CHAPTER 3

In this chapter we introduce the concepts of a process and concurrent execution; These concepts are at the very heart of modern operating systems. A process is a program in execution and is the unit of work in a modern time-sharing system. Such a system consists of a collection of processes: Operating-system processes executing system code and user processes executing user code. All these processes can potentially execute concurrently, with the CPU (or CPUs) multiplexed among them. By switching the CPU between processes, the operating system can make the computer more productive. We also introduce the notion of a thread (lightweight process) and interprocess communication (IPC). Threads are discussed in more detail in Chapter 4.

Exercises

3.1 Describe the differences among short-term, medium-term, and long-term scheduling.
Answer:
   a. Short-term (CPU scheduler)—selects from jobs in memory those jobs that are ready to execute and allocates the CPU to them.
   b. Medium-term—used especially with time-sharing systems as an intermediate scheduling level. A swapping scheme is implemented to remove partially run programs from memory and reinstate them later to continue where they left off.
   c. Long-term (job scheduler)—determines which jobs are brought into memory for processing.
The primary difference is in the frequency of their execution. The short-term must select a new process quite often. Long-term is used much less often since it handles placing jobs in the system and may wait a while for a job to finish before it admits another one.

3.2 Describe the actions taken by a kernel to context-switch between processes.
Answer:
In general, the operating system must save the state of the currently running process and restore the state of the process scheduled to be run next. Saving the state of a process typically includes the values of all the CPU registers in addition to memory allocation. Context switches must also perform many architecture-specific operations, including flushing data and instruction caches.

3.3 Construct a process tree similar to Figure 3.8. To obtain process information for the UNIX or Linux system, use the command `ps  -ael`. Use the command `man ps` to get more information about the `ps` command. The task manager on Windows systems does not provide the parent process id, yet the **process monitor** tool available from `technet.microsoft.com` provides a process tree tool.
Answer:
Answer: Results will vary widely.

3.4 Explain the role of the `init` process on UNIX and Linux systems in regards to process termination.
Answer:
When a process is terminated, it briefly moves to the zombie state and remains in that state until the parent invokes a call to `wait()`. When this occurs, the process id as well as entry in the process table are both released. However, if a parent does not invoke `wait()`, the child process remains a zombie as long as the parent remains alive. Once the parent process terminates, the `init` process becomes the new parent of the zombie. Periodically, the `init` process calls `wait()` which ultimately releases the `pid` and entry in the process table of the zombie process.

3.5 Including the initial parent process, how many processes are created by the program shown in Figure Figure 3.32?
Answer:
16 processes are created. The program online includes `printf()` statements to better understand how many processes have been created.

3.6 Explain the circumstances when the line of code marked printf("`LINE J`") in Figure 3.33 is reached.
Answer:
The call to `exec()` replaces the address space of the process with the program specified as the parameter to `exec()`. If the call to `exec()` succeeds, the new program is now running and control from the call to `exec()` never returns. In this scenario, the line `printf("Line J");` would never be performed. However, if an error occurs in the call to `exec()`, the function returns control and therefor the line `printf("Line J");` would be performed.

3.7 Using the program in Figure Figure 3.34, identify the values of `pid` at lines A, B, C, and D. (Assume that the actual pids of the parent and child are 2600 and 2603, respectively.)
Answer:
Answer: A = 0, B = 2603, C = 2603, D = 2600

3.8 Give an example of a situation in which ordinary pipes are more suitable than named pipes and an example of a situation in which named pipes are more suitable than ordinary pipes.
Answer:
Simple communication works well with ordinary pipes. For example, assume we have a process that counts characters in a file. An ordinary pipe can be used where the producer writes the file to the pipe and the consumer reads the files and counts the number of characters in the file. Next, for an example where named pipes are more suitable, consider the situation where several processes may write messages to a log. When processes wish to write a message to the log, they write it to the named pipe. A server reads the messages from the named pipe and writes them to the log file.

3.9 Consider the RPC mechanism. Describe the undesirable consequences that could arise from not enforcing either the "at most once" or "exactly once" semantic. Describe possible uses for a mechanism that has neither of these guarantees.
Answer:
If an RPC mechanism cannot support either the "at most once" or "at least once" semantics, then the RPC server cannot guarantee that a remote procedure will not be invoked multiple occurrences. Consider if a remote procedure were withdrawing money from a bank account on a system that did not support these semantics. It is possible that a single invocation of the remote procedure might lead to multiple withdrawals on the server.

      For a system to support either of these semantics generally requires the server maintain some form of client state such as the timestamp described in the text.

      If a system were unable to support either of these semantics, then such a system could only safely provide remote procedures that do not alter data or provide time-sensitive results. Using our bank account as an example, we certainly require "at most once" or "at least once" semantics for performing a withdrawal (or deposit!). However, an inquiry into an account balance or other account information such as name, address, etc. does not require these semantics.

3.10 Using the program shown in Figure 3.35, explain what the output will be at lines X and Y.
Answer:

Because the child is a copy of the parent, any changes the child makes will occur in its copy of the data and won't be reflected in the parent. As a result, the values output by the child at line X are 0, -1, -4, -9, -16. The values output by the parent at line Y are 0, 1, 2, 3, 4

3.11 What are the benefits and the disadvantages of each of the following? Consider both the system level and the programmer level.
   a. Synchronous and asynchronous communication
   b. Automatic and explicit buffering
   c. Send by copy and send by reference
   d. Fixed-sized and variable-sized messages

Answer:
   a. **Synchronous and asynchronous communication**—A benefit of synchronous communication is that it allows a rendezvous between the sender and receiver. A disadvantage of a blocking send is that a rendezvous may not be required and the message could be delivered asynchronously. As a result, message-passing systems often provide both forms of synchronization.
   b. **Automatic and explicit buffering**—Automatic buffering provides a queue with indefinite length, thus ensuring the sender will never have to block while waiting to copy a message. There are no specifications on how automatic buffering will be provided; one scheme may reserve sufficiently large memory where much of the memory is wasted. Explicit buffering specifies how large the buffer is. In this situation, the sender may be blocked while waiting for available space in the queue. However, it is less likely that memory will be wasted with explicit buffering.
   c. **Send by copy and send by reference**—Send by copy does not allow the receiver to alter the state of the parameter; send by reference does allow it. A benefit of send by reference is that it allows the programmer to write a distributed version of a centralized application. Java's RMI provides both; however, passing a parameter by reference requires declaring the parameter as a remote object as well.
   d. **Fixed-sized and variable-sized messages**—The implications of this are mostly related to buffering issues; with fixed-size messages, a buffer with a specific size can hold a known number of messages. The number of variable-sized messages that can be held by such a buffer is unknown. Consider how Windows 2000 handles this situation: with fixed-sized messages (anything < 256 bytes), the messages are copied from the address space of the sender to the address space of the receiving process. Larger messages (i.e. variable-sized messages) use shared memory to pass the message.