

CS 390

Chapter 1 Homework Solutions

1.1 What are the three main ...

The three main purposes of an operating system are

1. To simplify the use of the hardware by providing a convenient environment for the computer user to execute and interact with programs;
2. To enforce resource allocation policies;
3. To protect the hardware from accidental or malicious misuse.

1.5 How does the distinction between ...

Certain machine instructions can only be executed when the CPU is in kernel mode. These instructions typically (depending on the CPU) are those that allow for direct access to hardware. Thus, the CPU has limited capabilities when executing in user mode, thereby enforcing the protection of critical resources.

1.6 Which of the following instructions ...

- a. Yes (Here, the term “timer” refers to the hardware component that generates periodic interrupts, and not the system clock.)
- b. No
- c. Yes (This question means “clear the entire contents of the computer’s memory”.)
- d. No
- e. Yes
- f. Yes
- g. No (think carefully about why no is the only correct answer)
- h. Yes

1.7 Some early computers protected the ...

The data required by the operating system (passwords, filenames, accounting information) could not be stored in the same memory as

the operating system, because that memory could not be altered. Thus, it had to be stored in regular memory, where it was readable and writable by regular user processes. This would allow user processes to see information that they should not be able to see (like passwords) and modify data (like the contents of other processes' files) that they should not be able to modify.

The operating system would not be modifiable at run-time. If a new keyboard or mouse were added, the OS would have to be reconfigured to include the new code, and then recompiled and reinstalled.

1.15 Describe the differences between symmetric ...

In symmetric multiprocessing, all the processors in the system are peers, and each can perform all the tasks of the others.

In asymmetric multiprocessing, processors have different responsibilities and may have different capabilities. For example, in some systems, one processor controls the system and supervises all the other processors. (The Sony CELL processor of the PS3 game system uses this architecture.) In other architectures, some processors might be dedicated to a specific type of task (say, performing I/O or encryption, or video or sound processing). General purpose computer systems with high-performance video cards can also be classified as a type of asymmetric multiprocessing system.

Advantages of multiprocessor systems:

Increased throughput: A single multiprocessor machine can complete more work per unit of time.

Economy of scale: It is usually cheaper to construct a single machine with N CPU's than it is to build N machines each with a single CPU.

Increased reliability: Multiprocessor systems may be designed to that if a single processor fails, the system can continue to operate with the remaining good processors.

Disadvantages: Unless they are carefully designed and programmed, multiprocessor systems are an inefficient use of hardware resources. In addition, systems which run at most one job at a time will not benefit from multiprocessing.

1.19 What is the purpose of

An interrupt is an electrical signal sent from one component of the computer to another, signaling that an event has occurred which requires some action by the component receiving the interrupt. A trap is a specific type of interrupt which is generated by software and signals either

1. some unusual condition, such as an array index out of bounds, or
2. a request for an operating system service (a system call.)

User programs can generate traps, although not all high-level programming languages make this facility directly available to the programmer. (Ada, C and Java do allow the programmer to generate traps. Traps are called “exceptions” in Ada and Java, and “signals” in C. In addition, C allows programmers to bypass the API and directly make system calls.)

Traps are used to notify the CPU that some unusual software event has taken place, or that some special service has been requested by the program. The CPU can then either terminate the program, or interrupt the normal execution sequence of the program and transfer control to a special segment of code.

1.20 Direct memory access is used ...

- a. The CPU sets up a series of buffers in memory where the I/O device will copy the data to/from, along with a pointer or pointers to the buffers, and counters indicating the number of bytes to be transferred. It then notifies the I/O device to begin the transfer.
- b. The I/O device interrupts the CPU.
- c. Yes, it can interfere with the execution of other user programs. If a user program wants to modify some memory that is being used for DMA I/O, the program must be blocked until the DMA completes. In addition, since the I/O device is accessing memory, it can interfere directly with the CPU's memory access, since both typically use the same memory bus (in many microcomputers, at least). This might cause the CPU to complete fewer instructions per unit of time.

1.25 Describe a mechanism for enforcing ...

The term “memory protection” means preventing one process from reading or writing the memory of another process without explicit permission. Since the kernel is responsible for allocating memory to processes, the kernel knows, for each memory address, which process the address belongs to. Before a process reads or writes a memory location, the kernel could inspect the address and stop the process when it attempts to access memory that does not belong to it.

1.27 Describe some of the challenges ...

- Mobile devices may have limited memory. If so, the size of the kernel might be limited, and this might limit the features it can offer. In addition, the kernel might need to be more concerned with the allocation of memory to processes than it would in a general purpose system.
- Power consumption is an important issue with battery powered devices, so the kernel would need to control power usage by processes and I/O devices.
- A mobile device will probably have a slower or less-sophisticated CPU than a desktop or server machine. That means a mobile OS might need to be more concerned with the efficient allocation of CPU-cycles among processes than the OS of a larger system.