



# Topic 3

## (Textbook - Chapter 3)

### Processes

---







# Chapter 3: Processes

---

- Process Concept
- Process Scheduling
- Operations on Processes
- Interprocess Communication







# Objectives

---

- To introduce the notion of a process -- a program in execution, which forms the basis of all computation
- To describe the various features of processes, including scheduling, creation and termination, and communication






# Process Concept


---

- An operating system executes a variety of programs:
  - Batch system – jobs
  - Time-shared systems – user programs or tasks
- Textbook uses the terms *job* and *process* almost interchangeably
- Process – a program in execution; process execution must progress in sequential fashion
- A process includes:
  - program counter
  - stack
  - data section






# Process in Memory



A process is more than the program code, which is sometimes known as the **text section**. It also includes the current activity, as represented by the value of the **program counter** and the contents of the processor's registers. A process generally also includes the process **stack**, which contains temporary data (such as function parameters, return addresses, and local variables), and a **data section**, which contains global variables. A process may also include a **heap**, which is memory that is dynamically allocated during process run time. The structure of a process in memory is shown in Figure 3.1.






# Process State


---


- As a process executes, it changes *state*
  - **new**: The process is being created
  - **running**: Instructions are being executed
  - **waiting**: The process is waiting for some event to occur
  - **ready**: The process is waiting to be assigned to a processor
  - **terminated**: The process has finished execution





# Diagram of Process State






# Process Control Block (PCB)

---

Information associated with each process


- Process state
- Program counter
- CPU registers
- CPU scheduling information
- Memory-management information
- Accounting information
- I/O status information





# Process Control Block (PCB)

---





**Process state.** The state may be new, ready, running, waiting, halted, and so on.

• **Program counter.** The counter indicates the address of the next instruction to be executed for this process.

• **CPU registers.** The registers vary in number and type, depending on the computer architecture. They include accumulators, index registers, stack pointers, and general-purpose registers, plus any condition-code information. Along with the program counter, this state information must be saved when an interrupt occurs, to allow the process to be continued correctly afterward (Figure 3.4).

• **CPU-scheduling information.** This information includes a process priority, pointers to scheduling queues, and any other scheduling parameters.


(Chapter 6 describes process scheduling.)


• **Memory-management information.** This information may include such items as the value of the base and limit registers and the page tables, or the segment tables, depending on the memory system used by the operating system (Chapter 8).

**Accounting information.** This information includes the amount of CPU and real time used, time limits, account numbers, job or process numbers, and so on.


• **I/O status information.** This information includes the list of I/O devices allocated to the process, a list of open files, and so on.


In brief, the PCB simply serves as the repository for any information that may vary from process to process.





# CPU Switch From Process to Process






# Process Scheduling Queues


---


- **Job queue** – set of all processes in the system
- **Ready queue** – set of all processes residing in main memory, ready and waiting to execute
- **Device queues** – set of processes waiting for an I/O device
- Processes migrate among the various queues






# Ready Queue And Various I/O Device Queues





# Representation of Process Scheduling







# Schedulers

---

- **Long-term scheduler** (or job scheduler) – selects which processes should be brought into the ready queue
- **Short-term scheduler** (or CPU scheduler) – selects which process should be executed next and allocates CPU







■ These processes are spooled to a mass-storage device (typically a disk), where they are kept for later execution. The **long-term scheduler**, or **job**


- **scheduler**, selects processes from this pool and loads them into memory for execution. The **short-term scheduler**, or **CPU scheduler**, selects from among
- the processes that are ready to execute and allocates the CPU to one of them. The primary distinction between these two schedulers lies in frequency
- of execution. The short-term scheduler must select a new process for the CPU frequently. A process may execute for only a few milliseconds before waiting
- for an I/O request. Often, the short-term scheduler executes at least once every 100 milliseconds. Because of the short time between executions, the short-term
- scheduler must be fast. If it takes 10 milliseconds to decide to execute a process for 100 milliseconds, then  $10/(100 + 10) = 9$  percent of the CPU is being used
- (wasted) simply for scheduling the work.





# Addition of Medium Term Scheduling





Some operating systems, such as time-sharing systems, may introduce an additional, intermediate level of scheduling. This **medium-term scheduler** is


diagrammed in Figure 3.7. The key idea behind a medium-term scheduler is that sometimes it can be advantageous to remove a process from memory

(and from active contention for the CPU) and thus reduce the degree of multiprogramming. Later, the process can be reintroduced into memory, and its

execution can be continued where it left off. This scheme is called **swapping**.

The process is swapped out, and is later swapped in, by the medium-term scheduler. Swapping may be necessary to improve the process mix or because a change in memory requirements has overcommitted available memory, requiring memory to be freed up. Swapping is discussed in Chapter 8.






# Schedulers (Cont)

---

- Short-term scheduler is invoked very frequently (milliseconds)  $\Rightarrow$  (must be fast)
- Long-term scheduler is invoked very infrequently (seconds, minutes)  $\Rightarrow$  (may be slow)
- The long-term scheduler controls the *degree of multiprogramming* (*the number of processes in memory*)
- Processes can be described as either:
  - **I/O-bound process** – spends more time doing I/O than computations, many short CPU bursts
  - **CPU-bound process** – spends more time doing computations; few very long CPU bursts







# Context Switch

---

- When CPU switches to another process, the system must save the state of the old process and load the saved state for the new process via a **context switch**
- **Context** of a process represented in the PCB
- Context-switch time is overhead; the system does no useful work while switching
- Time dependent on hardware support






# Process Creation

---

- **Parent** process create **children** processes, which, in turn create other processes, forming a tree of processes
- Generally, process identified and managed via a **process identifier (pid)**
- Resource sharing
  - Parent and children share all resources
  - Children share subset of parent's resources
  - Parent and child share no resources
- Execution
  - Parent and children execute concurrently
  - Parent waits until children terminate





# Process Creation (Cont.)





Figure 3.8 A tree of processes on a typical Linux system.





# Process Creation (Cont)


---

- Address space either
  - Child duplicate of parent
  - Child has a program loaded into it
- UNIX examples
  - **fork** system call creates new process
  - **exec** system call used after a **fork** to replace the process' memory space with a new program



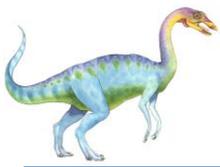


# Process Creation



**Figure 3.10** Process creation using the `fork()` system call.






# C Program Forking Separate Process


---


```
int main()
{
    pid_t pid;
    /* fork another process */
    pid = fork();
    if (pid < 0) { /* error occurred */
        fprintf(stderr, "Fork Failed");
        exit(-1);
    }
    else if (pid == 0) { /* child process */
        execlp("/bin/ls", "ls", NULL);
    }
    else { /* parent process */
        /* parent will wait for the child to complete */
        wait (NULL);
        printf ("Child Complete");
        exit(0);
    }
}
```





# A tree of processes on a typical Solaris






# Process Termination

- Process executes last statement and asks the operating system to delete it (**exit**)
  - Output data from child to parent (via **wait**)
  - Process' resources are deallocated by operating system
- Parent may terminate execution of children processes (**abort**)
  - Child has exceeded allocated resources
  - Task assigned to child is no longer required
  - If parent is exiting
    - ▶ Some operating system do not allow child to continue if its parent terminates
      - All children terminated - **cascading termination**
- A process that has terminated, but whose parent has not yet called `wait()`, is known as a **zombie** process.
- If a parent did not invoke `wait()` and instead terminated, thereby leaving its child processes as **orphans**, What would happen?






# Interprocess Communication


---


- Processes within a system may be **independent** or **cooperating**
- Cooperating process can affect or be affected by other processes, including sharing data
- Reasons for cooperating processes:
  - Information sharing
  - Computation speedup
  - Modularity
  - Convenience
- Cooperating processes need **interprocess communication (IPC)** mechanism to exchange data and information
- Two models of IPC
  - Shared memory
  - Message passing





# Communications Models






# Producer-Consumer Problem

---

- OS normally prevent one process to access another process's memory. However,
- Shared memory require that two or more processes agree to remove this restriction.
- A common paradigm for cooperating processes, *producer* process produces information that is consumed by a *consumer* process
  - *unbounded-buffer* places no practical limit on the size of the buffer
  - *bounded-buffer* assumes that there is a fixed buffer size





# Bounded-Buffer – Shared-Memory Solution


- Shared data

```
#define BUFFER_SIZE 10
typedef struct {
    . . .
} item;

item buffer[BUFFER_SIZE];
int in = 0;
int out = 0;
```

- Solution is correct, but can only use  $BUFFER\_SIZE-1$  elements






# Bounded-Buffer – Producer

---

```
item next_produced;
```

```
while (true) {  
    /* Produce an item next_produced */  
    while (((in + 1) % BUFFER_SIZE) == out)  
        ; /* do nothing -- no free buffers */  
  
    buffer[in] = next_produced;  
    in = (in + 1) % BUFFER_SIZE;  
}
```





# Bounded Buffer – Consumer

---

```
Item next_consumed;
```

```
while (true) {  
    while (in == out)  
        ; // do nothing -- nothing to consume  
  
    // remove an item from the buffer  
    next_consumed = buffer[out];  
    out = (out + 1) % BUFFER_SIZE;  
    return item;  
}
```



# End of Chapter 3

---

