



Department of Computer Engineering Techniques (Stage: 4)

Advance Computer Technologies

Dr.: Mayas Aljibawi

Segment and Offset Addressing Scheme Allows Relocation

A relocatable program is one that can be placed into any area of memory and executed without change. Relocatable data are data that can be placed in any area of memory and used without any change to the program. The segment and offset addressing scheme allows both programs and data to be relocated without changing anything in the programs or data. This is ideal for use in a general-purpose computer system where not all machines contain the same memory areas. The personal computer memory structure is different from machine to machine, requiring relocatable software and data.

Because memory is addressed within a segment by an offset address, the memory segment can be moved to any place in the memory system without changing any of the offset addresses. This is accomplished by moving the entire program, as a block, to a new area and then changing only the contents of the segment registers.

Protected Mode Memory Addressing

- Protected mode allows access to data and programs located within and above the first 1MB of memory.
 - Microsoft Windows operates in protected mode.
- The segment registers are used differently in protected mode.
 - They no longer store the segment address.
 - Instead, they contain a selector that selects a descriptor from a descriptor table.

Protected Mode Memory Addressing

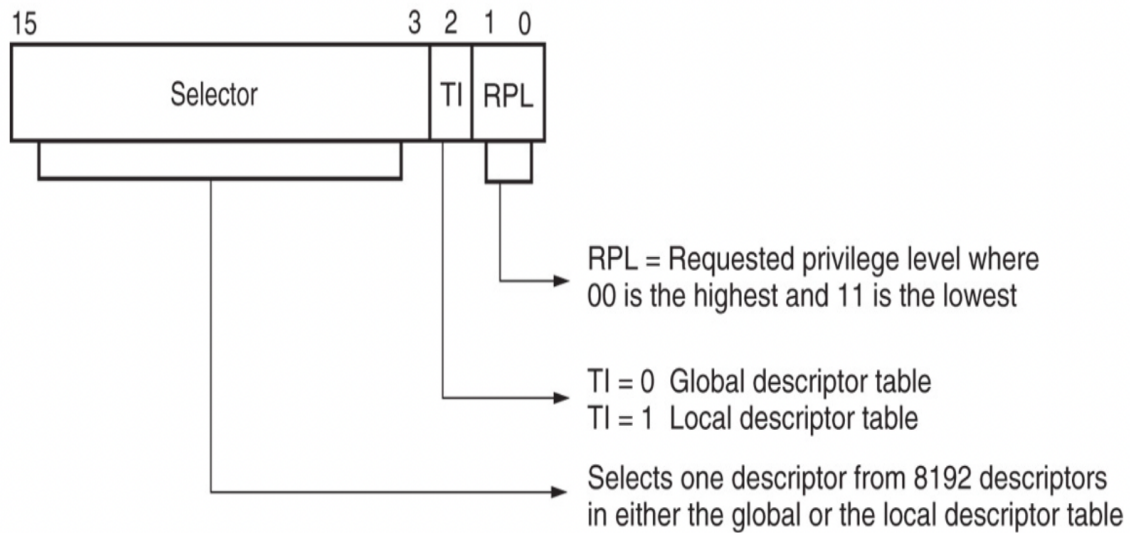
- A descriptor table can contain up to 8192 descriptors.
- There are two types of descriptor tables:
 - **Global descriptors**: contain segment definitions that apply to all programs.
 - **Local descriptors**: unique to an application.
- A descriptor contains information about:
 - Memory segment's location
 - Length of the segment
 - Access rights

Selectors and Descriptors

The selector, located in the segment register, selects one of 8,192 descriptors from one of two tables of descriptors. The descriptor describes the location, length, and access rights of the segment of memory. Indirectly, the segment register still selects a memory segment, but not directly as in the real mode. For example, in the real mode, if CS = 0008H, the code segment begins at location 00080H. In the protected mode, this segment number can address any memory location in the entire system for the code segment, as explained shortly.

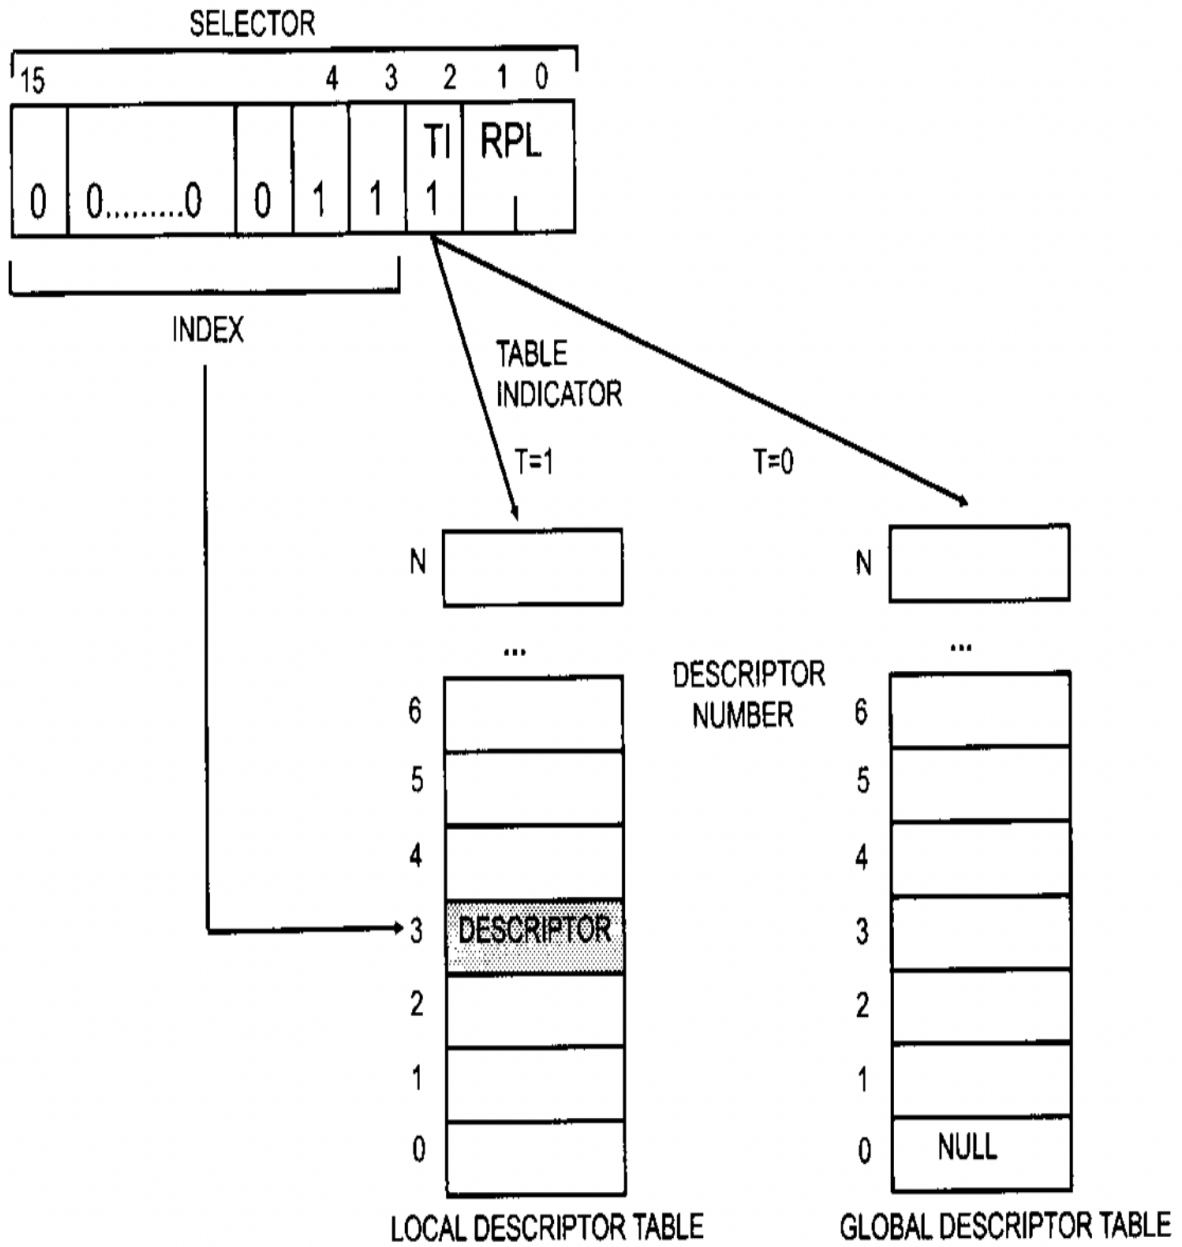
There are two descriptor tables used with the segment registers: one contains global descriptors and the other contains local descriptors. The **global descriptors** contain segment definitions that apply to all programs, while the **local descriptors** are usually unique to an application. Each descriptor table contains 8,192 descriptors, so a total of 16,384 descriptors are available to an application at any time. Because the descriptor describes a memory segment, this allows up to 16,384 memory segments to be described for each application.

Protected Mode Memory Addressing



The contents of a segment register during protected mode operation of the 80286 through Core2 microprocessors.

Local and global descriptor tables






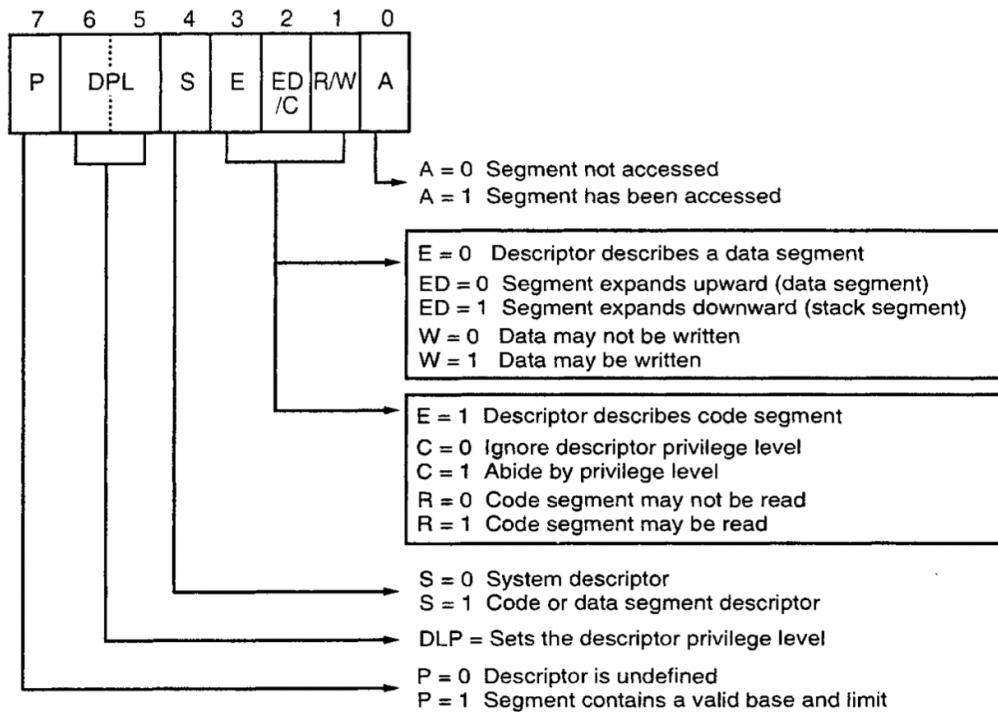
FIGURE 2-6 The descriptor formats for the 80286 and 80386/80486/Pentium/Pentium Pro microprocessors

The **segment limit** contains the last offset address found in a segment. For example, if a segment begins at memory location F00000H and ends at location F000FFH, the base address is F00000H and the limit is FFH. For the 80286 microprocessor, the base address is F00000H and the limit is 00FFH. For the 80386 and above, the base address is 00F00000H and the limit is 000FFH. Notice that the 80286 has a 16-bit limit, and the 80386 through the Pentium Pro have a 20-bit limit. The 80286 accesses memory segments that are between 1 and 64K bytes in length. The 80386 and above access memory segments that are between 1 and 1M byte or 4K and 4G bytes in length.

There is another feature found in the 80386 through the Pentium Pro descriptor that is not found in the 80286 descriptor: the G bit or **granularity bit**. If $G = 0$, the limit specifies a segment limit of from 1 to 1M byte in length. If $G = 1$, the value of the limit is multiplied by 4K bytes (appended with 000H). If $G = 1$, the limit is any multiple of 4K bytes. This allows a segment length of 4K to 4G bytes in steps of 4K bytes. The reason that the segment length is 64K bytes in the 80286 is that the offset address is always 16-bits because of its 16-bit internal architecture. The 80386 and above use a 32-bit architecture, which allows an offset address, in the protected mode operation, of the 32-bits. This 32-bit offset address allows segment lengths of 4G bytes, and the 16-bit offset address allows segment lengths of 64K bytes. Operating systems operate in either a 16- or 32-bit environment.

The AV bit, in the 80386 and above descriptor, is used by some operating systems to indicate that the segment is available ($AV = 1$) or not available ($AV = 0$). The D bit indicates how the 80386 through the Pentium Pro instructions access register and memory data in the protected or real mode. If $D = 0$, the instructions are 16-bit instructions compatible with the 8086–80286 microprocessors. This means that the instructions use 16-bit offset addresses and 16-bit registers by default. This mode is often called the 16-bit instruction mode. If $D = 1$, the instructions are 32-bit instructions. By default, the 32-bit instruction mode assumes that all offset addresses as well as all registers are 32-bits.

The **access rights byte** (see Figure 2-7) controls access to the protected mode memory segment. This byte describes how the segment functions in the system. The access rights byte allows complete control over the segment. If the segment is a data segment, the direction of growth is specified. If the segment grows beyond its limit, the microprocessor's program is interrupted, indicating a general protection fault. You can even specify if a data segment can be written or is write-protected. The code segment is also controlled in a similar fashion and can have reading inhibited to protect software.



Note: Some of the letters used to describe the bits in the access rights bytes vary in Intel documentation.

FIGURE 2-7 The access rights byte for the 80286, 80386, 80486, Pentium, and Pentium Pro descriptor

EXAMPLE 8.9

The access rights byte of a segment descriptor contains FE_{16} . What type of segment descriptor does it describe, and what are its characteristics?

Solution

Expressing the access rights byte in binary form, we get

$$FE_{16} = 11111110_2$$

Since bit 4 is 1, the access rights byte is for a code/data segment descriptor. This segment has the characteristics that follow:

$P = 1$ = segment is mapped into physical memory

$DPL = 11$ = privilege level 3

$E = 1$ = executable code segment

$C = 1$ = conforming code segment

$R = 1$ = readable code segment

$A = 0$ = segment has not been accessed

Example 21-13

From Figure 21-8 we have the following access byte for code and data.

P DPL 1 1 A (access byte for code segment)

P DPL 1 0 A (access byte for data segment)

Discuss the following access bytes.

(a) 10011011 (b) 10010111 (c) 11110001

Solution:

- (a) This is an access byte for code segment, present, accessed, and privilege level of 00 (highest).
- (b) This is an access byte for data segment, present, accessed, privilege level of 00 (highest), and both read and write accessible.
- (c) This is an access byte for data segment, present, accessed, privilege level of 11 (lowest), and write protected.

Q1- State the characteristic of each of the following access bytes. State for each if it is for code or data.

1. 10010001

2. 11111011