

Concurrency: mutual exclusion and synchronization

Slides are mainly taken from «*Operating Systems: Internals and Design Principles*», 8/E William Stallings (Chapter 5).

Sistemi di Calcolo (II semestre) – Roberto Baldoni

Multiple Processes

- Operating System design is concerned with the management of processes and threads:
 - Multiprogramming
 - Multiprocessing
 - Distributed Processing



Concurrency

Arises in Three Different Contexts:

Multiple Applications

invented to allow processing time to be shared among active applications

Structured Applications

extension of modular design and structured programming

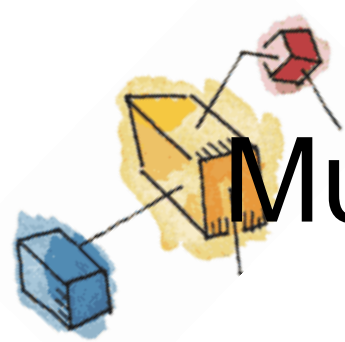
Operating System Structure

OS themselves implemented as a set of processes or threads

atomic operation	A function or action implemented as a sequence of one or more instructions that appears to be indivisible; that is, no other process can see an intermediate state or interrupt the operation. The sequence of instruction is guaranteed to execute as a group, or not execute at all, having no visible effect on system state. Atomicity guarantees isolation from concurrent processes.
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.

Table 5.1

**Some Key
Terms
Related
to
Concurrency**



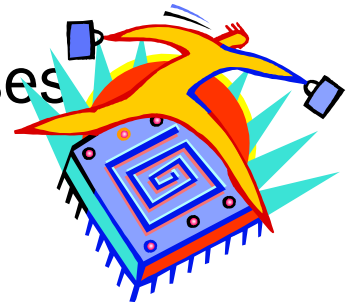
Multiprogramming Concerns

- **Output of process must be independent of the speed of execution of other concurrent processes**



Principles of Concurrency

- Interleaving and overlapping
 - can be viewed as examples of concurrent processing
 - both present the same problems
- Uniprocessor – the relative speed of execution of processes cannot be predicted
 - depends on activities of other processes
 - the way the OS handles interrupts
 - scheduling policies of the OS



Difficulties of Concurrency

- Sharing of global resources
- Difficult for the OS to manage the allocation of resources optimally
- Difficult to locate programming errors as results are not deterministic and reproducible



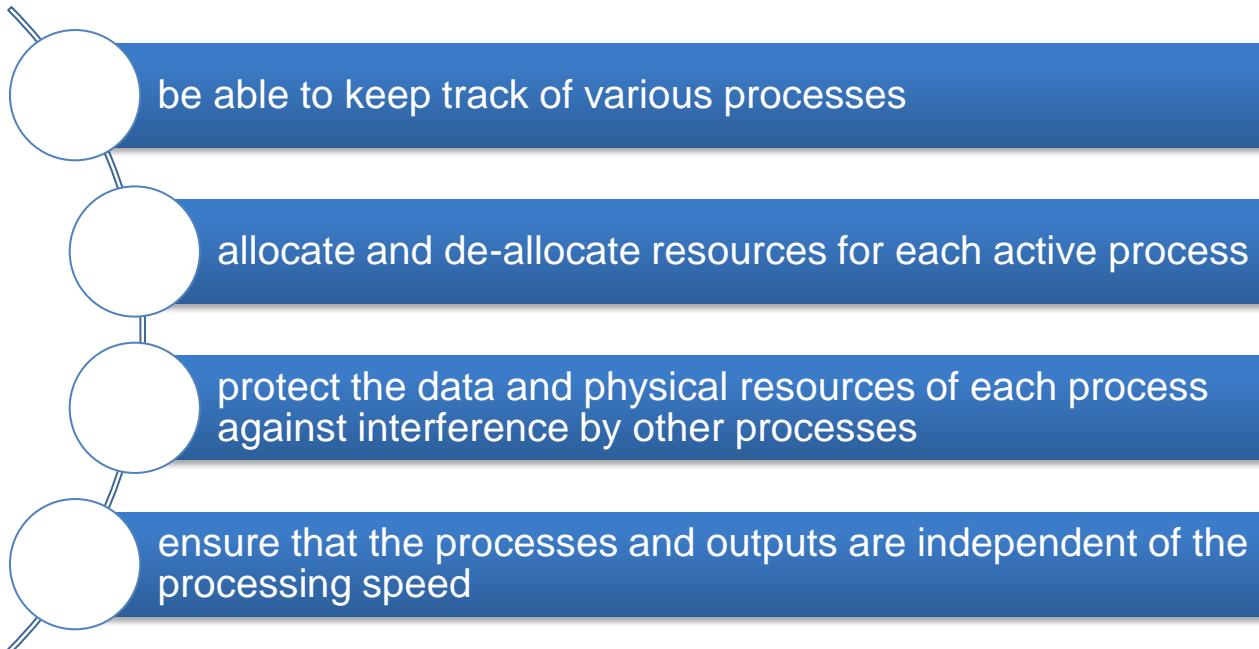
Race Condition

- Occurs when multiple processes or threads read and write data items
- The final result depends on the order of execution
 - the “loser” of the race is the process that updates last and will determine the final value of the variable



Operating System Concerns

- Design and management issues raised by the existence of concurrency:
 - The OS must:



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

Table 5.2

Process Interaction

Resource Competition

- Concurrent processes come into conflict when they are competing for use of the same resource
 - for example: I/O devices, memory, processor time, clock

In the case of competing processes three control problems must be faced:



- the need for mutual exclusion
- deadlock
- starvation

Figure 5.1

Illustration of Mutual Exclusion

PROCESS 1 */	/* PROCESS 2 */	...	/* PROCESS n */
<pre>void P1 { while (true) { /* preceding code */; entercritical (Ra); /* critical section */; exitcritical (Ra); /* following code */; } }</pre>	<pre>void P2 { while (true) { /* preceding code */; entercritical (Ra); /* critical section */; exitcritical (Ra); /* following code */; } }</pre>		<pre>void Pn { while (true) { /* preceding code */; entercritical (Ra); /* critical section */; exitcritical (Ra); /* following code */; } }</pre>

Requirements for Mutual Exclusion

- Must be enforced
- A process that halts must do so without interfering with other processes
- No deadlock or starvation
- A process must not be denied 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



Mutual Exclusion: Hardware Support

■ Interrupt Disabling

- uniprocessor system
- disabling interrupts guarantees mutual exclusion

■ Disadvantages:

- the efficiency of execution could be noticeably degraded
- this approach will not work in a multiprocessor architecture

Mutual Exclusion: Hardware Support

- Compare&Swap Instruction

» also

instr

» a co

mer

» if th

occur

» carr

```
int compare_and_swap(int* reg, int oldval, int newval)
{
    ATOMIC();
    int old_reg_val = *reg;
    if (old_reg_val == oldval)
        *reg = newval;
    END_ATOMIC();
    return old_reg_val;
}
```

Figure 5.2

Hardware Support for Mutual Exclusion

```
/* 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;
    par begin ( P(1), P(2), . . . , P(n));
}
}
```

```
int compare_and_swap(int* reg, int oldval, int newval)
{
    ATOMIC();
    int old_reg_val = *reg;
    if (old_reg_val == oldval)
        *reg = newval;
    END_ATOMIC();
    return old_reg_val;
}
```

(a) Compare and swap instruction

(b) Exchange instruction

Figure 5.2

Hardware Support for Mutual Exclusion

```
/* program mutualexclusion */
const int n = /* number of processes */;

• Exchange instruction

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

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

(a) Compare and swap instruction

(b) Exchange instruction

Special Machine Instruction: Advantages

- ↑ Applicable to any number of processes on either a single processor or multiple processors sharing main memory
- ↑ Simple and easy to verify
 - It can be used to support multiple critical sections; each critical section can be defined by its own variable



Special Machine Instruction: Disadvantages

- ↓ Busy-waiting is employed, thus while a process is waiting for access to a critical section it continues to consume processor time
- ↓ Starvation is possible when a process leaves a critical section and more than one process is waiting
- ↓ Deadlock is possible



Semaphor e	An integer value used for signaling among processes. Only three operations may be performed on a semaphore, all of which are atomic: initialize, decrement, and increment. The decrement operation may result in the blocking of a process, and the increment operation may result in the unblocking of a process. Also known as a counting semaphore or a general semaphore
Binary Semaphore	A semaphore that takes on only the values 0 and 1.
Mutex	Similar to a binary semaphore. A key difference between the two is that the process that locks the mutex (sets the value to zero) must be the one to unlock it (sets the value to 1).
Condition Variable	A data type that is used to block a process or thread until a particular condition is true.
Monitor	A programming language construct that encapsulates variables, access procedures and initialization code within an abstract data type. The monitor's variable may only be accessed via its access procedures and only one process may be actively accessing the monitor at any one time. The access procedures are <i>critical sections</i> . A monitor may have a queue of processes that are waiting to access it.
Event Flags	A memory word used as a synchronization mechanism. Application code may associate a different event with each bit in a flag. A thread can wait for either a single event or a combination of events by checking one or multiple bits in the corresponding flag. The thread is blocked until all of the required bits are set (AND) or until at least one of the bits is set (OR).
Mailboxes/Messages	A means for two processes to exchange information and that may be used for synchronization.
Spinlocks	Mutual exclusion mechanism in which a process executes in an infinite loop waiting for the value of a lock variable to indicate availability.

Table 5.3

Common

Concurrenc

y

Mechanisms

Semaphore

A variable that has an integer value upon which only three operations are defined:

- There is no way to inspect or manipulate semaphores other than these three operations

- 1) May be initialized to a nonnegative integer value
- 2) The semWait operation decrements the value
- 3) The semSignal operation increments the value

Consequences

There is no way to know before a process decrements a semaphore whether it will block or not

There is no way to know which process will continue immediately on a uniprocessor system when two processes are running concurrently

You don't know whether another process is waiting so the number of unblocked processes may be zero or one

**Figure
5.3**

**A
Definition
of
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 */;  
    }  
}
```

```

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 Semaphores

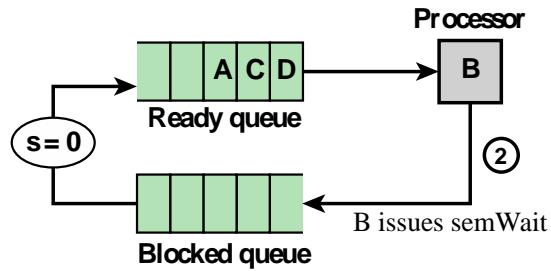
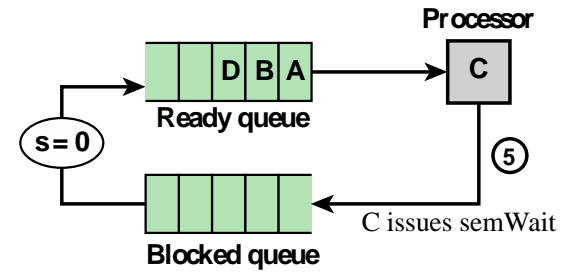
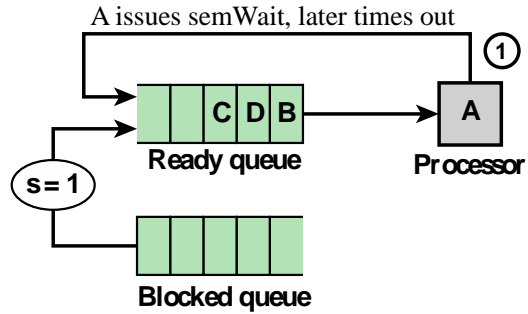
☹ A queue is used to hold processes waiting on the semaphore

Strong Semaphores

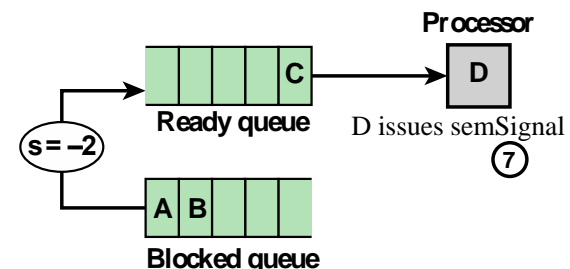
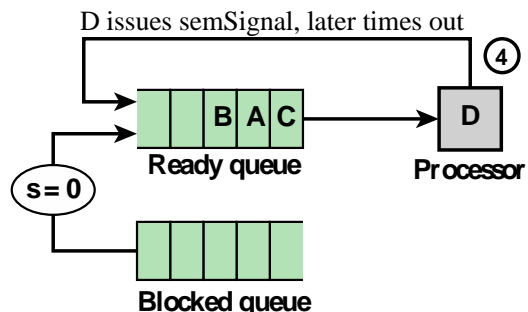
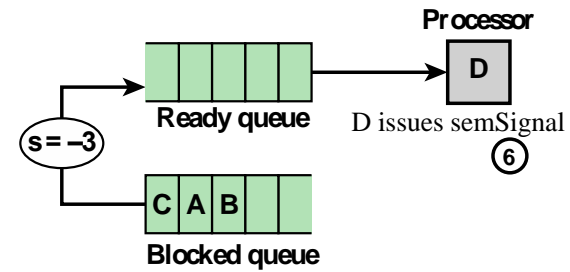
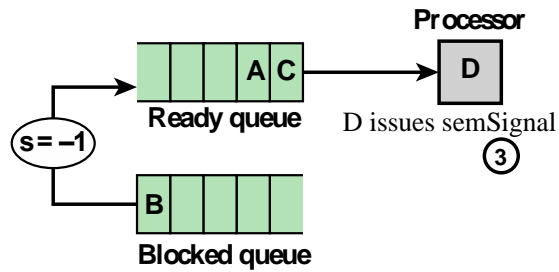
- the process that has been blocked the longest is released from the queue first (FIFO)

Weak Semaphores

- the order in which processes are removed from the queue is not specified



•
•
•



```

/* 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

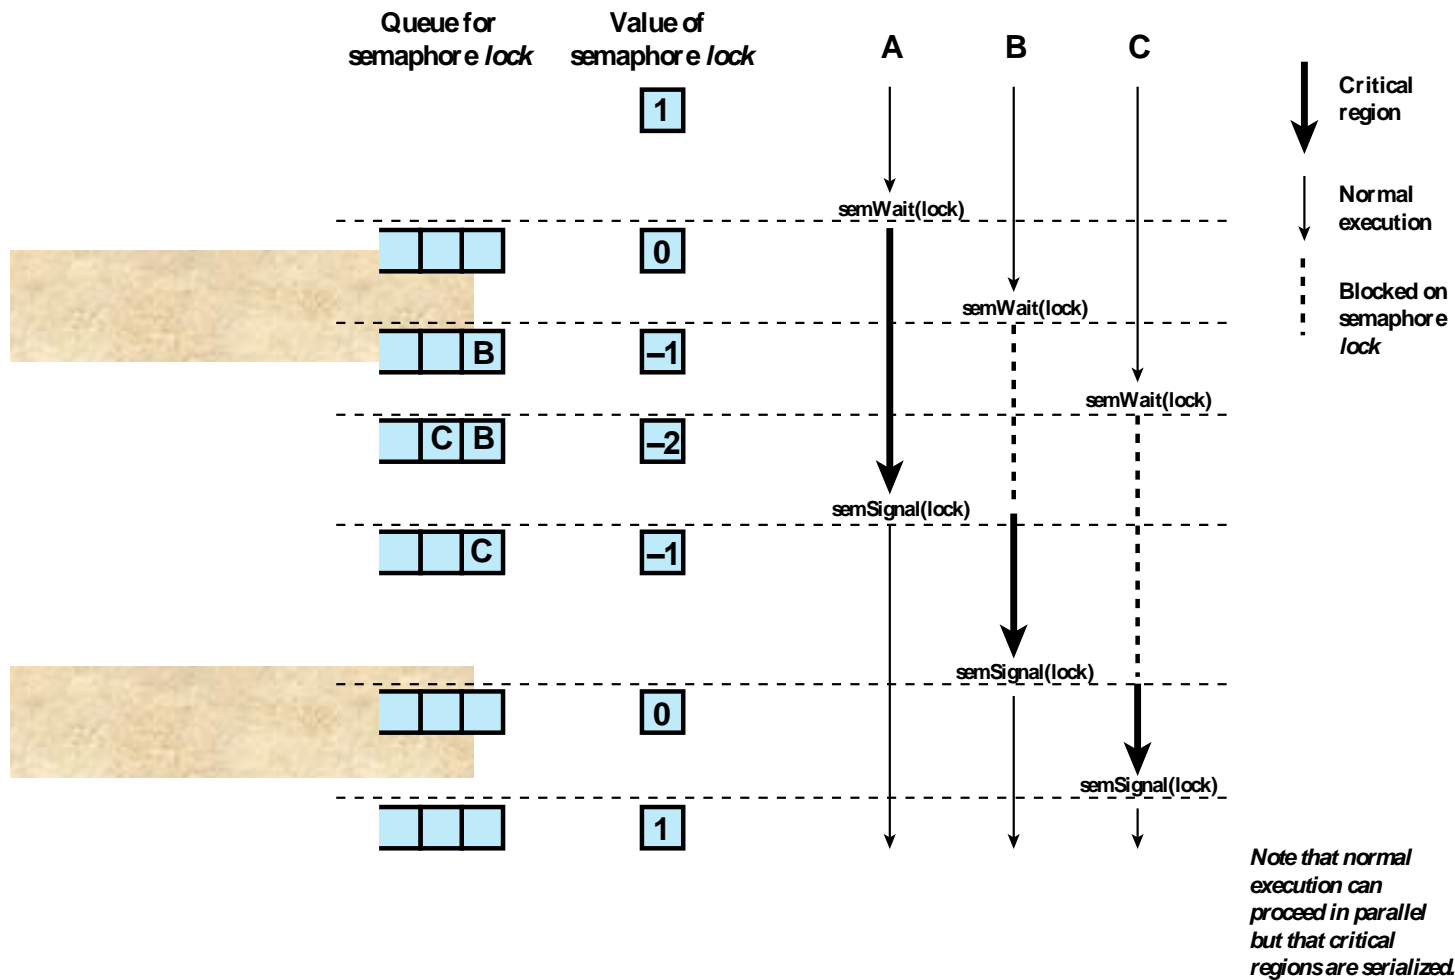


Figure 5.7 Processes Accessing Shared Data Protected by a Semaphore

Producer/Consumer Problem

General Statement:

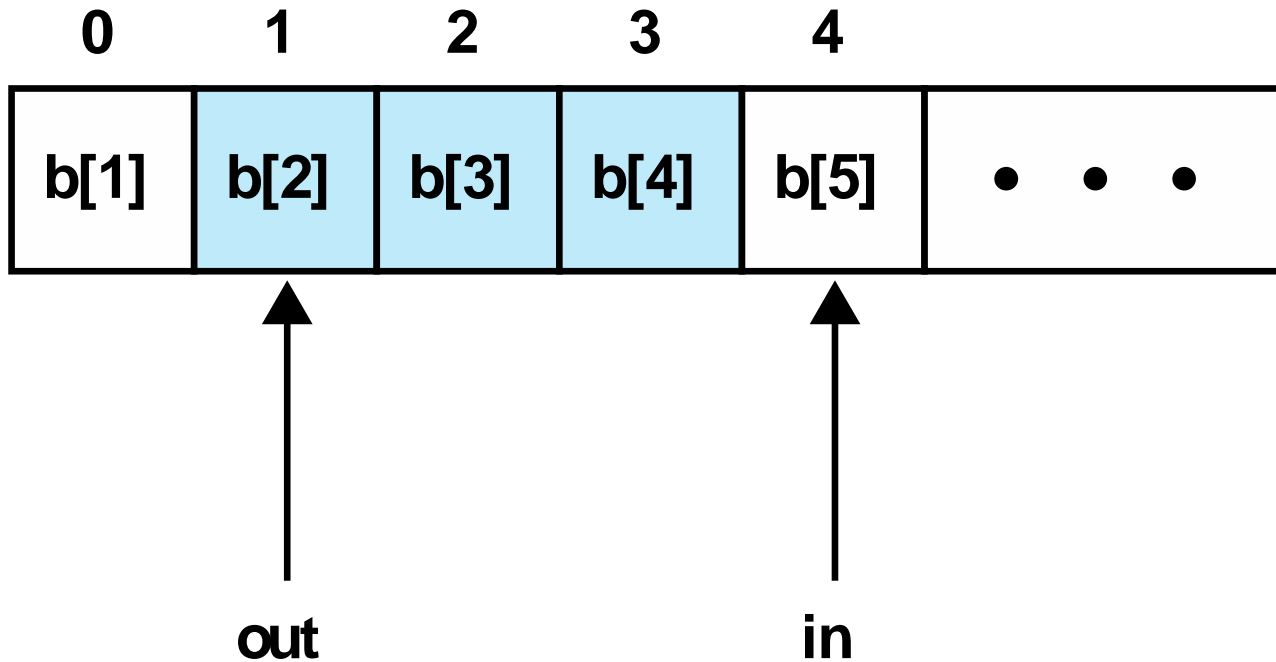
one or more producers are generating data and placing these in a buffer

a single consumer is taking items out of the buffer one at a 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 an empty buffer



Note: shaded area indicates portion of buffer that is occupied

Figure 5.8 Infinite Buffer for the Producer/Consumer Problem

```

/* pr ogr am producerconsumer */
i nt n;
binary semaphore s = 1, delay = 0;
voi d producer()
{
    whi l e (true) {
        produce();
        semWaitB(s);
        append();
        n++;
        i f (n==1) semSignalB(delay);
        semSignalB(s);
    }
}
voi d consumer()
{
    semWaitB(delay);
    whi l e (true) {
        semWaitB(s);
        take();
        n--;
        semSignalB(s);
        consume();
        i f (n==0) semWaitB(delay);
    }
}
voi d main()
{
    n = 0;
    par begi n (producer, consumer);
}

```

Figure 5.9

**An Incorrect
Solution
to the
Infinite-Buffer
Producer/Consu
mer
Problem Using
Binary
Semaphores**

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		semSignalB(s)	1	-1	0

Note: White areas represent the critical section controlled by semaphore s.


```

/* 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);
}

```

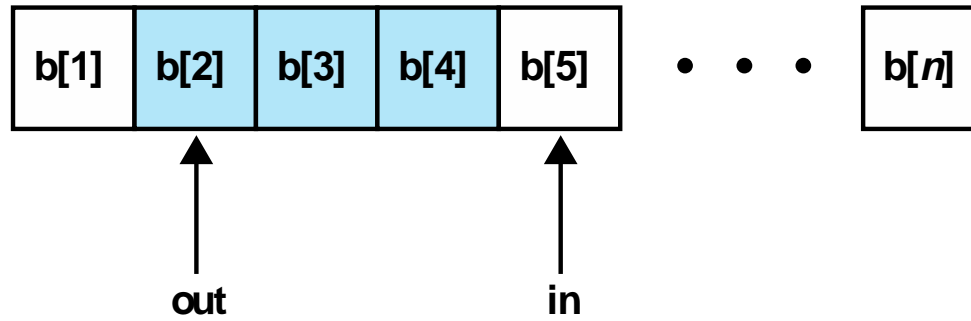
Figure 5.10

**A Correct
Solution to the
Infinite-Buffer
Producer/Consumer
Problem
Using Binary
Semaphores**

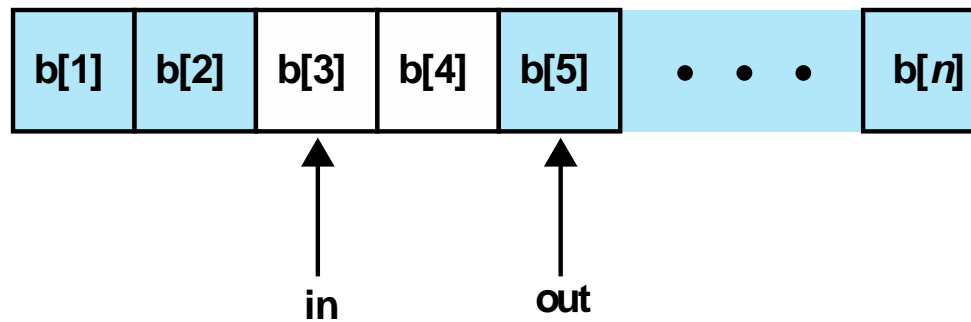
**Figure
5.11**

**A Solution
to the
Infinite-
Buffer
Producer/C
onsumer
Problem
Using
Semaphore
s**

```
/* pr ogr am  producerconsumer */
semaphore n = 0, s = 1;
voi d producer()
{
    whi l e (true) {
        produce();
        semWait(s);
        append();
        semSignal(s);
        semSignal(n);
    }
}
voi d consumer()
{
    whi l e (true) {
        semWait(n);
        semWait(s);
        take();
        semSignal(s);
        consume();
    }
}
voi d main()
{
    par begi n (producer, consumer);
}
```



(a)



(b)

Figure 5.12 Finite Circular Buffer for the Producer/Consumer Problem

Figure 5.13

A Solution to the Bounded- Buffer Producer/Consumer Problem Using 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);
}
```

Implementation of Semaphores

- Imperative that the `semWait` and `semSignal` operations be implemented as atomic primitives
- Can be implemented in hardware or firmware
- Software schemes such as Dekker's or Peterson's algorithms can be used
- Use one of the hardware-supported schemes for mutual exclusion

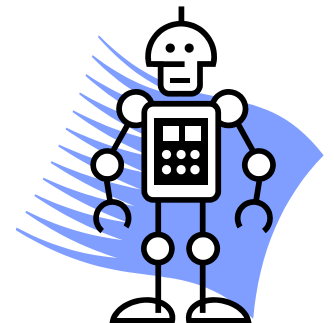


Figure 5.14

Two Possible Implementations of Semaphores

```
semWait(s)
{
    while (compare_and_swap(s.flag, 0 , 1) == 1)
        /* do nothing */;
    s.count--;
    if (s.count < 0) {
        /* place this process in s.queue*/;
        /* block this process (must also set s.flag to 0)
    */;
    }
    s.flag = 0;
}

semSignal(s)
{
    while (compare_and_swap(s.flag, 0 , 1) == 1)
        /* do nothing */;
    s.count++;
    if (s.count <= 0) {
        /* remove a process P from s.queue */;
        /* place process P on ready list */;
    }
    s.flag = 0;
}
```

```
semWait(s)
{
    inhibit interrupts;
    s.count--;
    if (s.count < 0) {
        /* place this process in s.queue*/;
        /* block this process and allow interrupts*/;
    }
    else
        allow interrupts;
}

semSignal(s)
{
    inhibit interrupts;
    s.count++;
    if (s.count <= 0) {
        /* remove a process P from s.queue*/;
        /* place process P on ready list*/;
    }
    allow interrupts;
}
```

(a) Compare and Swap Instruction

(b) Interrupts


Monitors



- Programming language construct that provides equivalent functionality to that of semaphores and is easier to control
- Implemented in a number of programming languages
 - including Concurrent Pascal, Pascal-Plus, Modula-2, Modula-3, and Java
- Has also been implemented as a program library
- Software module consisting of one or more procedures, an initialization sequence, and local data

Monitor Characteristics

Local data variables are accessible only by the monitor's procedures and not by any external procedure



Process enters monitor by invoking one of its procedures



Only one process may be executing in the monitor at a time

Synchronization

- Achieved by the use of **condition variables** that are contained within the monitor and accessible only within the monitor
 - Condition variables are operated on by two 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



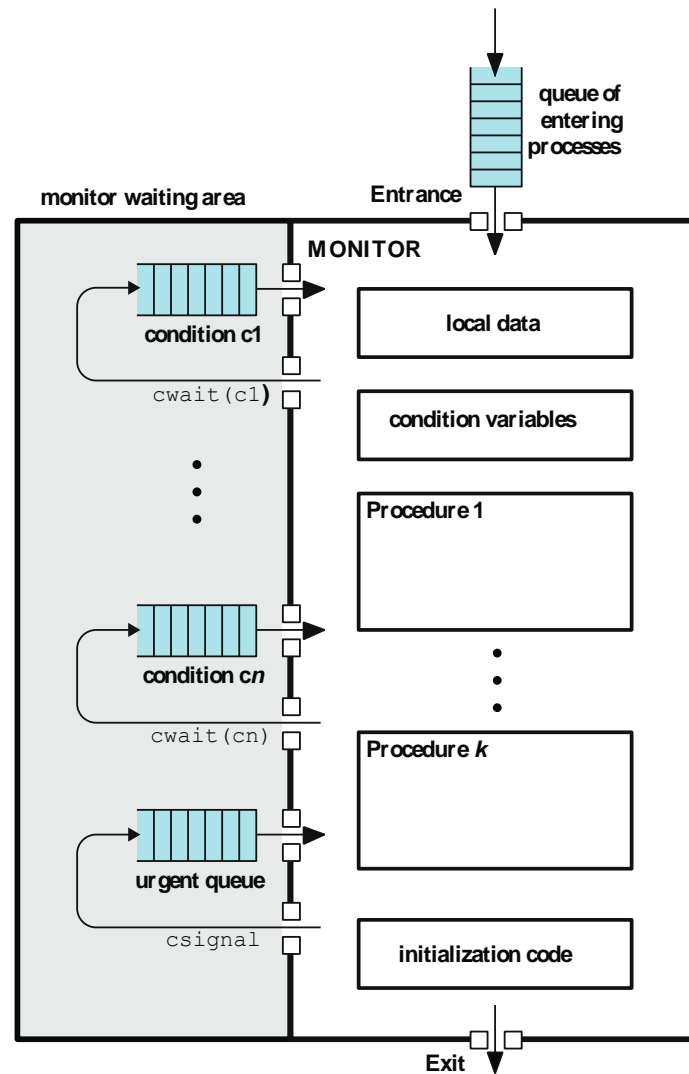


Figure 5.15 Structure of a Monitor

```

/* 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--;
    /* one fewer item in buffer */
    csignal(notfull);                         /* resume any waiting producer */
}

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

```

```

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);
}

```

Figure 5.16

**A Solution to
the Bounded-
Buffer
Producer/Cons
umer Problem
Using a Monitor**

Message Passing

- When processes interact with one another two fundamental requirements must be satisfied:

synchronization

- to enforce mutual exclusion

communication

- to exchange information

- Message Passing is one approach to providing both of these functions
 - works with distributed systems *and* shared memory multiprocessor and uniprocessor systems

Message Passing



- The actual function is normally provided in the form of a pair of primitives:
 - send (destination, message)
 - receive (source, message)
- » A process sends information in the form of a *message* to another process designated by a *destination*
- » A process receives information by executing the `receive` primitive, indicating the *source* and the *message*

Synchronization	Format
Send	Content
blocking	Length
nonblocking	fixed
Receive	variable
blocking	Queueing Discipline
nonblocking	FIFO
test for arrival	Priority
Addressing	
Direct	
send	
receive	
explicit	
implicit	
Indirect	
static	
dynamic	
ownership	

Table 5.5
Design Characteristics of Message Systems for
Interprocess Communication and Synchronization

Synchronization

Communication of a message between two processes implies synchronization between the two

When a receive primitive is executed in a process there are two possibilities:

if there is no waiting message the process is blocked until a message arrives or the process continues to execute, abandoning the attempt to receive

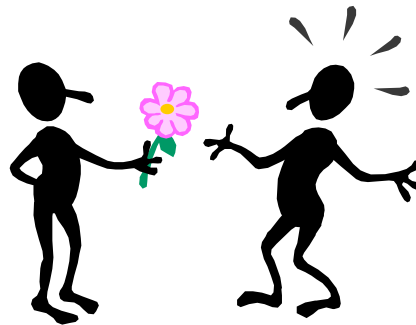
the receiver cannot receive a message until it has been sent by another process

if a message has previously been sent the message is received and execution continues



Blocking Send, Blocking Receive

- Both sender and receiver are blocked until the message is delivered
- Sometimes referred to as a *rendezvous*
- Allows for tight synchronization between processes



Nonblocking Send

Nonblocking send, blocking receive

- sender continues on but receiver is blocked until the requested message arrives
- most useful combination
- sends one or more messages to a variety of destinations as quickly as possible
- example -- a service process that exists to provide a service or resource to other processes

Nonblocking send, nonblocking receive

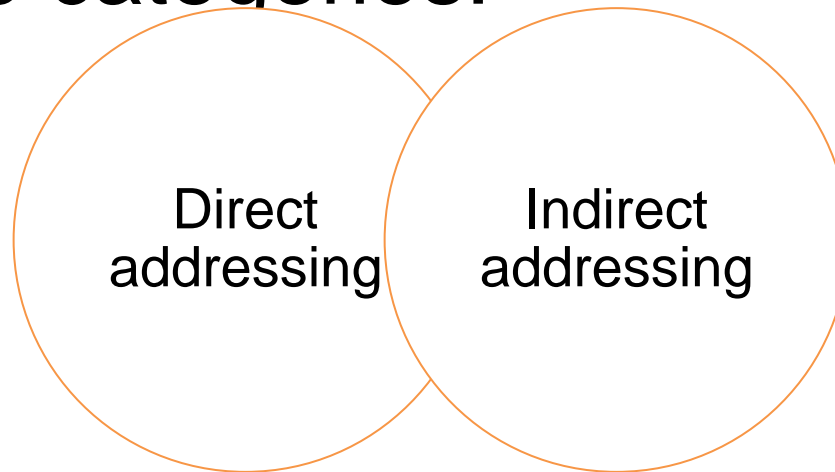
- neither party is required to wait





Addressing

- ★ Schemes for specifying processes in `send` and `receive` primitives fall into two categories:





Direct Addressing

- Send primitive includes a specific identifier of the destination process
- Receive primitive can be handled in one of two ways:
 - require that the process explicitly designate a sending process
 - effective for cooperating concurrent processes
 - implicit addressing
 - source parameter of the receive primitive possesses a value returned when the receive operation has been performed



Indirect Addressing

Messages are sent to a shared data structure consisting of queues that can temporarily hold messages



Queues are referred to as *mailboxes*



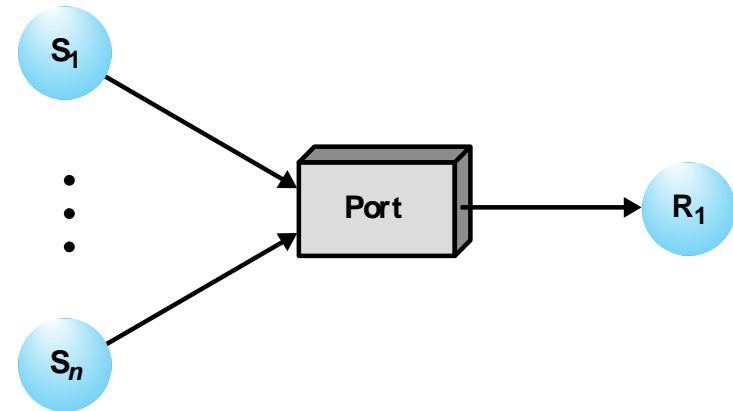
Allows for greater flexibility in the use of messages



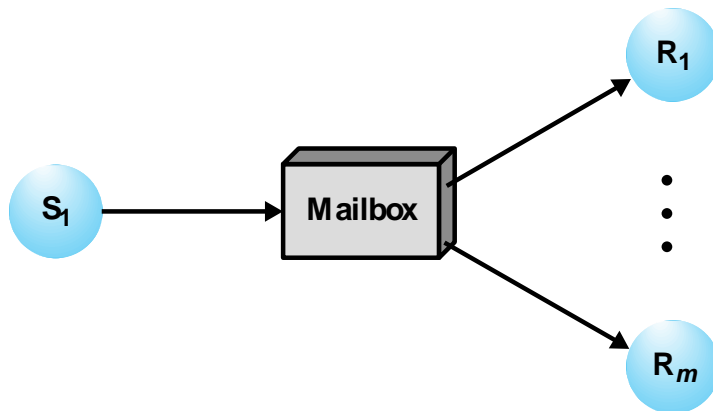
One process sends a message to the mailbox and the other process picks up the message from the mailbox



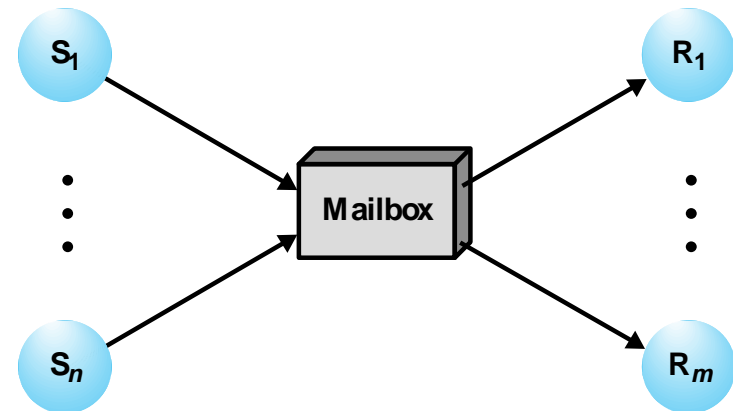
(a) One to one



(b) Many to one



(c) One to many



(d) Many to many

Figure 5.18 Indirect Process Communication

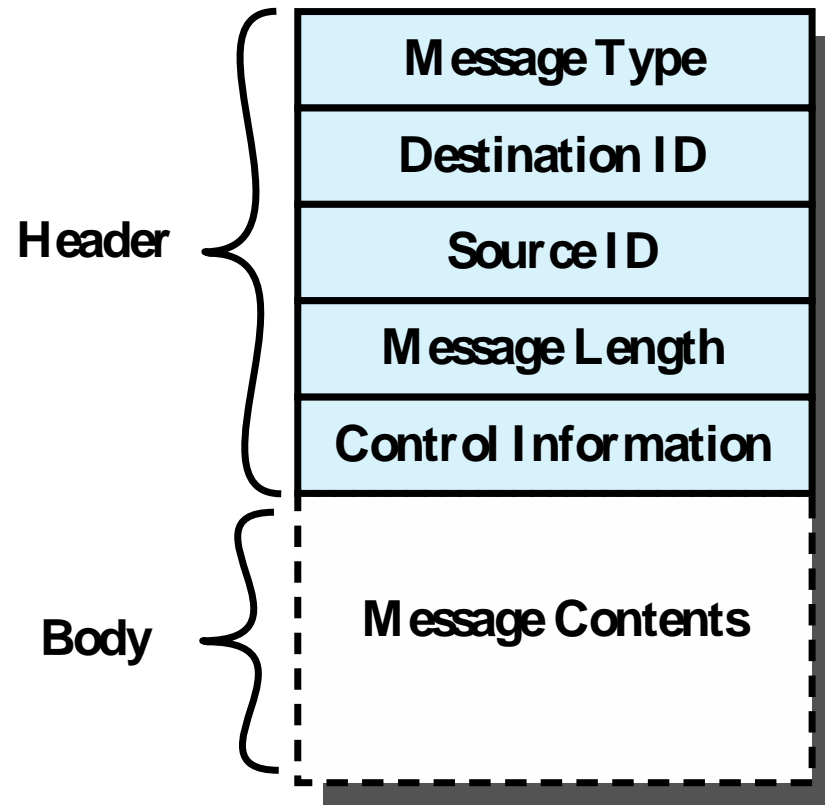


Figure 5.19 General Message Format

```

/* 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

```

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);
}

```

Figure 5.21

A Solution to the Bounded-Buffer Producer/Consumer Problem Using Messages

Init. Mayproduce is full

Init. Mayconsume is empty

Readers/Writers Problem

- A data area is shared among many processes
 - some processes only read the data area, (readers) and some only write to the data area (writers)
- Conditions that must be satisfied:
 1. any number of readers may simultaneously read the file
 2. only one writer at a time may write to the file
 3. if a writer is writing to the file, no reader may read it

```

/* 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);
}

```

Figure 5.22

**A Solution to
the
Readers/Write
rs Problem
Using
Semaphores:
Readers Have
Priority**

Readers only in the system	<ul style="list-style-type: none"> • <i>wsem</i> set • no queues
Writers only in the system	<ul style="list-style-type: none"> • <i>wsem</i> and <i>rsem</i> set • writers queue on <i>wsem</i>
Both readers and writers with read first	<ul style="list-style-type: none"> • <i>wsem</i> set by reader • <i>rsem</i> set by writer • all writers queue on <i>wsem</i> • one reader queues on <i>rsem</i> • other readers queue on <i>z</i>
Both readers and writers with write first	<ul style="list-style-type: none"> • <i>wsem</i> set by writer • <i>rsem</i> set by writer • writers queue on <i>wsem</i> • one reader queues on <i>rsem</i> • other readers queue on <i>z</i>

Table 5.6
State of the Process Queues for Program of Figure 5.23

```

/*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);
    }
}
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);
}

```

Figure 5.23

A Solution to the Readers/Writers Problem Using Semaphores: Writers Have Priority

```

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++;
        }
    }
}

```

Figure 5.24

A Solution to the Readers/Writers Problem Using Message Passing

Summary

- Principles of concurrency
 - Race condition
 - OS concerns
 - Process interaction
 - Requirements for mutual exclusion
- Mutual exclusion: hardware support
 - Interrupt disabling
 - Special machine instructions
- Semaphores
 - Mutual exclusion
 - Producer/consumer problem
 - Implementation of semaphores
- Monitors
 - Monitor with signal
 - Alternate model of monitors with notify and broadcast
- Message passing
 - Synchronization
 - Addressing
 - Message format
 - Queueing discipline
 - Mutual exclusion
- Readers/writers problem
 - Readers have priority
 - Writers have priority