

Unit III

CONCURRENCY AND SCHEDULING

Principles of Concurrency

Mutual Exclusion

Semaphores

Monitors

Readers/Writers problem.

Deadlocks –

Prevention


Avoidance

Detection

Scheduling

Types of Scheduling

Scheduling algorithms.



Roadmap

- Principals of Concurrency
- • Mutual Exclusion: Hardware Support
- Semaphores
- Monitors
- Message Passing
- Readers/Writers Problem



Multiple Processes



- Central to the design of modern Operating Systems is managing multiple processes
 - Multiprogramming(multiple processes within a uniprocessor system)
 - Multiprocessing(multiple processes within a multiprocessor)
 - Distributed Processing(multiple processes executing on multiple ,distributed processing)
- Big Issue is Concurrency
 - Managing the interaction of all of these processes

Concurrency

Concurrency arises in:

- Multiple applications
 - (processing time)Sharing time
- Structured applications
 - Extension of modular design(structured programming)
- Operating system structure
 - OS themselves implemented as a set of processes or threads



Key Terms



Table 5.1 Some Key Terms Related to Concurrency

atomic operation	A sequence of one or more statements that appears to be indivisible; that is, no other process can see an intermediate state or interrupt the operation.
critical section	A section of code within a process that requires access to shared resources and that must not be executed while another process is in a corresponding section of code.
deadlock	A situation in which two or more processes are unable to proceed because each is waiting for one of the others to do something.
livelock	A situation in which two or more processes continuously change their states in response to changes in the other process(es) without doing any useful work.
mutual exclusion	The requirement that when one process is in a critical section that accesses shared resources, no other process may be in a critical section that accesses any of those shared resources.
race condition	A situation in which multiple threads or processes read and write a shared data item and the final result depends on the relative timing of their execution.
starvation	A situation in which a runnable process is overlooked indefinitely by the scheduler; although it is able to proceed, it is never chosen.

Interleaving and Overlapping Processes



- Earlier (Ch2) we saw that processes may be interleaved on uniprocessors





Figure 2.12 Multiprogramming and Multiprocessing

Interleaving and Overlapping Processes



- And not only interleaved but overlapped on multi-processors



(b) Interleaving and overlapping (multiprocessing; two processors)

■ Blocked ■ Running

Figure 2.12 Multiprogramming and Multiprocessing

Difficulties of Concurrency



- Sharing of global resources
- Optimally managing the allocation of resources
- Difficult to locate programming errors as results are not deterministic and reproducible.

A Simple Example

```
void echo()
{
    chin = getchar();
    chout = chin;
    putchar(chout);
}
```



A Simple Example: On a Multiprocessor



Process P1

```
.  
chin = getchar();  
.   
chout = chin;  
putchar(chout);  
.   
.
```

Process P2

```
.   
.   
chin = getchar();  
chout = chin;  
.   
putchar(chout);  
.   
.
```

Enforce Single Access

- If we enforce a rule that only one process may enter the function at a time then:
- P1 & P2 run on separate processors
- P1 enters echo first,
 - P2 tries to enter but is blocked – P2 suspends
- P1 completes execution
 - P2 resumes and executes echo



Race Condition



- A race condition occurs when
 - Multiple processes or threads read and write data items
 - They do so in a way where the final result depends on the order of execution of the processes.
- The output depends on who finishes the race last.

Operating System Concerns



- What design and management issues are raised by the existence of concurrency?
- The OS must
 - Keep track of various processes(PCB)
 - Allocate and de-allocate resources(PROCESSOR TIME, MEMORY, FILES, I/O DEVICES)
 - Protect the data and resources against interference by other processes.
 - Ensure that the processes and outputs are independent of the processing speed

Process Interaction



Table 5.2 Process Interaction

Degree of Awareness	Relationship	Influence That One Process Has on the Other	Potential Control Problems
Processes unaware of each other	Competition	<ul style="list-style-type: none">• Results of one process independent of the action of others• Timing of process may be affected	<ul style="list-style-type: none">• Mutual exclusion• Deadlock (renewable resource)• Starvation
Processes indirectly aware of each other (e.g., shared object)	Cooperation by sharing	<ul style="list-style-type: none">• Results of one process may depend on information obtained from others• Timing of process may be affected	<ul style="list-style-type: none">• Mutual exclusion• Deadlock (renewable resource)• Starvation• Data coherence
Processes directly aware of each other (have communication primitives available to them)	Cooperation by communication	<ul style="list-style-type: none">• Results of one process may depend on information obtained from others• Timing of process may be affected	<ul style="list-style-type: none">• Deadlock (consumable resource)• Starvation

Competition among Processes for Resources



Three main control problems:

- Need for Mutual Exclusion
 - Critical sections
- Deadlock
- Starvation

Requirements for Mutual Exclusion



- Only one process at a time is allowed in the critical section for a resource
- A process that halts in its noncritical section must do so without interfering with other processes
- No deadlock or starvation

Requirements for Mutual Exclusion



- A process must not be delayed access to a critical section when there is no other process using it
- No assumptions are made about relative process speeds or number of processes
- A process remains inside its critical section for a finite time only

Roadmap

- Principals of Concurrency
- Mutual Exclusion: Hardware Support
- • Semaphores
- Monitors
- Message Passing
- Readers/Writers Problem



Disabling Interrupts



- Uniprocessors only allow interleaving
- Interrupt Disabling
 - A process runs until it invokes an operating system service or until it is interrupted
 - Disabling interrupts guarantees mutual exclusion
 - Will not work in multiprocessor architecture

Pseudo-Code



```
while (true) {  
    /* disable interrupts */;  
    /* critical section */;  
    /* enable interrupts */;  
    /* remainder */;  
}
```

Special Machine Instructions



- Compare&Swap Instruction
 - also called a “compare and exchange instruction”
- Exchange Instruction
- COMPARE is made between a memory value and a test value, if the values are same ,a swap occurs.

Compare&Swap Instruction



```
int compare_and_swap (int *word,  
    int testval, int newval)  
{  
    int oldval;  
    oldval = *word;  
    if (oldval == testval) *word = newval;  
    return oldval;  
}
```

Mutual Exclusion (fig 5.2)



```
/* program mutualexclusion */
const int n = /* number of processes */;
int bolt;
void P(int i)
{
    while (true) {
        while (compare_and_swap(bolt, 0, 1) == 1)
            /* do nothing */;
        /* critical section */;
        bolt = 0;
        /* remainder */;
    }
}
void main()
{
    bolt = 0;
    parbegin (P(1), P(2), ... ,P(n));
}
}
```

(a) Compare and swap instruction

Exchange instruction



```
void exchange (int register, int
memory)
{
    int temp;
    temp = memory;
    memory = register;
    register = temp;
}
```


Exchange Instruction (fig 5.2)



```
    /* program mutualexclusion */
int const n = /* number of processes**/;
int bolt;
void P(int i)
{
    int keyi = 1;
    while (true) {
        do exchange (keyi, bolt)
        while (keyi != 0);
        /* critical section */;
        bolt = 0;
        /* remainder */;
    }
}
void main()
{
    bolt = 0;
    parbegin (P(1), P(2), ..., P(n));
}
```

(b) Exchange instruction

Hardware Mutual Exclusion: Advantages



- Applicable to any number of processes on either a single processor or multiple processors sharing main memory
- It is simple and therefore easy to verify
- It can be used to support multiple critical sections

Hardware Mutual Exclusion: Disadvantages



- Busy-waiting consumes processor time
- Starvation is possible when a process leaves a critical section and more than one process is waiting.
 - Some process could indefinitely be denied access.
- Deadlock is possible

Roadmap

- Principals of Concurrency
- Mutual Exclusion: Hardware Support
- Semaphores
- • Monitors
- Message Passing
- Readers/Writers Problem



Semaphore



- Semaphore:
 - An integer value used for signalling among processes.
- Only three operations may be performed on a semaphore, all of which are atomic:
 - initialize,
 - Decrement (`semWait`)
 - increment. (`semSignal`)

Semaphore Primitives

```
struct semaphore {
    int count;
    queueType queue;
};
void semWait(semaphore s)
{
    s.count--;
    if (s.count < 0) {
        /* place this process in s.queue */;
        /* block this process */;
    }
}
void semSignal(semaphore s)
{
    s.count++;
    if (s.count <= 0) {
        /* remove a process P from s.queue */;
        /* place process P on ready list */;
    }
}
```

Figure 5.3 A Definition of Semaphore Primitives

Binary Semaphore Primitives



```
struct binary_semaphore {
    enum {zero, one} value;
    queueType queue;
};
void semWaitB(binary_semaphore s)
{
    if (s.value == one)
        s.value = zero;
    else {
        /* place this process in s.queue */;
        /* block this process */;
    }
}
void semSignalB(semaphore s)
{
    if (s.queue is empty())
        s.value = one;
    else {
        /* remove a process P from s.queue */;
        /* place process P on ready list */;
    }
}
```

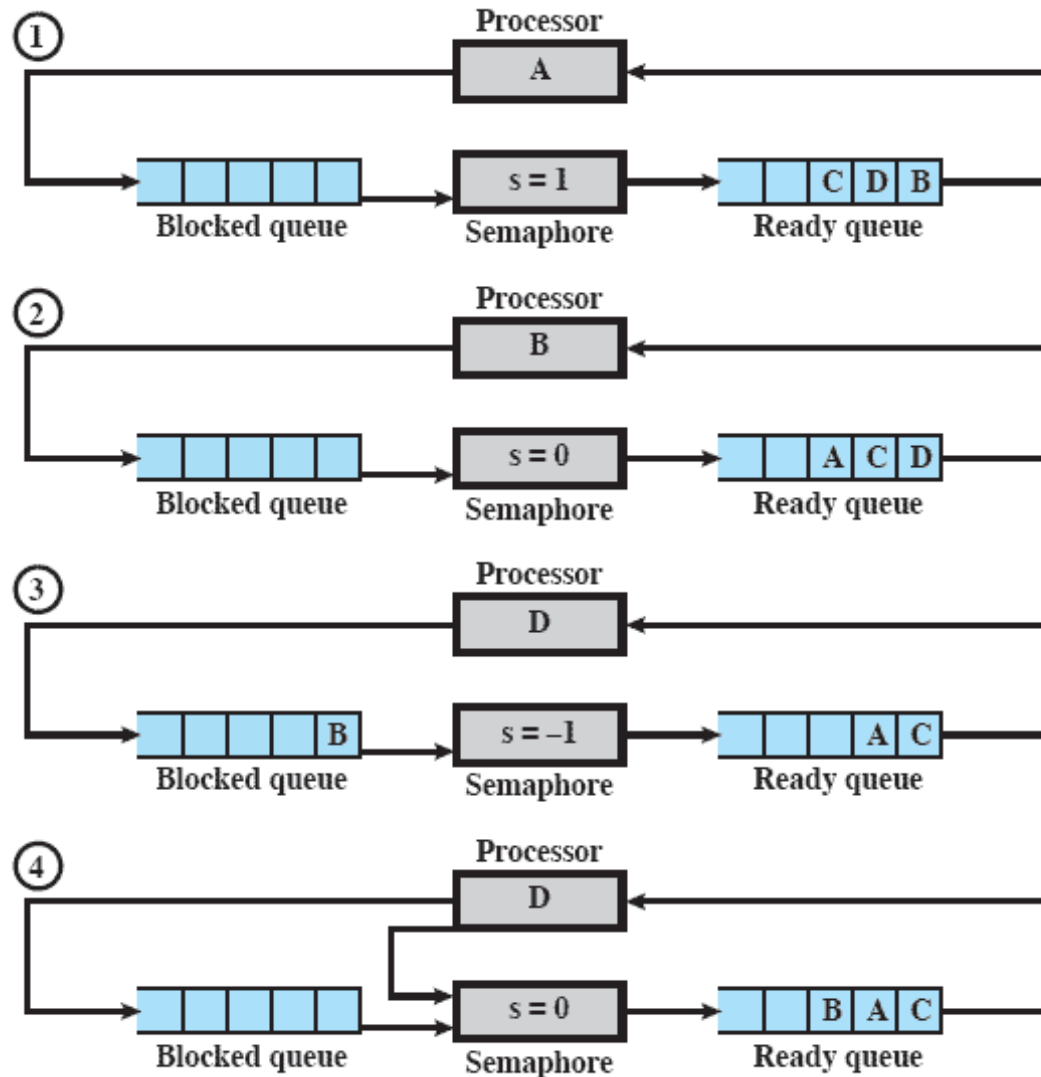
Figure 5.4 A Definition of Binary Semaphore Primitives

Strong/Weak Semaphore



- A queue is used to hold processes waiting on the semaphore
 - In what order are processes removed from the queue?
- **Strong Semaphores** use FIFO
- **Weak Semaphores** don't specify the order of removal from the queue

Example of Strong Semaphore Mechanism



Example of Semaphore Mechanism

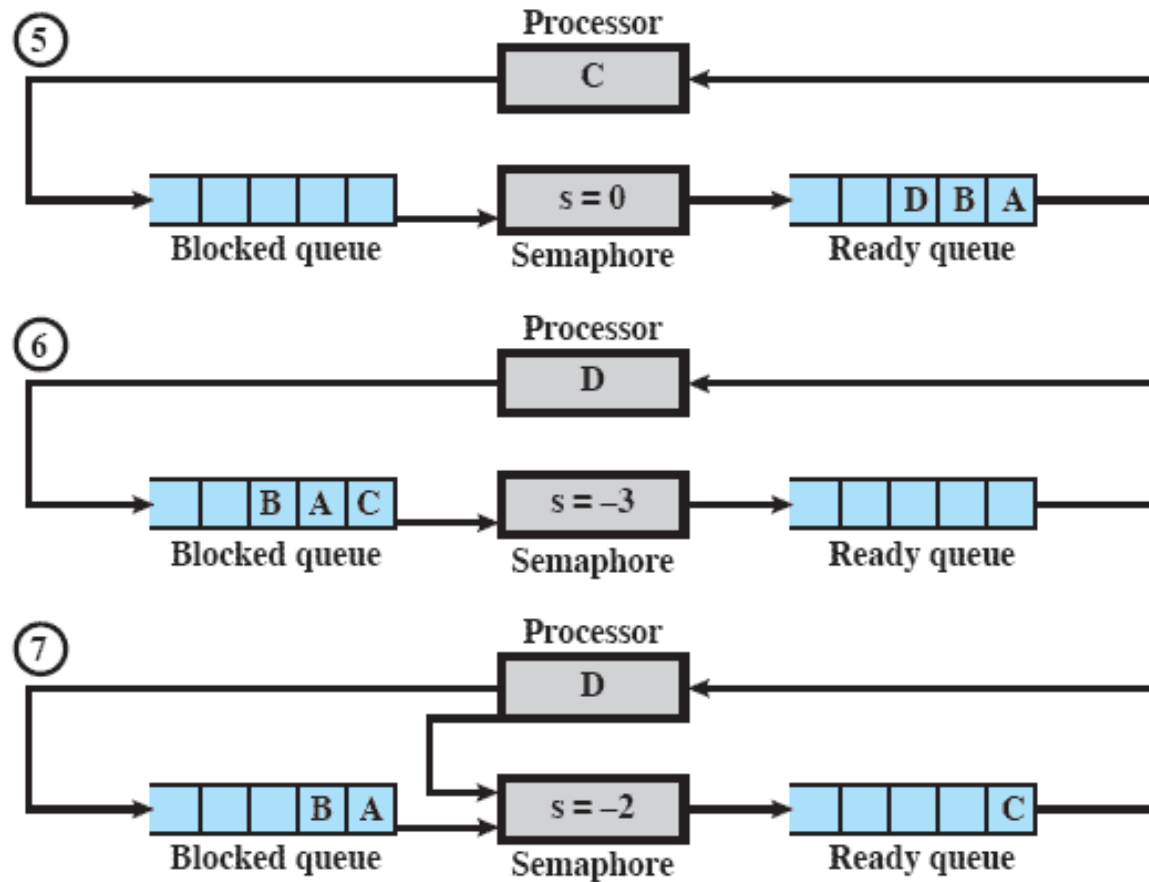


Figure 5.5 Example of Semaphore Mechanism

Mutual Exclusion Using Semaphores



```
/* program mutualexclusion */
const int n = /* number of processes */;
semaphore s = 1;
void P(int i)
{
    while (true) {
        semWait(s);
        /* critical section */;
        semSignal(s);
        /* remainder */;
    }
}
void main()
{
    parbegin (P(1), P(2), . . . , P(n));
}
```

Figure 5.6 Mutual Exclusion Using Semaphores

Processes Using Semaphore

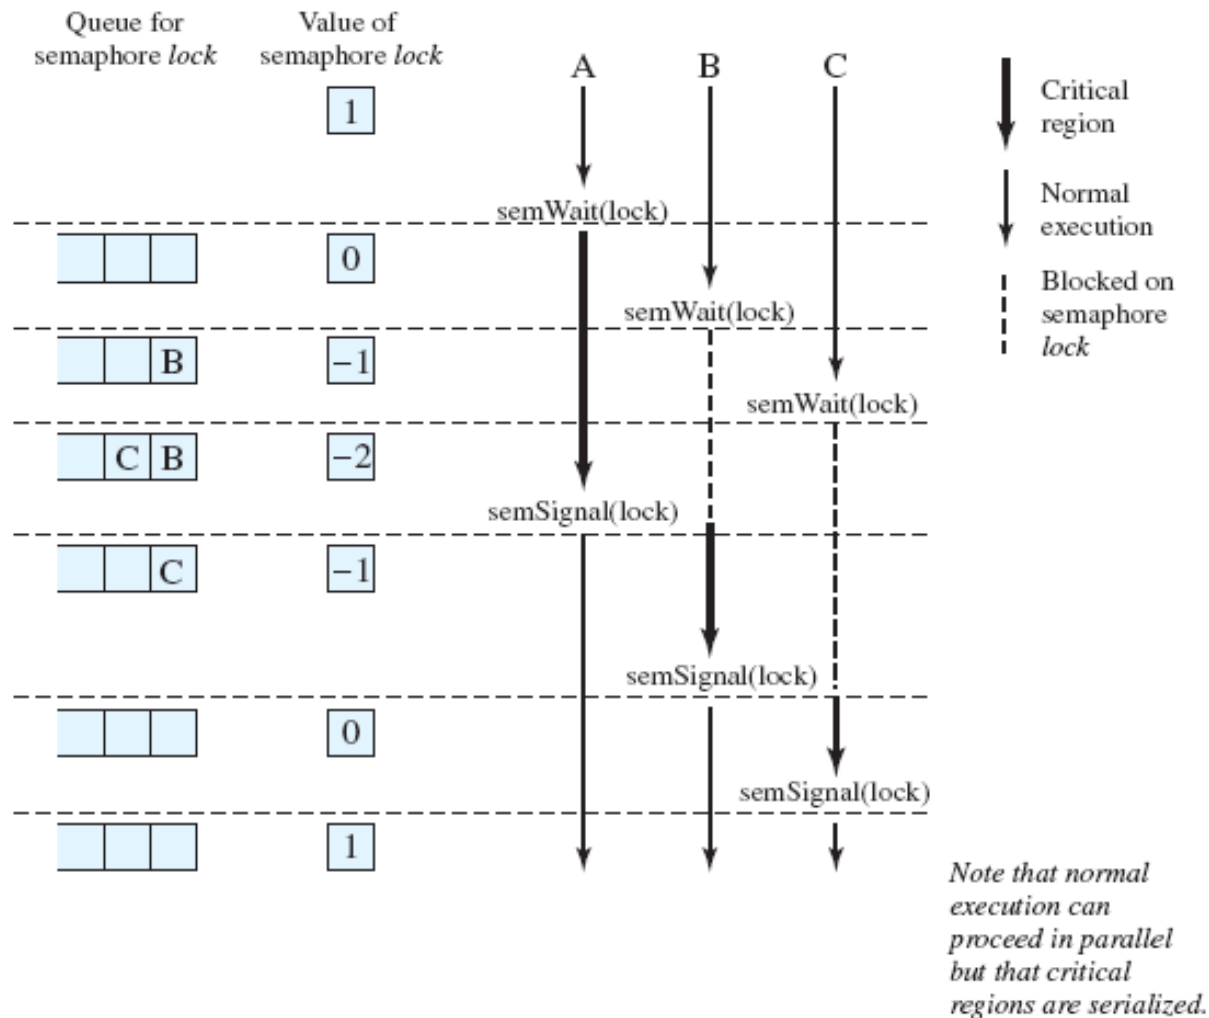


Figure 5.7 Processes Accessing Shared Data Protected by a Semaphore

Producer/Consumer Problem



- General Situation:
 - One or more producers are generating data and placing these in a buffer
 - A single consumer is taking items out of the buffer one at time
 - Only one producer or consumer may access the buffer at any one time
- The Problem:
 - Ensure that the Producer can't add data into full buffer and consumer can't remove data from empty buffer




Functions



- Assume an infinite buffer **b** with a linear array of elements

Producer	Consumer
<pre>while (true) { /* produce item v */ b[in] = v; in++; }</pre>	<pre>while (true) { while (in <= out) /*do nothing */; w = b[out]; out++; /* consume item w */ }</pre>

Buffer



Note: shaded area indicates portion of buffer that is occupied

Figure 5.8 Infinite Buffer for the Producer/Consumer Problem

Incorrect Solution

```
/* program producerconsumer */
int n;
binary_semaphore s = 1, delay = 0;
void producer()
{
    while (true) {
        produce();
        semWaitB(s);
        append();
        n++;
        if (n==1) semSignalB(delay);
        semSignalB(s);
    }
}
void consumer()
{
    semWaitB(delay);
    while (true) {
        semWaitB(s);
        take();
        n--;
        semSignalB(s);
        consume();
        if (n==0) semWaitB(delay);
    }
}
void main()
{
    n = 0;
    parbegin (producer, consumer);
}
```



Possible Scenario



Table 5.4 Possible Scenario for the Program of Figure 5.9

	Producer	Consumer	s	n	Delay
1			1	0	0
2	semWaitB(s)		0	0	0
3	n++		0	1	0
4	if (n==1) (semSignalB(delay))		0	1	1
5	semSignalB(s)		1	1	1
6		semWaitB(delay)	1	1	0
7		semWaitB(s)	0	1	0
8		n--	0	0	0
9		semSignalB(s)	1	0	0
10	semWaitB(s)		0	0	0
11	n++		0	1	0
12	if (n==1) (semSignalB(delay))		0	1	1
13	semSignalB(s)		1	1	1
14		if (n==0) (semWaitB(delay))	1	1	1
15		semWaitB(s)	0	1	1
16		n--	0	0	1
17		semSignalB(s)	1	0	1
18		if (n==0) (semWaitB(delay))	1	0	0
19		semWaitB(s)	0	0	0
20		n--	0	-1	0
21		semiSignlaB(s)	1	-1	0

NOTE: White areas represent the critical section controlled by semaphore s

Correct Solution



```
/* program producerconsumer */
int n;
binary_semaphore s = 1, delay = 0;
void producer()
{
    while (true) {
        produce();
        semWaitB(s);
        append();
        n++;
        if (n==1) semSignalB(delay);
        semSignalB(s);
    }
}
void consumer()
{
    int m; /* a local variable */
    semWaitB(delay);
    while (true) {
        semWaitB(s);
        take();
        n--;
        m = n;
        semSignalB(s);
        consume();
        if (m==0) semWaitB(delay);
    }
}
void main()
{
    n = 0;
    parbegin (producer, consumer);
}
```

Semaphores

```
/* program producerconsumer */
semaphore n = 0, s = 1;
void producer()
{
    while (true) {
        produce();
        semWait(s);
        append();
        semSignal(s);
        semSignal(n);
    }
}
void consumer()
{
    while (true) {
        semWait(n);
        semWait(s);
        take();
        semSignal(s);
        consume();
    }
}
void main()
{
    parbegin (producer, consumer);
}
```

Figure 5.11 A Solution to the Infinite-Buffer Producer/Consumer Problem Using Semaphores

Bounded Buffer



Block on:	Unblock on:
Producer: insert in full buffer	Consumer: item inserted
Consumer: remove from empty buffer	Producer: item removed





Figure 5.12 Finite Circular Buffer for the Producer/Consumer Problem

Semaphores

```
/* program boundedbuffer */
const int sizeofbuffer = /* buffer size */;
semaphore s = 1, n = 0, e = sizeofbuffer;
void producer()
{
    while (true) {
        produce();
        semWait(e);
        semWait(s);
        append();
        semSignal(s);
        semSignal(n);
    }
}
void consumer()
{
    while (true) {
        semWait(n);
        semWait(s);
        take();
        semSignal(s);
        semSignal(e);
        consume();
    }
}
void main()
{
    parbegin (producer, consumer);
}
```



Functions in a Bounded Buffer



- .

Producer	Consumer
<pre>while (true) { /* produce item v */ while ((in + 1) % n == out) /* do nothing */; b[in] = v; in = (in + 1) % n</pre>	<pre>while (true) { while (in == out) /* do nothing */; w = b[out]; out = (out + 1) % n; /* consume item</pre>

Demonstration Animations



- Producer/Consumer
 - Illustrates the operation of a producer-consumer buffer.
- Bounded-Buffer Problem Using Semaphores
 - Demonstrates the bounded-buffer consumer/producer problem using semaphores.

Roadmap

- Principals of Concurrency
- Mutual Exclusion: Hardware Support
- Semaphores
- Monitors
- • Message Passing
- Readers/Writers Problem



Monitors

- The monitor is a programming-language construct that provides equivalent functionality to that of semaphores and that is easier to control.
- Implemented in a number of programming languages, including
 - Concurrent Pascal, Pascal-Plus,
 - Modula-2, Modula-3, and Java.



Chief characteristics

- Local data variables are accessible only by the monitor
- Process enters monitor by invoking one of its procedures
- Only one process may be executing in the monitor at a time

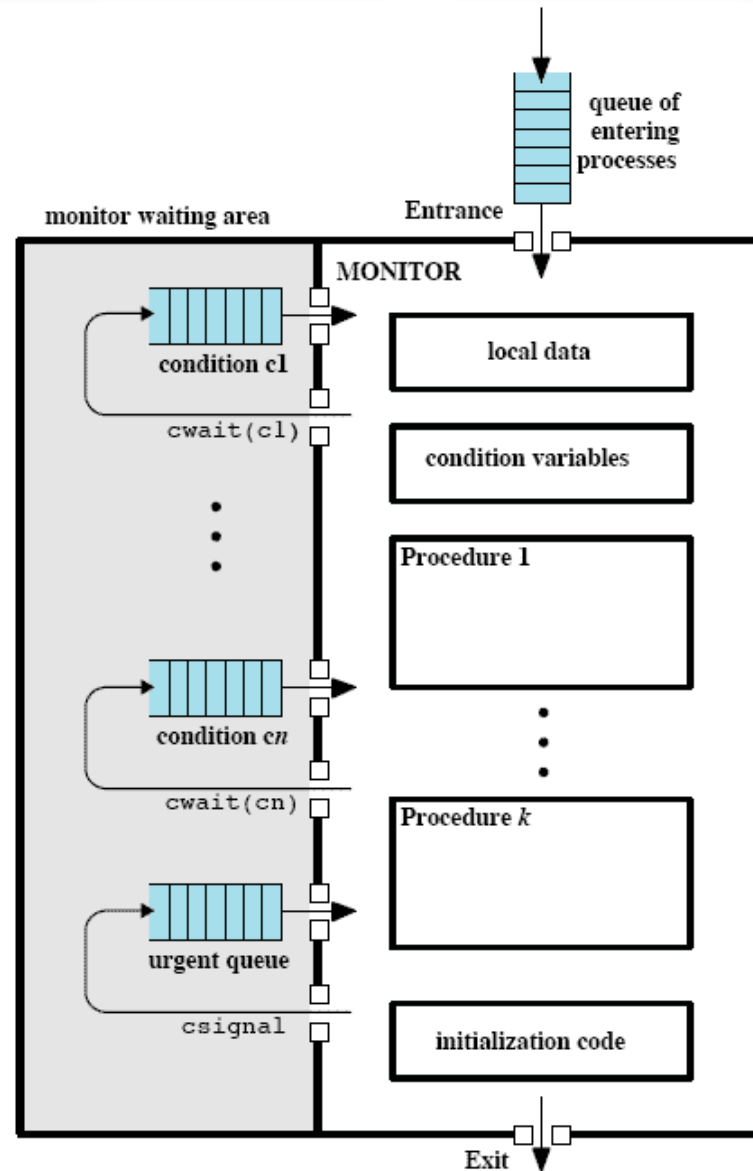


Synchronization



- Synchronisation achieved by **condition variables** within a monitor
 - only accessible by the monitor.
- Monitor Functions:
 - **Cwait(c)**: Suspend execution of the calling process on condition *c*
 - **Csignal(c)** Resume execution of some process blocked after a cwait on the same condition

Structure of a Monitor



Bounded Buffer Solution Using Monitor



```
/* program producerconsumer */
monitor boundedbuffer;
char buffer [N];                               /* space for N items */
int nextin, nextout;                            /* buffer pointers */
int count;                                     /* number of items in buffer */
cond notfull, notempty;                       /* condition variables for synchronization */

void append (char x)
{
    if (count == N) cwait(notfull);           /* buffer is full; avoid overflow */
    buffer[nextin] = x;
    nextin = (nextin + 1) % N;
    count++;
    /* one more item in buffer */
    csignal(notempty);                        /* resume any waiting consumer */
}

void take (char x)
{
    if (count == 0) cwait(notempty);          /* buffer is empty; avoid underflow */
    x = buffer[nextout];
    nextout = (nextout + 1) % N;
    count--;
    csignal(notfull);                         /* one fewer item in buffer */
                                           /* resume any waiting producer */
}

/* monitor body */
nextin = 0; nextout = 0; count = 0;          /* buffer initially empty */
}
```

Solution Using Monitor



```
void producer()
{
    char x;
    while (true) {
        produce(x);
        append(x);
    }
}
void consumer()
{
    char x;
    while (true) {
        take(x);
        consume(x);
    }
}
void main()
{
    parbegin (producer, consumer);
}
```

Bounded Buffer Monitor



```
void append (char x)
{
    while(count == N) cwait(notfull);    /* buffer is full; avoid overflow */
    buffer[nextin] = x;
    nextin = (nextin + 1) % N;
    count++;                             /* one more item in buffer */
    cnotify(notempty);                   /* notify any waiting consumer */
}

void take (char x)
{
    while(count == 0) cwait(notempty); /* buffer is empty; avoid underflow */
    x = buffer[nextout];
    nextout = (nextout + 1) % N;
    count--;                             /* one fewer item in buffer */
    cnotify(notfull);                    /* notify any waiting producer */
}
```

Figure 5.17 Bounded Buffer Monitor Code for Mesa Monitor

Roadmap

- Principals of Concurrency
- Mutual Exclusion: Hardware Support
- Semaphores
- Monitors
- Message Passing
- • Readers/Writers Problem



Process Interaction



- When processes interact with one another, two fundamental requirements must be satisfied:
 - synchronization and
 - communication.
- Message Passing is one solution to the second requirement
 - Added bonus: It works with shared memory *and* with distributed systems

Message Passing



- The actual function of message passing is normally provided in the form of a pair of primitives:
 - send (destination, message)
 - receive (source, message)

Synchronization



- Communication requires synchronization
 - Sender must send before receiver can receive
- What happens to a process after it issues a send or receive primitive?
 - Sender and receiver may or may not be blocking (waiting for message)

Blocking send, Blocking receive

- Both sender and receiver are blocked until message is delivered
- Known as a *rendezvous*
- Allows for tight synchronization between processes.




Non-blocking Send



- More natural for many concurrent programming tasks.
- Nonblocking send, blocking receive
 - Sender continues on
 - Receiver is blocked until the requested message arrives
- Nonblocking send, nonblocking receive
 - Neither party is required to wait


Addressing

- Sendin process need to be able to specify which process should receive the message
 - Direct addressing
 - Indirect Addressing



Direct Addressing

- Send primitive includes a specific identifier of the destination process
- Receive primitive could know ahead of time which process a message is expected
- Receive primitive could use source parameter to return a value when the receive operation has been performed



Indirect addressing

- Messages are sent to a shared data structure consisting of queues
- Queues are called *mailboxes*
- One process sends a message to the mailbox and the other process picks up the message from the mailbox



Indirect Process Communication

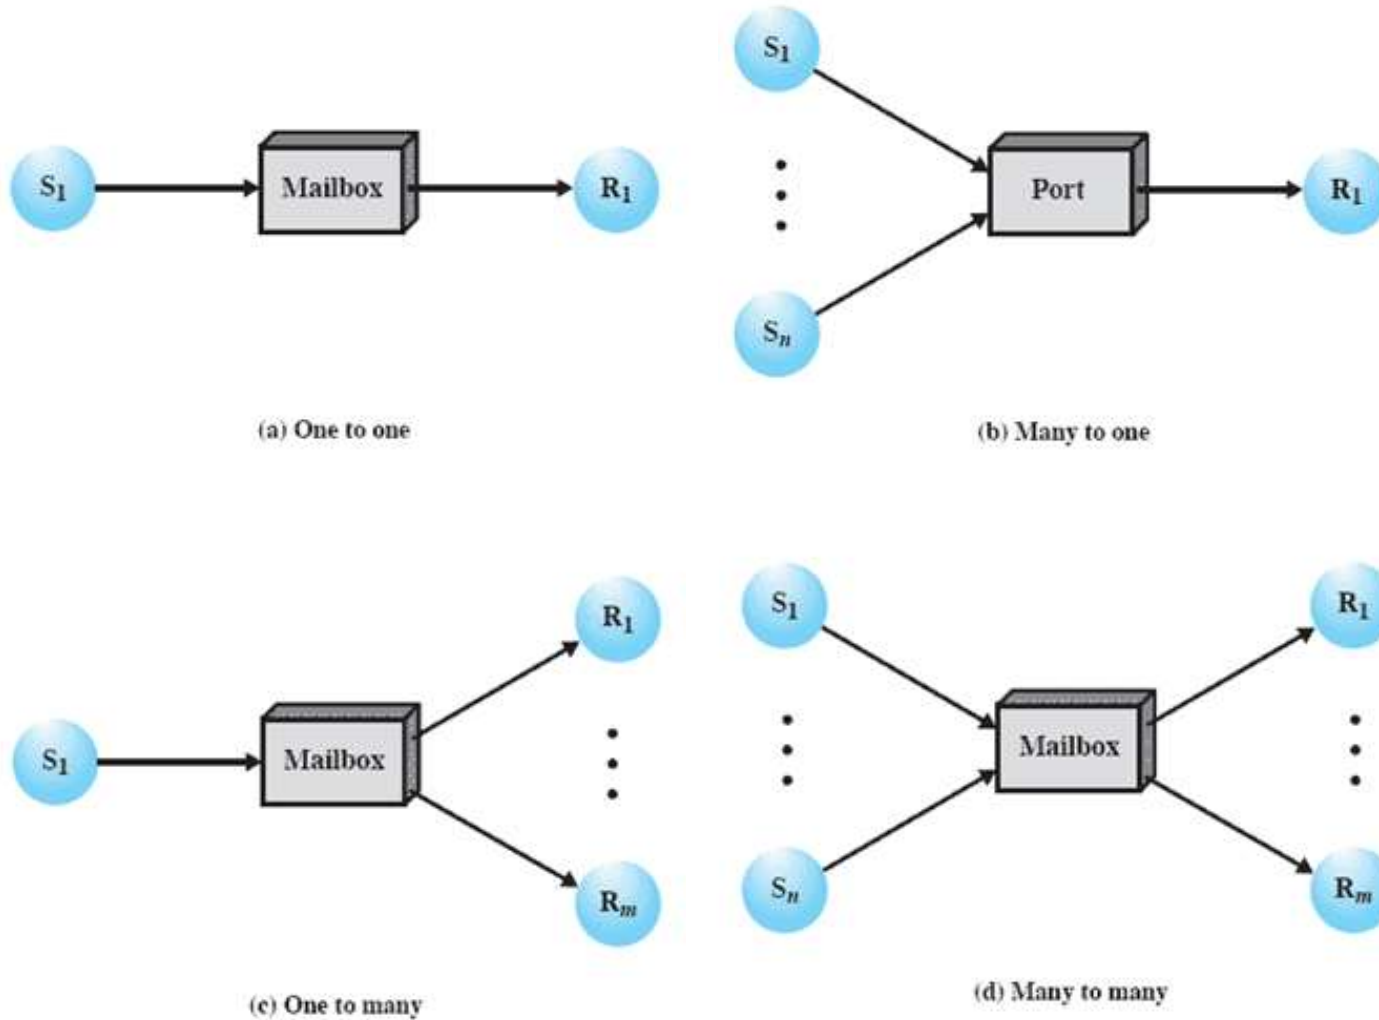


Figure 5.18 Indirect Process Communication

General Message Format




Figure 5.19 General Message Format

Mutual Exclusion Using Messages



```
/* program mutualexclusion */
const int n = /* number of processes */;
void P(int i)
{
    message msg;
    while (true) {
        receive (box, msg);
        /* critical section */;
        send (box, msg);
        /* remainder */;
    }
}
void main()
{
    create mailbox (box);
    send (box, null);
    parbegin (P(1), P(2), . . . , P(n));
}
```

Figure 5.20 Mutual Exclusion Using Messages

Producer/Consumer Messages



```
const int
    capacity = /* buffering capacity */ ;
    null = /* empty message */ ;
int i;
void producer()
{   message pmsg;
    while (true) {
        receive (mayproduce, pmsg);
        pmsg = produce();
        send (mayconsume, pmsg);
    }
}
void consumer()
{   message cmsg;
    while (true) {
        receive (mayconsume, cmsg);
        consume (cmsg);
        send (mayproduce, null);
    }
}

void main()
{
    create_mailbox (mayproduce);
    create_mailbox (mayconsume);
    for (int i = 1; i <= capacity; i++) send (mayproduce, null);
    parbegin (producer, consumer);
}
```

Roadmap

- Principals of Concurrency
- Mutual Exclusion: Hardware Support
- Semaphores
- Monitors
- Message Passing
- Readers/Writers Problem



Readers/Writers Problem



- A data area is shared among many processes
 - Some processes only read the data area, some only write to the area
- Conditions to satisfy:
 1. Multiple readers may read the file at once.
 2. Only one writer at a time may write
 3. If a writer is writing to the file, no reader may read it.

interaction of readers and writers.



Readers have Priority



```
/* program readersandwriters */
int readcount;
semaphore x = 1, wsem = 1;
void reader()
{
    while (true) {
        semWait (x);
        readcount++;
        if (readcount == 1) semWait (wsem);
        semSignal (x);
        READUNIT();
        semWait (x);
        readcount--;
        if (readcount == 0) semSignal (wsem);
        semSignal (x);
    }
}
void writer()
{
    while (true) {
        semWait (wsem);
        WRITEUNIT();
        semSignal (wsem);
    }
}

void main()
{
    readcount = 0;
    parbegin (reader, writer);
}
```

Writers have Priority



```
/* program readersandwriters */
int  readcount, writecount;
semaphore x = 1, y = 1, z = 1, wsem = 1, rsem = 1;
void reader()
{
    while (true) {
        semWait (z);
        semWait (rsem);
        semWait (x);
        readcount++;
        if (readcount == 1) semWait (wsem);
        semSignal (x);
        semSignal (rsem);
        semSignal (z);
        READUNIT();
        semWait (x);
        readcount--;
        if (readcount == 0) semSignal (wsem);
        semSignal (x);
    }
}
```


Writers have Priority



```
void writer ()
{
    while (true) {
        semWait (y);
        writecount++;
        if (writecount == 1) semWait (rsem);
        semSignal (y);
        semWait (wsem);
        WRITEUNIT();
        semSignal (wsem);
        semWait (y);
        writecount--;
        if (writecount == 0) semSignal (rsem);
        semSignal (y);
    }
}

void main()
{
    readcount = writecount = 0;
    parbegin (reader, writer);
}
```

Message Passing

```
void reader(int i)
{
    message rmsg;
    while (true) {
        rmsg = i;
        send (readrequest, rmsg);
        receive (mbox[i], rmsg);
        READUNIT ();
        rmsg = i;
        send (finished, rmsg);
    }
}

void writer(int j)
{
    message rmsg;
    while(true) {
        rmsg = j;
        send (writerequest, rmsg);
        receive (mbox[j], rmsg);
        WRITEUNIT ();
        rmsg = j;
        send (finished, rmsg);
    }
}
```

```
void controller()
{
    while (true)
    {
        if (count > 0) {
            if (!empty (finished)) {
                receive (finished, msg);
                count++;
            }
            else if (!empty (writerequest)) {
                receive (writerequest, msg);
                writer_id = msg.id;
                count = count - 100;
            }
            else if (!empty (readrequest)) {
                receive (readrequest, msg);
                count--;
                send (msg.id, "OK");
            }
        }
        if (count == 0) {
            send (writer id, "OK");
            receive (finished, msg);
            count = 100;
        }
        while (count < 0) {
            receive (finished, msg);
            count++;
        }
    }
}
```

Message Passing



```
void controller()
{
    while (true)
    {
        if (count > 0) {
            if (!empty (finished)) {
                receive (finished, msg);
                count++;
            }
            else if (!empty (writerequest)) {
                receive (writerequest, msg);
                writer_id = msg.id;
                count = count - 100;
            }
            else if (!empty (readrequest)) {
                receive (readrequest, msg);
                count--;
                send (msg.id, "OK");
            }
        }
        if (count == 0) {
            send (writer id, "OK");
            receive (finished, msg);
            count = 100;
        }
        while (count < 0) {
            receive (finished, msg);
            count++;
        }
    }
}
```

Deadlock

- Permanent blocking of a set of processes that either compete for system resources or communicate with each other
- No efficient solution
- Involve conflicting needs for resources by two or more processes



Deadlock




Figure 6.1 Illustration of Deadlock

Deadlock




Figure 6.2 Example of Deadlock

Deadlock




Figure 6.3 Example of No Deadlock [BACO03]

Reusable Resources

- Used by only one process at a time and not depleted by that use
- Processes obtain resources that they later release for reuse by other processes



Reusable Resources



- Processors, I/O channels, main and secondary memory, devices, and data structures such as files, databases, and semaphores
- Deadlock occurs if each process holds one resource and requests the other

Reusable Resources




Process P		Process Q	
Step	Action	Step	Action
p ₀	Request (D)	q ₀	Request (T)
p ₁	Lock (D)	q ₁	Lock (T)
p ₂	Request (T)	q ₂	Request (D)
p ₃	Lock (T)	q ₃	Lock (D)
p ₄	Perform function	q ₄	Perform function
p ₅	Unlock (D)	q ₅	Unlock (T)
p ₆	Unlock (T)	q ₆	Unlock (D)

Figure 6.4 Example of Two Processes Competing for Reusable Resources

Reusable Resources



- Space is available for allocation of 200Kbytes, and the following sequence of events occur



- ~~Deadlock occurs if both processes~~ progress to their second request

Consumable Resources





- Created (produced) and destroyed (consumed)
- Interrupts, signals, messages, and information in I/O buffers
- Deadlock may occur if a Receive message is blocking
- May take a rare combination of events to cause deadlock

Example of Deadlock



- Deadlock occurs if receives blocking



Resource Allocation Graphs

- Directed graph that depicts a state of the system of resources and processes



(a) Resource is requested



(b) Resource is held

Conditions for Deadlock




- Mutual exclusion
 - Only one process may use a resource at a time
- Hold-and-wait
 - A process may hold allocated resources while awaiting assignment of others

Conditions for Deadlock




- No preemption
 - No resource can be forcibly removed from a process holding it
- Circular wait
 - A closed chain of processes exists, such that each process holds at least one resource needed by the next process in the chain

Resource Allocation Graphs



(c) Circular wait



(d) No deadlock

Resource Allocation Graphs




Figure 6.6 Resource Allocation Graph for Figure 6.1b

Possibility of Deadlock

- Mutual Exclusion
- No preemption
- Hold and wait



Existence of Deadlock

- Mutual Exclusion
- No preemption
- Hold and wait
- Circular wait



Deadlock Prevention



- Mutual Exclusion
 - Must be supported by the OS
- Hold and Wait
 - Require a process request all of its required resources at one time

Deadlock Prevention



- No Preemption
 - Process must release resource and request again
 - OS may preempt a process to require it releases its resources
- Circular Wait
 - Define a linear ordering of resource types

Deadlock Avoidance



- A decision is made dynamically whether the current resource allocation request will, if granted, potentially lead to a deadlock
- Requires knowledge of future process requests

Two Approaches to Deadlock Avoidance



- Do not start a process if its demands might lead to deadlock
- Do not grant an incremental resource request to a process if this allocation might lead to deadlock

Resource Allocation Denial



- Referred to as the banker's algorithm
- State of the system is the current allocation of resources to process
- Safe state is where there is at least one sequence that does not result in deadlock
- Unsafe state is a state that is not safe

Determination of a Safe State



	R1	R2	R3
P1	3	2	2
P2	6	1	3
P3	3	1	4
P4	4	2	2

Claim matrix **C**

	R1	R2	R3
P1	1	0	0
P2	6	1	2
P3	2	1	1
P4	0	0	2

Allocation matrix **A**

	R1	R2	R3
P1	2	2	2
P2	0	0	1
P3	1	0	3
P4	4	2	0

C - A

R1	R2	R3
9	3	6

Resource vector **R**

R1	R2	R3
0	1	1

Available vector **V**

(a) Initial state

Determination of a Safe State



	R1	R2	R3
P1	3	2	2
P2	0	0	0
P3	3	1	4
P4	4	2	2

Claim matrix C

	R1	R2	R3
P1	1	0	0
P2	0	0	0
P3	2	1	1
P4	0	0	2

Allocation matrix A

	R1	R2	R3
P1	2	2	2
P2	0	0	0
P3	1	0	3
P4	4	2	0

C - A

R1	R2	R3
9	3	6

Resource vector R

R1	R2	R3
6	2	3

Available vector V

(b) P2 runs to completion

Determination of a Safe State



	R1	R2	R3
P1	0	0	0
P2	0	0	0
P3	3	1	4
P4	4	2	2

Claim matrix **C**

	R1	R2	R3
P1	0	0	0
P2	0	0	0
P3	2	1	1
P4	0	0	2

Allocation matrix **A**

	R1	R2	R3
P1	0	0	0
P2	0	0	0
P3	1	0	3
P4	4	2	0

C - A

R1	R2	R3
9	3	6

Resource vector **R**

R1	R2	R3
7	2	3

Available vector **V**

(c) **P1 runs to completion**

Determination of a Safe State



	R1	R2	R3
P1	0	0	0
P2	0	0	0
P3	0	0	0
P4	4	2	2

Claim matrix **C**

	R1	R2	R3
P1	0	0	0
P2	0	0	0
P3	0	0	0
P4	0	0	2

Allocation matrix **A**

	R1	R2	R3
P1	0	0	0
P2	0	0	0
P3	0	0	0
P4	4	2	0

C - A

R1	R2	R3
9	3	6

Resource vector **R**

R1	R2	R3
9	3	4

Available vector **V**

(d) **P3 runs to completion**

Determination of an Unsafe State



	R1	R2	R3
P1	3	2	2
P2	6	1	3
P3	3	1	4
P4	4	2	2

Claim matrix **C**

	R1	R2	R3
P1	1	0	0
P2	5	1	1
P3	2	1	1
P4	0	0	2

Allocation matrix **A**

	R1	R2	R3
P1	2	2	2
P2	1	0	2
P3	1	0	3
P4	4	2	0

C - A

R1	R2	R3
9	3	6

Resource vector **R**

R1	R2	R3
1	1	2

Available vector **V**

(a) Initial state

	R1	R2	R3
P1	3	2	2
P2	6	1	3
P3	3	1	4
P4	4	2	2

Claim matrix **C**

	R1	R2	R3
P1	2	0	1
P2	5	1	1
P3	2	1	1
P4	0	0	2

Allocation matrix **A**

	R1	R2	R3
P1	1	2	1
P2	1	0	2
P3	1	0	3
P4	4	2	0

C - A

R1	R2	R3
9	3	6

Resource vector **R**

R1	R2	R3
0	1	1

Available vector **V**

(b) P1 requests one unit each of R1 and R3

Deadlock Avoidance Logic

```
struct state {
    int resource[m];
    int available[m];
    int claim[n][m];
    int alloc[n][m];
}
```

(a) global data structures

```
if (alloc [i,*] + request [*] > claim [i,*])
    < error >; /* total request > claim*/
else if (request [*] > available [*])
    < suspend process >;
else { /* simulate alloc */
    < define newstate by:
    alloc [i,*] = alloc [i,*] + request [*];
    available [*] = available [*] - request [*] >;
}
if (safe (newstate))
    < carry out allocation >;
else {
    < restore original state >;
    < suspend process >;
}
```

(b) resource alloc algorithm

Deadlock Avoidance Logic



```
boolean safe (state S) {
    int currentavail[m];
    process rest[<number of processes>];
    currentavail = available;
    rest = {all processes};
    possible = true;
    while (possible) {
        <find a process Pk in rest such that
            claim [k,*] - alloc [k,*] <= currentavail;>
        if (found) {                                /* simulate execution of Pk */
            currentavail = currentavail + alloc [k,*];
            rest = rest - {Pk};
        }
        else possible = false;
    }
    return (rest == null);
}
```

(c) test for safety algorithm (banker's algorithm)

Figure 6.9 Deadlock Avoidance Logic

Deadlock Avoidance



- Maximum resource requirement must be stated in advance
- Processes under consideration must be independent; no synchronization requirements
- There must be a fixed number of resources to allocate
- No process may exit while holding resources

Deadlock Detection



	R1	R2	R3	R4	R5
P1	0	1	0	0	1
P2	0	0	1	0	1
P3	0	0	0	0	1
P4	1	0	1	0	1

Request matrix Q

	R1	R2	R3	R4	R5
P1	1	0	1	1	0
P2	1	1	0	0	0
P3	0	0	0	1	0
P4	0	0	0	0	0

Allocation matrix A

R1	R2	R3	R4	R5
2	1	1	2	1

Resource vector

R1	R2	R3	R4	R5
0	0	0	0	1

Allocation vector

Figure 6.10 Example for Deadlock Detection

Strategies Once Deadlock Detected



- Abort all deadlocked processes
- Back up each deadlocked process to some previously defined checkpoint, and restart all process
 - Original deadlock may occur

Strategies Once Deadlock Detected



- Successively abort deadlocked processes until deadlock no longer exists
- Successively preempt resources until deadlock no longer exists

Advantages and Disadvantages



Table 6.1 Summary of Deadlock Detection, Prevention, and Avoidance Approaches for Operating Systems [ISLO80]

Approach	Resource Allocation Policy	Different Schemes	Major Advantages	Major Disadvantages
Prevention	Conservative; undercommits resources	Requesting all resources at once	<ul style="list-style-type: none"> • Works well for processes that perform a single burst of activity • No preemption necessary 	<ul style="list-style-type: none"> • Inefficient • Delays process initiation • Future resource requirements must be known by processes
		Preemption	<ul style="list-style-type: none"> • Convenient when applied to resources whose state can be saved and restored easily 	<ul style="list-style-type: none"> • Preempts more often than necessary
		Resource ordering	<ul style="list-style-type: none"> • Feasible to enforce via compile-time checks • Needs no run-time computation since problem is solved in system design 	<ul style="list-style-type: none"> • Disallows incremental resource requests
Avoidance	Midway between that of detection and prevention	Manipulate to find at least one safe path	<ul style="list-style-type: none"> • No preemption necessary 	<ul style="list-style-type: none"> • Future resource requirements must be known by OS • Processes can be blocked for long periods
Detection	Very liberal; requested resources are granted where possible	Invoke periodically to test for deadlock	<ul style="list-style-type: none"> • Never delays process initiation • Facilitates online handling 	<ul style="list-style-type: none"> • Inherent preemption losses

Dining Philosophers Problem




Figure 6.11 Dining Arrangement for Philosophers

Dining Philosophers Problem



```
/* program      diningphilosophers */
semaphore fork [5] = {1};
int i;
void philosopher (int i)
{
    while (true) {
        think();
        wait (fork[i]);
        wait (fork [(i+1) mod 5]);
        eat();
        signal(fork [(i+1) mod 5]);
        signal(fork[i]);
    }
}
void main()
{
    parbegin (philosopher (0), philosopher (1), philosopher
(2),
            philosopher (3), philosopher (4));
}
```

Figure 6.12 A First Solution to the Dining Philosophers Problem

Dining Philosophers Problem

```
/* program diningphilosophers */
semaphore fork[5] = {1};
semaphore room = {4};
int i;
void philosopher (int i)
{
    while (true) {
        think();
        wait (room);
        wait (fork[i]);
        wait (fork [(i+1) mod 5]);
        eat();
        signal (fork [(i+1) mod 5]);
        signal (fork[i]);
        signal (room);
    }
}

void main()
{
    parbegin (philosopher (0), philosopher (1), philosopher (2),
             philosopher (3), philosopher (4));
}
```

Figure 6.13 A Second Solution to the Dining Philosophers Problem

Dining Philosophers Problem



```
monitor dining_controller;
cond ForkReady[5];          /* condition variable for synchronization */
boolean fork[5] = {true};   /* availability status of each fork */

void get_forks(int pid)     /* pid is the philosopher id number */
{
    int left = pid;
    int right = (++pid) % 5;
    /*grant the left fork*/
    if (!fork(left))
        cwait(ForkReady[left]);          /* queue on condition variable */
    fork(left) = false;
    /*grant the right fork*/
    if (!fork(right))
        cwait(ForkReady[right]);        /* queue on condition variable */
    fork(right) = false;
}

void release_forks(int pid)
{
    int left = pid;
    int right = (++pid) % 5;
    /*release the left fork*/
    if (empty(ForkReady[left])          /*no one is waiting for this fork */
        fork(left) = true;
    else                                  /* awaken a process waiting on this fork */
        csignal(ForkReady[left]);
    /*release the right fork*/
    if (empty(ForkReady[right])         /*no one is waiting for this fork */
        fork(right) = true;
    else                                  /* awaken a process waiting on this fork */
        csignal(ForkReady[right]);
}
```

Dining Philosophers Problem



```
void philosopher[k=0 to 4]          /* the five philosopher clients */
{
    while (true) {
        <think>;
        get forks(k);                /* client requests two forks via monitor */
        <eat spaghetti>;
        release forks(k);           /* client releases forks via the monitor */
    }
}
```

Figure 6.14 A Solution to the Dining Philosophers Problem Using a Monitor

UNIX Concurrency Mechanisms



- Pipes
- Messages
- Shared memory
- Semaphores
- Signals

UNIX Signals



Value	Name	Description
01	SIGHUP	Hang up; sent to process when kernel assumes that the user of that process is doing no useful work
02	SIGINT	Interrupt
03	SIGQUIT	Quit; sent by user to induce halting of process and production of core dump
04	SIGILL	Illegal instruction
05	SIGTRAP	Trace trap; triggers the execution of code for process tracing
06	SIGIOT	IOT instruction
07	SIGEMT	EMT instruction
08	SIGFPE	Floating-point exception
09	SIGKILL	Kill; terminate process
10	SIGBUS	Bus error
11	SIGSEGV	Segmentation violation; process attempts to access location outside its virtual address space
12	SIGSYS	Bad argument to system call
13	SIGPIPE	Write on a pipe that has no readers attached to it
14	SIGALRM	Alarm clock; issued when a process wishes to receive a signal after a period of time
15	SIGTERM	Software termination
16	SIGUSR1	User-defined signal 1
17	SIGUSR2	User-defined signal 2
18	SIGCHLD	Death of a child
19	SIGPWR	Power failure

Linux Kernel Concurrency Mechanism



- Includes all the mechanisms found in UNIX
- Atomic operations execute without interruption and without interference

Linux Atomic Operations



Table 6.3 Linux Atomic Operations

Atomic Integer Operations	
<code>ATOMIC_INIT (int i)</code>	At declaration: initialize an <code>atomic_t</code> to <code>i</code>
<code>int atomic_read(atomic_t *v)</code>	Read integer value of <code>v</code>
<code>void atomic_set(atomic_t *v, int i)</code>	Set the value of <code>v</code> to integer <code>i</code>
<code>void atomic_add(int i, atomic_t *v)</code>	Add <code>i</code> to <code>v</code>
<code>void atomic_sub(int i, atomic_t *v)</code>	Subtract <code>i</code> from <code>v</code>
<code>void atomic_inc(atomic_t *v)</code>	Add 1 to <code>v</code>
<code>void atomic_dec(atomic_t *v)</code>	Subtract 1 from <code>v</code>
<code>int atomic_sub_and_test(int i, atomic_t *v)</code>	Subtract <code>i</code> from <code>v</code> ; return 1 if the result is zero; return 0 otherwise
<code>int atomic_add_negative(int i, atomic_t *v)</code>	Add <code>i</code> to <code>v</code> ; return 1 if the result is negative; return 0 otherwise (used for implementing semaphores)
<code>int atomic_dec_and_test(atomic_t *v)</code>	Subtract 1 from <code>v</code> ; return 1 if the result is zero; return 0 otherwise
<code>int atomic_inc_and_test(atomic_t *v)</code>	Add 1 to <code>v</code> ; return 1 if the result is zero; return 0 otherwise

Linux Atomic Operations



Atomic Bitmap Operations	
<code>void set_bit(int nr, void *addr)</code>	Set bit nr in the bitmap pointed to by addr
<code>void clear_bit(int nr, void *addr)</code>	Clear bit nr in the bitmap pointed to by addr
<code>void change_bit(int nr, void *addr)</code>	Invert bit nr in the bitmap pointed to by addr
<code>int test_and_set_bit(int nr, void *addr)</code>	Set bit nr in the bitmap pointed to by addr; return the old bit value
<code>int test_and_clear_bit(int nr, void *addr)</code>	Clear bit nr in the bitmap pointed to by addr; return the old bit value
<code>int test_and_change_bit(int nr, void *addr)</code>	Invert bit nr in the bitmap pointed to by addr; return the old bit value
<code>int test_bit(int nr, void *addr)</code>	Return the value of bit nr in the bitmap pointed to by addr

Linux Spinlocks



<code>void spin_lock(spinlock_t *lock)</code>	Acquires the specified lock, spinning if needed until it is available
<code>void spin_lock_irq(spinlock_t *lock)</code>	Like <code>spin_lock</code> , but also disables interrupts on the local processor
<code>void spin_lock_irqsave(spinlock_t *lock, unsigned long flags)</code>	Like <code>spin_lock_irq</code> , but also saves the current interrupt state in flags
<code>void spin_lock_bh(spinlock_t *lock)</code>	Like <code>spin_lock</code> , but also disables the execution of all bottom halves
<code>void spin_unlock(spinlock_t *lock)</code>	Releases given lock
<code>void spin_unlock_irq(spinlock_t *lock)</code>	Releases given lock and enables local interrupts
<code>void spin_unlock_irqrestore(spinlock_t *lock, unsigned long flags)</code>	Releases given lock and restores local interrupts to given previous state
<code>void spin_unlock_bh(spinlock_t *lock)</code>	Releases given lock and enables bottom halves
<code>void spin_lock_init(spinlock_t *lock)</code>	Initializes given spinlock
<code>int spin_trylock(spinlock_t *lock)</code>	Tries to acquire specified lock; returns nonzero if lock is currently held and zero otherwise
<code>int spin_is_locked(spinlock_t *lock)</code>	Returns nonzero if lock is currently held and zero otherwise

Linux Semaphores



Traditional Semaphores	
<code>void sema_init(struct semaphore *sem, int count)</code>	Initializes the dynamically created semaphore to the given count
<code>void init_MUTEX(struct semaphore *sem)</code>	Initializes the dynamically created semaphore with a count of 1 (initially unlocked)
<code>void init_MUTEX_LOCKED(struct semaphore *sem)</code>	Initializes the dynamically created semaphore with a count of 0 (initially locked)
<code>void down(struct semaphore *sem)</code>	Attempts to acquire the given semaphore, entering uninterruptible sleep if semaphore is unavailable
<code>int down_interruptible(struct semaphore *sem)</code>	Attempts to acquire the given semaphore, entering interruptible sleep if semaphore is unavailable; returns <code>-EINTR</code> value if a signal other than the result of an up operation is received.
<code>int down_trylock(struct semaphore *sem)</code>	Attempts to acquire the given semaphore, and returns a nonzero value if semaphore is unavailable
<code>void up(struct semaphore *sem)</code>	Releases the given semaphore
Reader-Writer Semaphores	
<code>void init_rwsem(struct rw_semaphore, *rwsem)</code>	Initializes the dynamically created semaphore with a count of 1
<code>void down_read(struct rw_semaphore, *rwsem)</code>	Down operation for readers
<code>void up_read(struct rw_semaphore, *rwsem)</code>	Up operation for readers
<code>void down_write(struct rw_semaphore, *rwsem)</code>	Down operation for writers
<code>void up_write(struct rw_semaphore, *rwsem)</code>	Up operation for writers

Linux Memory Barrier Operations



Table 6.6 Linux Memory Barrier Operations

<code>rmb()</code>	Prevents loads from being reordered across the barrier
<code>wmb()</code>	Prevents stores from being reordered across the barrier
<code>mb()</code>	Prevents loads and stores from being reordered across the barrier
<code>Barrier()</code>	Prevents the compiler from reordering loads or stores across the barrier
<code>smp_rmb()</code>	On SMP, provides a <code>rmb()</code> and on UP provides a <code>barrier()</code>
<code>smp_wmb()</code>	On SMP, provides a <code>wmb()</code> and on UP provides a <code>barrier()</code>
<code>smp_mb()</code>	On SMP, provides a <code>mb()</code> and on UP provides a <code>barrier()</code>

SMP = symmetric multiprocessor


UP = uniprocessor

Solaris Thread Synchronization Primitives




- Mutual exclusion (mutex) locks
- Semaphores
- Multiple readers, single writer (readers/writer) locks
- Condition variables

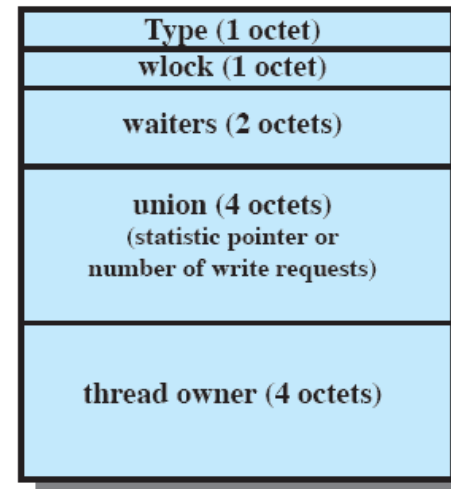
Solaris Synchronization Data Structures



(a) MUTEX lock



(b) Semaphore



(c) Reader/writer lock



(d) Condition variable

Figure 6.15 Solaris Synchronization Data Structures

Windows Synchronization Objects



Object Type	Definition	Set to Signaled State When	Effect on Waiting Threads
Notification Event	An announcement that a system event has occurred	Thread sets the event	All released
Synchronization event	An announcement that a system event has occurred.	Thread sets the event	One thread released
Mutex	A mechanism that provides mutual exclusion capabilities; equivalent to a binary semaphore	Owning thread or other thread releases the mutex	One thread released
Semaphore	A counter that regulates the number of threads that can use a resource	Semaphore count drops to zero	All released
Waitable timer	A counter that records the passage of time	Set time arrives or time interval expires	All released
File	An instance of an opened file or I/O device	I/O operation completes	All released
Process	A program invocation, including the address space and resources required to run the program	Last thread terminates	All released
Thread	An executable entity within a process	Thread terminates	All released

Note: Shaded rows correspond to objects that exist for the sole purpose of synchronization.

*Operating Systems:
Internals and Design Principles,
6/E*

William Stallings

Chapter 9
Uniprocessor Scheduling



Roadmap

- Types of Processor Scheduling
- • Scheduling Algorithms
- Traditional UNIX Scheduling



Scheduling

- An OS must allocate resources amongst competing processes.
- The resource provided by a processor is execution time
 - The resource is allocated by means of a schedule



Overall Aim of Scheduling

- The aim of processor scheduling is to assign processes to be executed by the processor over time,
 - in a way that meets system objectives, such as response time, throughput, and processor efficiency.



Scheduling Objectives



- The scheduling function should
 - Share time *fairly* among processes
 - Prevent starvation of a process
 - Use the processor efficiently
 - Have low overhead
 - Prioritise processes when necessary (e.g. real time deadlines)

Types of Scheduling




Table 9.1 Types of Scheduling

Long-term scheduling	The decision to add to the pool of processes to be executed
Medium-term scheduling	The decision to add to the number of processes that are partially or fully in main memory
Short-term scheduling	The decision as to which available process will be executed by the processor
I/O scheduling	The decision as to which process's pending I/O request shall be handled by an available I/O device

Two Suspend States

- Remember this diagram from Chapter 3



(b) With Two Suspend States

Scheduling and Process State Transitions


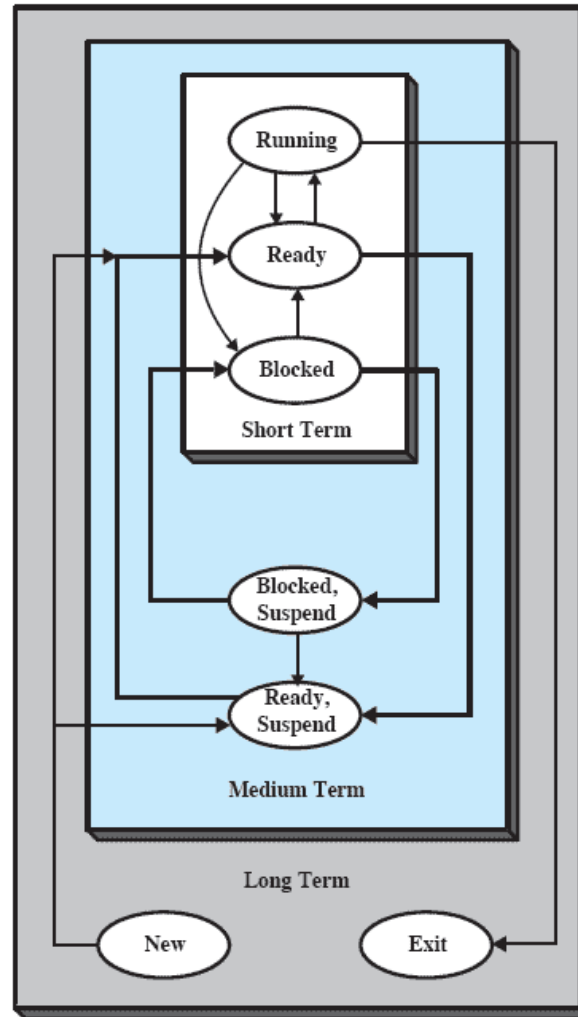


Figure 9.1 Scheduling and Process State Transitions

Nesting of Scheduling Functions



Queuing Diagram

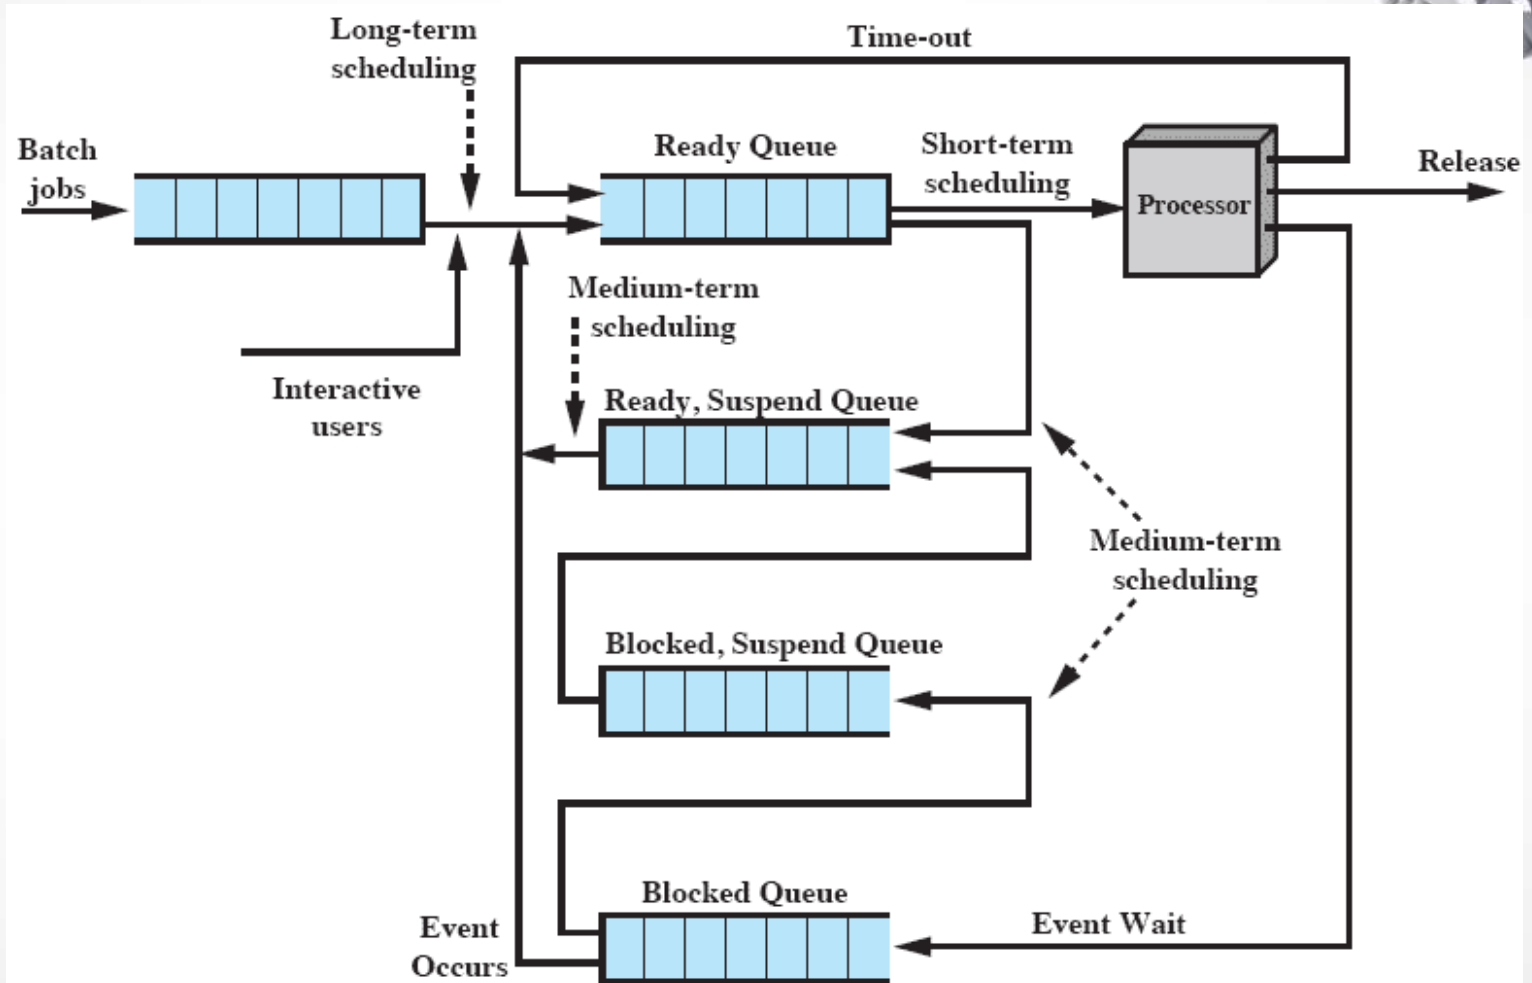


Figure 9.3 Queuing Diagram for Scheduling

Long-Term Scheduling



- Determines which programs are admitted to the system for processing
 - May be first-come-first-served
 - Or according to criteria such as priority, I/O requirements or expected execution time
- Controls the degree of multiprogramming
- More processes, smaller percentage of time each process is executed

Medium-Term Scheduling



- Part of the swapping function
- Swapping-in decisions are based on the need to manage the degree of multiprogramming

Short-Term Scheduling




- Known as the dispatcher
- Executes most frequently
- Invoked when an event occurs
 - Clock interrupts
 - I/O interrupts
 - Operating system calls
 - Signals

Roadmap

- Types of Processor Scheduling
- Scheduling Algorithms
- • Traditional UNIX Scheduling



Aim of Short Term Scheduling



- Main objective is to allocate processor time to optimize certain aspects of system behaviour.
- A set of criteria is needed to evaluate the scheduling policy.

Short-Term Scheduling

Criteria: User vs System



- We can differentiate between user and system criteria
- User-oriented
 - Response Time
 - Elapsed time between the submission of a request until there is output.
- System-oriented
 - Effective and efficient utilization of the processor

Short-Term Scheduling

Criteria: Performance



- We could differentiate between performance related criteria, and those unrelated to performance
- Performance-related
 - Quantitative, easily measured
 - E.g. response time and throughput
- Non-performance related
 - Qualitative
 - Hard to measure

Interdependent Scheduling Criteria



User Oriented, Performance Related

Turnaround time This is the interval of time between the submission of a process and its completion. Includes actual execution time plus time spent waiting for resources, including the processor. This is an appropriate measure for a batch job.

Response time For an interactive process, this is the time from the submission of a request until the response begins to be received. Often a process can begin producing some output to the user while continuing to process the request. Thus, this is a better measure than turnaround time from the user's point of view. The scheduling discipline should attempt to achieve low response time and to maximize the number of interactive users receiving acceptable response time.

Deadlines When process completion deadlines can be specified, the scheduling discipline should subordinate other goals to that of maximizing the percentage of deadlines met.

User Oriented, Other

Predictability A given job should run in about the same amount of time and at about the same cost regardless of the load on the system. A wide variation in response time or turnaround time is distracting to users. It may signal a wide swing in system workloads or the need for system tuning to cure instabilities.

Interdependent Scheduling Criteria cont.



System Oriented, Performance Related

Throughput The scheduling policy should attempt to maximize the number of processes completed per unit of time. This is a measure of how much work is being performed. This clearly depends on the average length of a process but is also influenced by the scheduling policy, which may affect utilization.

Processor utilization This is the percentage of time that the processor is busy. For an expensive shared system, this is a significant criterion. In single-user systems and in some other systems, such as real-time systems, this criterion is less important than some of the others.

System Oriented, Other


Fairness In the absence of guidance from the user or other system-supplied guidance, processes should be treated the same, and no process should suffer starvation.

Enforcing priorities When processes are assigned priorities, the scheduling policy should favor higher-priority processes.

Balancing resources The scheduling policy should keep the resources of the system busy. Processes that will underutilize stressed resources should be favored. This criterion also involves medium-term and long-term scheduling.

Priorities

- Scheduler will always choose a process of higher priority over one of lower priority
- Have multiple ready queues to represent each level of priority



Priority Queuing

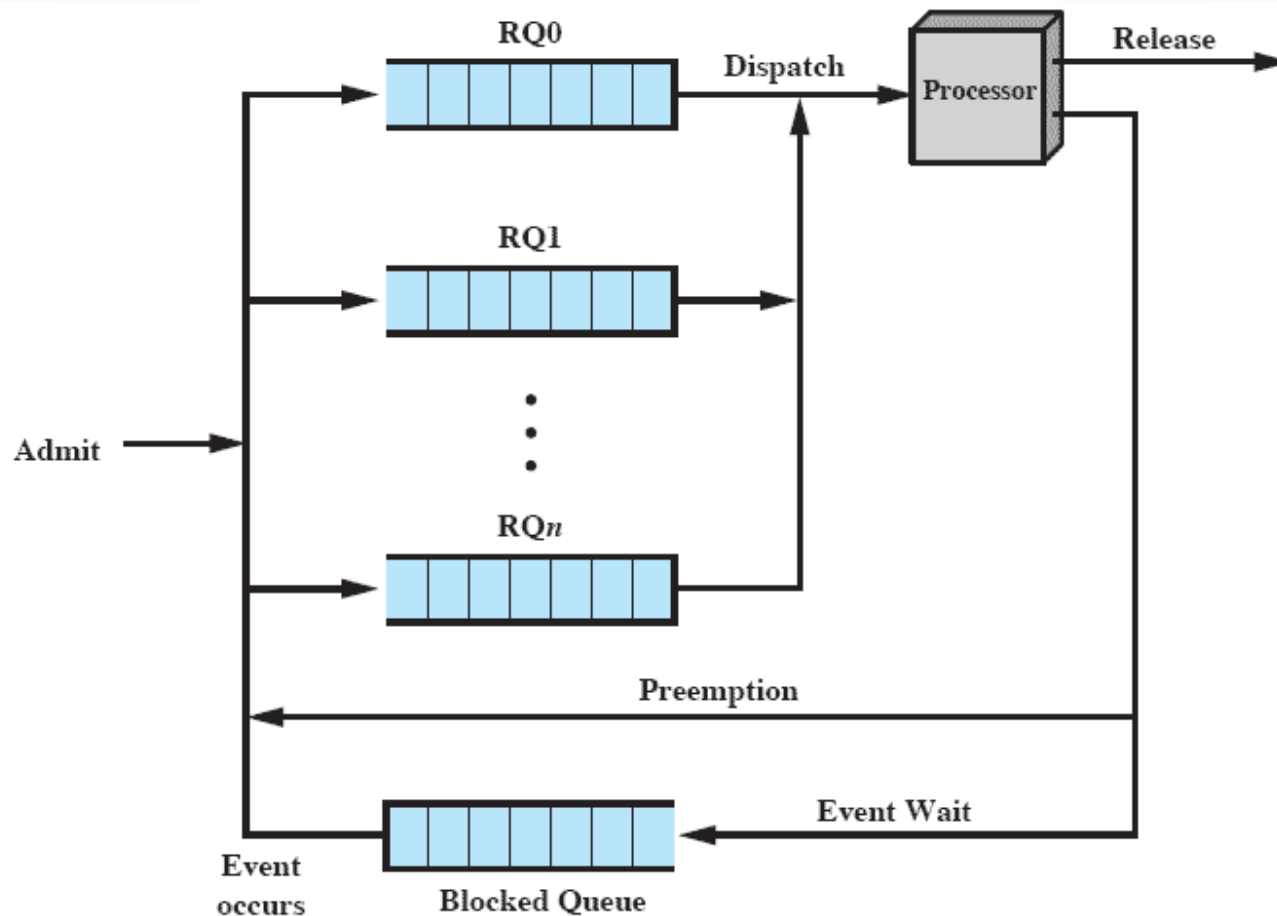


Figure 9.4 Priority Queuing

Starvation



- Problem:
 - Lower-priority may suffer starvation if there is a steady supply of high priority processes.
- Solution
 - Allow a process to change its priority based on its age or execution history

Alternative Scheduling Policies



Table 9.3 Characteristics of Various Scheduling Policies

	FCFS	Round robin	SPN	SRT	HRRN	Feedback
Selection function	$\max[w]$	constant	$\min[s]$	$\min[s - e]$	$\max\left(\frac{w + s}{s}\right)$	(see text)
Decision mode	Non-preemptive	Preemptive (at time quantum)	Non-preemptive	Preemptive (at arrival)	Non-preemptive	Preemptive (at time quantum)
Throughput	Not emphasized	May be low if quantum is too small	High	High	High	Not emphasized
Response time	May be high, especially if there is a large variance in process execution times	Provides good response time for short processes	Provides good response time for short processes	Provides good response time	Provides good response time	Not emphasized
Overhead	Minimum	Minimum	Can be high	Can be high	Can be high	Can be high
Effect on processes	Penalizes short processes; penalizes I/O bound processes	Fair treatment	Penalizes long processes	Penalizes long processes	Good balance	May favor I/O bound processes
Starvation	No	No	Possible	Possible	No	Possible


Selection Function



- Determines which process is selected for execution
- If based on execution characteristics then important quantities are:
 - w = time spent in system so far, waiting
 - e = time spent in execution so far
 - s = total service time required by the process, including e ;

Decision Mode

- Specifies the instants in time at which the selection function is exercised.
- Two categories:
 - Nonpreemptive
 - Preemptive



Nonpreemptive vs Preemptive



- Non-preemptive
 - Once a process is in the running state, it will continue until it terminates or blocks itself for I/O
- Preemptive
 - Currently running process may be interrupted and moved to ready state by the OS
 - Preemption may occur when new process arrives, on an interrupt, or periodically.

Process Scheduling Example



- Example set of processes, consider each a batch job

Table 9.4 Process Scheduling Example


Process	Arrival Time	Service Time
A	0	3
B	2	6
C	4	4
D	6	5
E	8	2

- Service time represents total execution time

First-Come-First-Served

- Each process joins the Ready queue
- When the current process ceases to execute, the longest process in the Ready queue is selected

First-Come-First Served (FCFS)



First-Come- First-Served

- A short process may have to wait a very long time before it can execute
- Favors CPU-bound processes
 - I/O processes have to wait until CPU-bound process completes




Round Robin




- Uses preemption based on a clock
 - also known as time slicing, because each process is given a slice of time before being preempted.

Round-Robin
(RR), $q = 1$

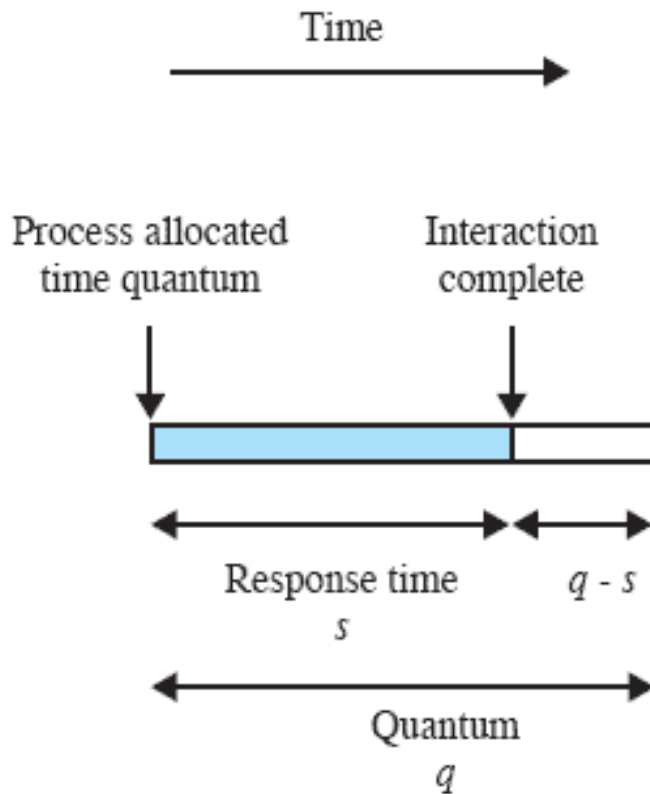


Round Robin

- Clock interrupt is generated at periodic intervals
- When an interrupt occurs, the currently running process is placed in the ready queue
 - Next ready job is selected

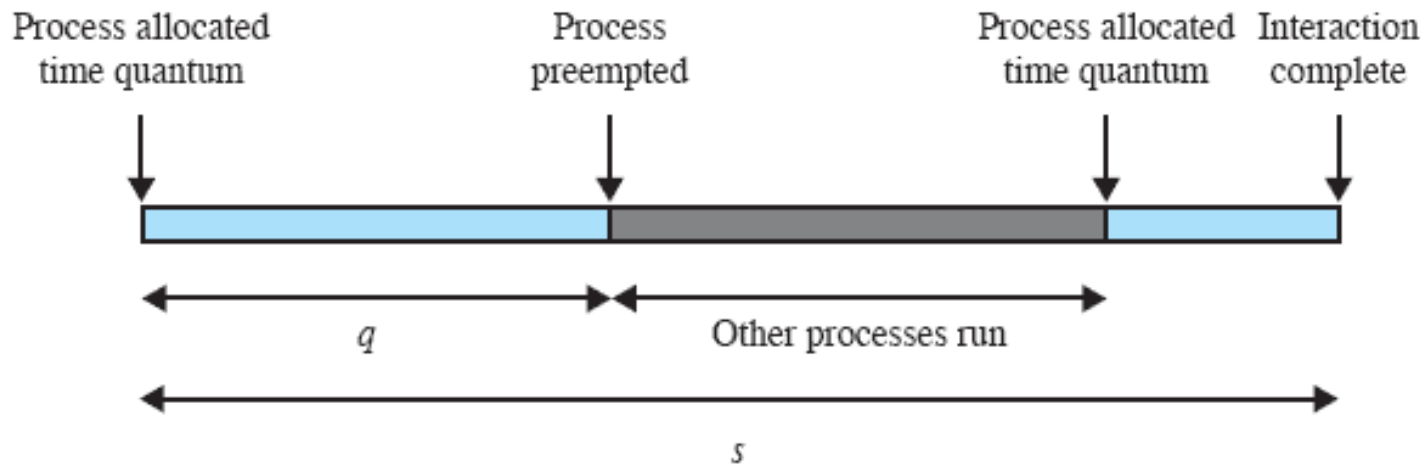


Effect of Size of Preemption Time Quantum



(a) Time quantum greater than typical interaction

Effect of Size of Preemption Time Quantum



(b) Time quantum less than typical interaction

Figure 9.6 Effect of Size of Preemption Time Quantum

'Virtual Round Robin'





Figure 9.7 Queuing Diagram for Virtual Round-Robin Scheduler

Shortest Process Next



- Nonpreemptive policy
- Process with shortest expected processing time is selected next
- Short process jumps ahead of longer processes

Shortest Process
Next (SPN)



Shortest Process Next



- Predictability of longer processes is reduced
- If estimated time for process not correct, the operating system may abort it
- Possibility of starvation for longer processes

Calculating Program 'Burst'



- Where:
 - T_i = processor execution time for the i th instance of this process
 - S_i = predicted value for the i th instance
 - S_1 = predicted value for first instance; not calculated

$$S_{n+1} = \frac{1}{n} \sum_{i=1}^n T_i$$

Exponential Averaging



- A common technique for predicting a future value on the basis of a time series of past values is exponential averaging

$$S_{n+1} = \alpha T_n + (1 - \alpha)S_n$$

Exponential Smoothing Coefficients






Figure 9.8 Exponential Smoothing Coefficients

Use Of Exponential Averaging



(a) Increasing function

Use Of Exponential Averaging




(b) Decreasing function

Shortest Remaining Time



- Preemptive version of shortest process next policy
- Must estimate processing time and choose the shortest

Shortest Remaining Time (SRT)




Highest Response Ratio Next



- Choose next process with the greatest ratio

$$\text{Ratio} = \frac{\text{time spent waiting} + \text{expected service time}}{\text{expected service time}}$$

Highest Response Ratio Next (HRRN)



Feedback Scheduling



- Penalize jobs that have been running longer
- Don't know remaining time process needs to execute





Figure 9.10 Feedback Scheduling

Feedback Performance

- Variations exist, simple version pre-empts periodically, similar to round robin
 - But can lead to starvation



Performance Comparison



- Any scheduling discipline that chooses the next item to be served independent of service time obeys the relationship:

$$\frac{T_r}{T_s} = \frac{1}{1 - \rho}$$

where

T_r = turnaround time or residence time; total time in system, waiting plus execution

T_s = average service time; average time spent in Running state

ρ = processor utilization

Formulas



Table 9.6 Formulas for Single-Server Queues with Two Priority Categories

- Assumptions:
1. Poisson arrival rate.
 2. Priority 1 items are serviced before priority 2 items.
 3. First-come-first-served dispatching for items of equal priority.
 4. No item is interrupted while being served.
 5. No items leave the queue (lost calls delayed).

(a) General formulas

$$\begin{aligned}\lambda &= \lambda_1 + \lambda_2 \\ \rho_1 &= \lambda_1 T_{s1}; \quad \rho_2 = \lambda_2 T_{s2} \\ \rho &= \rho_1 + \rho_2 \\ T_s &= \frac{\lambda_1}{\lambda} T_{s1} + \frac{\lambda_2}{\lambda} T_{s2} \\ T_r &= \frac{\lambda_1}{\lambda} T_{r1} + \frac{\lambda_2}{\lambda} T_{r2}\end{aligned}$$

b) No interrupts; exponential service times

$$\begin{aligned}T_{r1} &= T_{s1} + \frac{\rho_1 T_{s1} + \rho_2 T_{s2}}{1 - \rho_1} \\ T_{r2} &= T_{s2} + \frac{T_{r1} - T_{s1}}{1 - \rho}\end{aligned}$$

(c) Preemptive-resume queuing discipline; exponential service times

$$\begin{aligned}T_{r1} &= T_{s1} + \frac{\rho_1 T_{s1}}{1 - \rho_1} \\ T_{r2} &= T_{s2} + \frac{1}{1 - \rho_1} \left(\rho_1 T_{s2} + \frac{\rho T_s}{1 - \rho} \right)\end{aligned}$$

Overall Normalized Response Time




Figure 9.11 Overall Normalized Response Time

Normalized Response Time for Shorter Process




Figure 9.12 Normalized Response Time for Shorter Processes

Normalized Response Time for Longer Processes

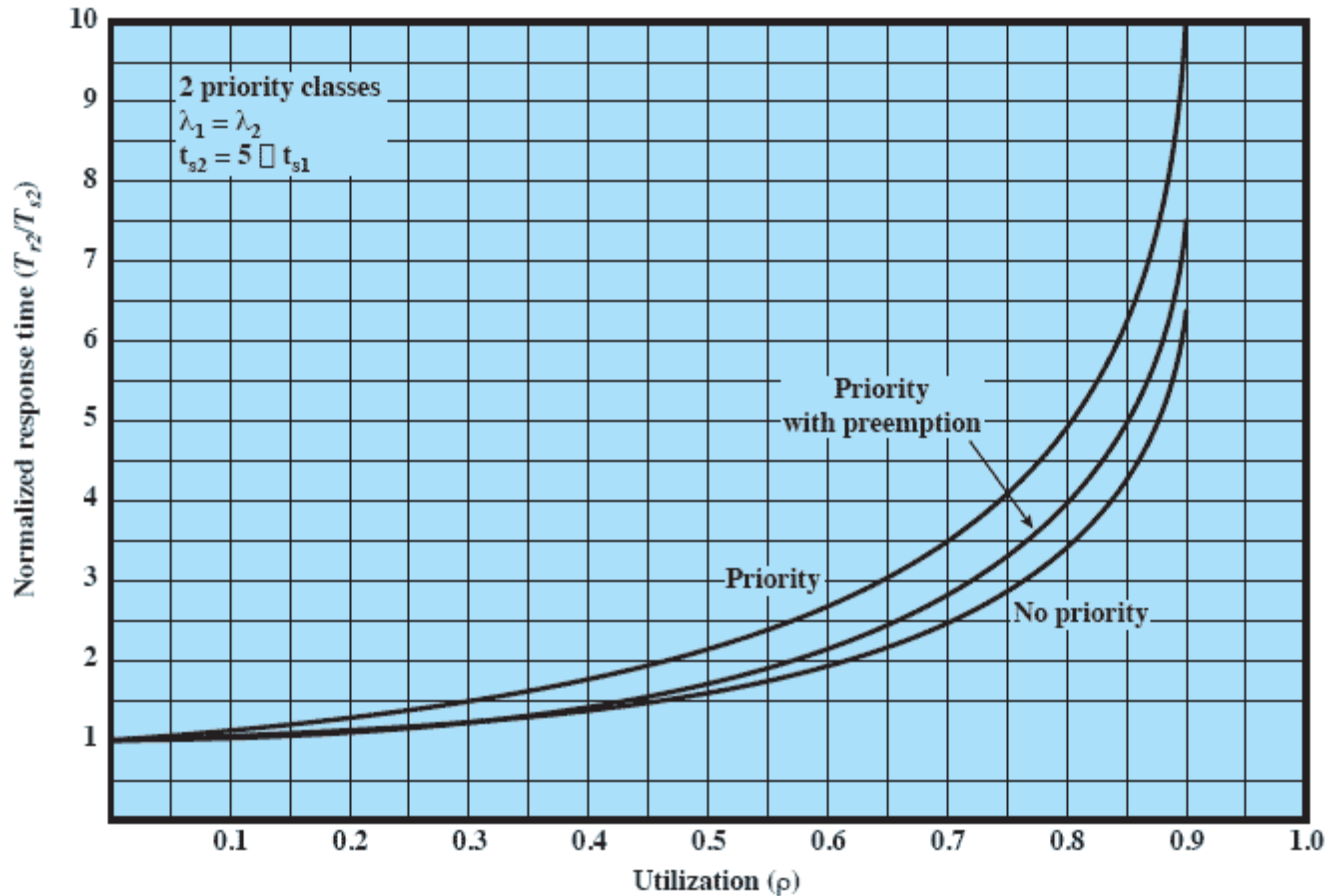


Figure 9.13 Normalized Response Time for Longer Processes

Normalized Turnaround Time




Figure 9.14 Simulation Results for Normalized Turnaround Time





Figure 9.15 Simulation Results for Waiting Time

Fair-Share Scheduling



- User's application runs as a collection of processes (threads)
- User is concerned about the performance of the application
- Need to make scheduling decisions based on process sets

Fair-Share Scheduler



Time	Process A			Process B			Process C			
	Priority	Process CPU count	Group CPU count	Priority	Process CPU count	Group CPU count	Priority	Process CPU count	Group CPU count	
0	60	0	0	60	0	0	60	0	0	
		1	1							
		2	2							
		•	•							
		60	60							
1	90	30	30	60	0	0	60	0	0	
					1	1			1	
					2	2			2	
					•	•			•	
					60	60			60	
2	74	15	15	90	30	30	75	0	30	
		16	16							
		17	17							
		•	•							
		75	75							
3	96	37	37	74	15	15	67	0	15	
						16		1	16	
						17		2	17	
						•		•	•	
						75		60	75	
4	78	18	18	81	7	37	93	30	37	
		19	19							
		20	20							
		•	•							
		78	78							
5	98	39	39	70	3	18	76	15	18	

Colored rectangle represents executing process

Figure 9.16 Example of Fair Share Scheduler — Three Processes, Two Groups