematic connecting these chips to the 8086 and determine the map.

3.24 Assume the memory and I/O maps of Table 3.2. Interface the following A/D to the 8086/2716/6116/8255 of Figure 3.14:
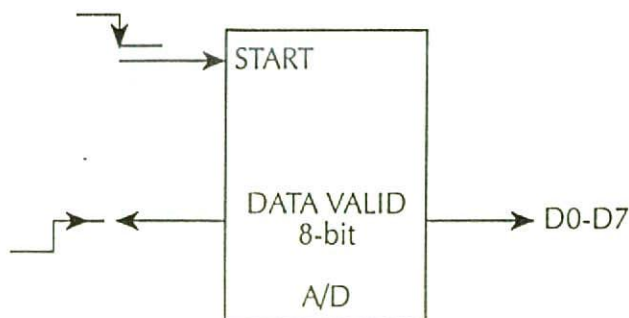
Write an 8086 assembly program to input the A/D converter and turn an LED ON connected to bit 5 of port A of the even 8255 if the number read from the A/D is odd; otherwise turn the LED OFF. Assume that the LED is turned ON by a HIGH and turned OFF by a LOW.

3.25   Repeat problem 3.24 using the 8086 INTR interrupt.

3.26   Write an 8086 assembly language program to add a 16-bit number stored in DX (bits 0 to 7 containing the high-order byte of the number and bits 8 to 15 containing the low-order byte) with another 16-bit number stored in BX (bits 0 to 7 containing the low-order 8 bits of the number and bits 8 through 15 containing the high-order 8 bits). Store the result in CX.

3.27   Assume an 8086/8255 based system. Write an 8086 assembly language program to input 16-bit data via Port B and Port C, and then divide this by the 8-bit input data at Port A. Assume all numbers to be signed.

3.28   An 8086/8255 based microcomputer is required to drive a common anode seven-segment display connected to Port C as follows:
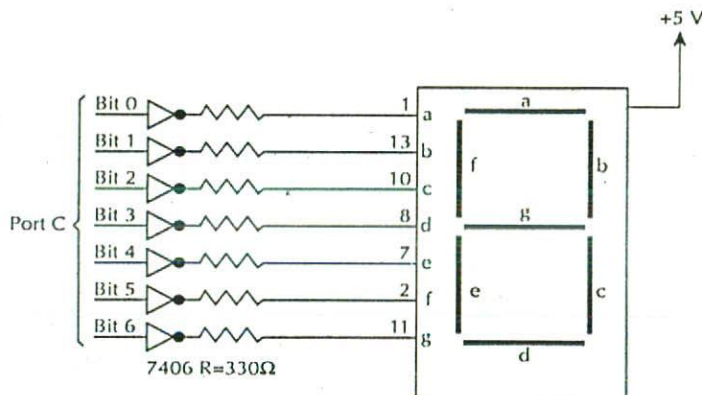
Write an 8086 assembly language program to display a single hexadecimal digit (0 to F) from offset 3000H at DS = 1000H. Use a look-up-table.

3.29   Assume an 8086/8255 based microcomputer. Suppose that four switches are connected to bits 0 through 3 of Port A, an LED is connected at bit 4 of Port B and another LED is connected at bit 2 of Port C. If the number of low switches is even, turn the Port B LED "ON" and Port C LED "OFF". If the number of low switches is ODD turn Port B LED "OFF" and Port C LED "ON". Write an 8086 assembly language program to accomplish the above.

3.30
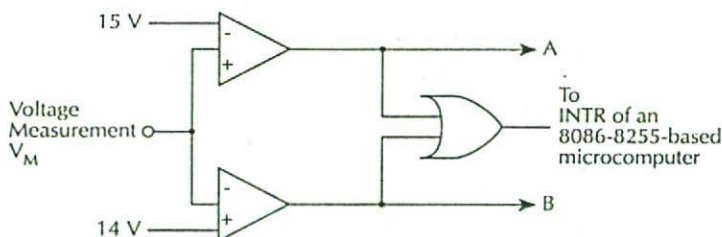
In the above figure, if $V_M > 15V$, turn an LED "ON" connected at bit 2 of Port C. On the other hand if $V_M < 14V$, turn the LED "OFF". Use registers and memory locations of your choice.

Draw a block diagram showing the microcomputer and the connections of the figure to its ports. Also, write 8086 assembly language programs to accomplish the above using:

a) Polled I/O by inputting one or more outputs in the figure.
b) INTR
c) NMI