CHAPTER 6 AND 7 - SYNCHRONIZATION (TOOLS AND EXAMPLES)

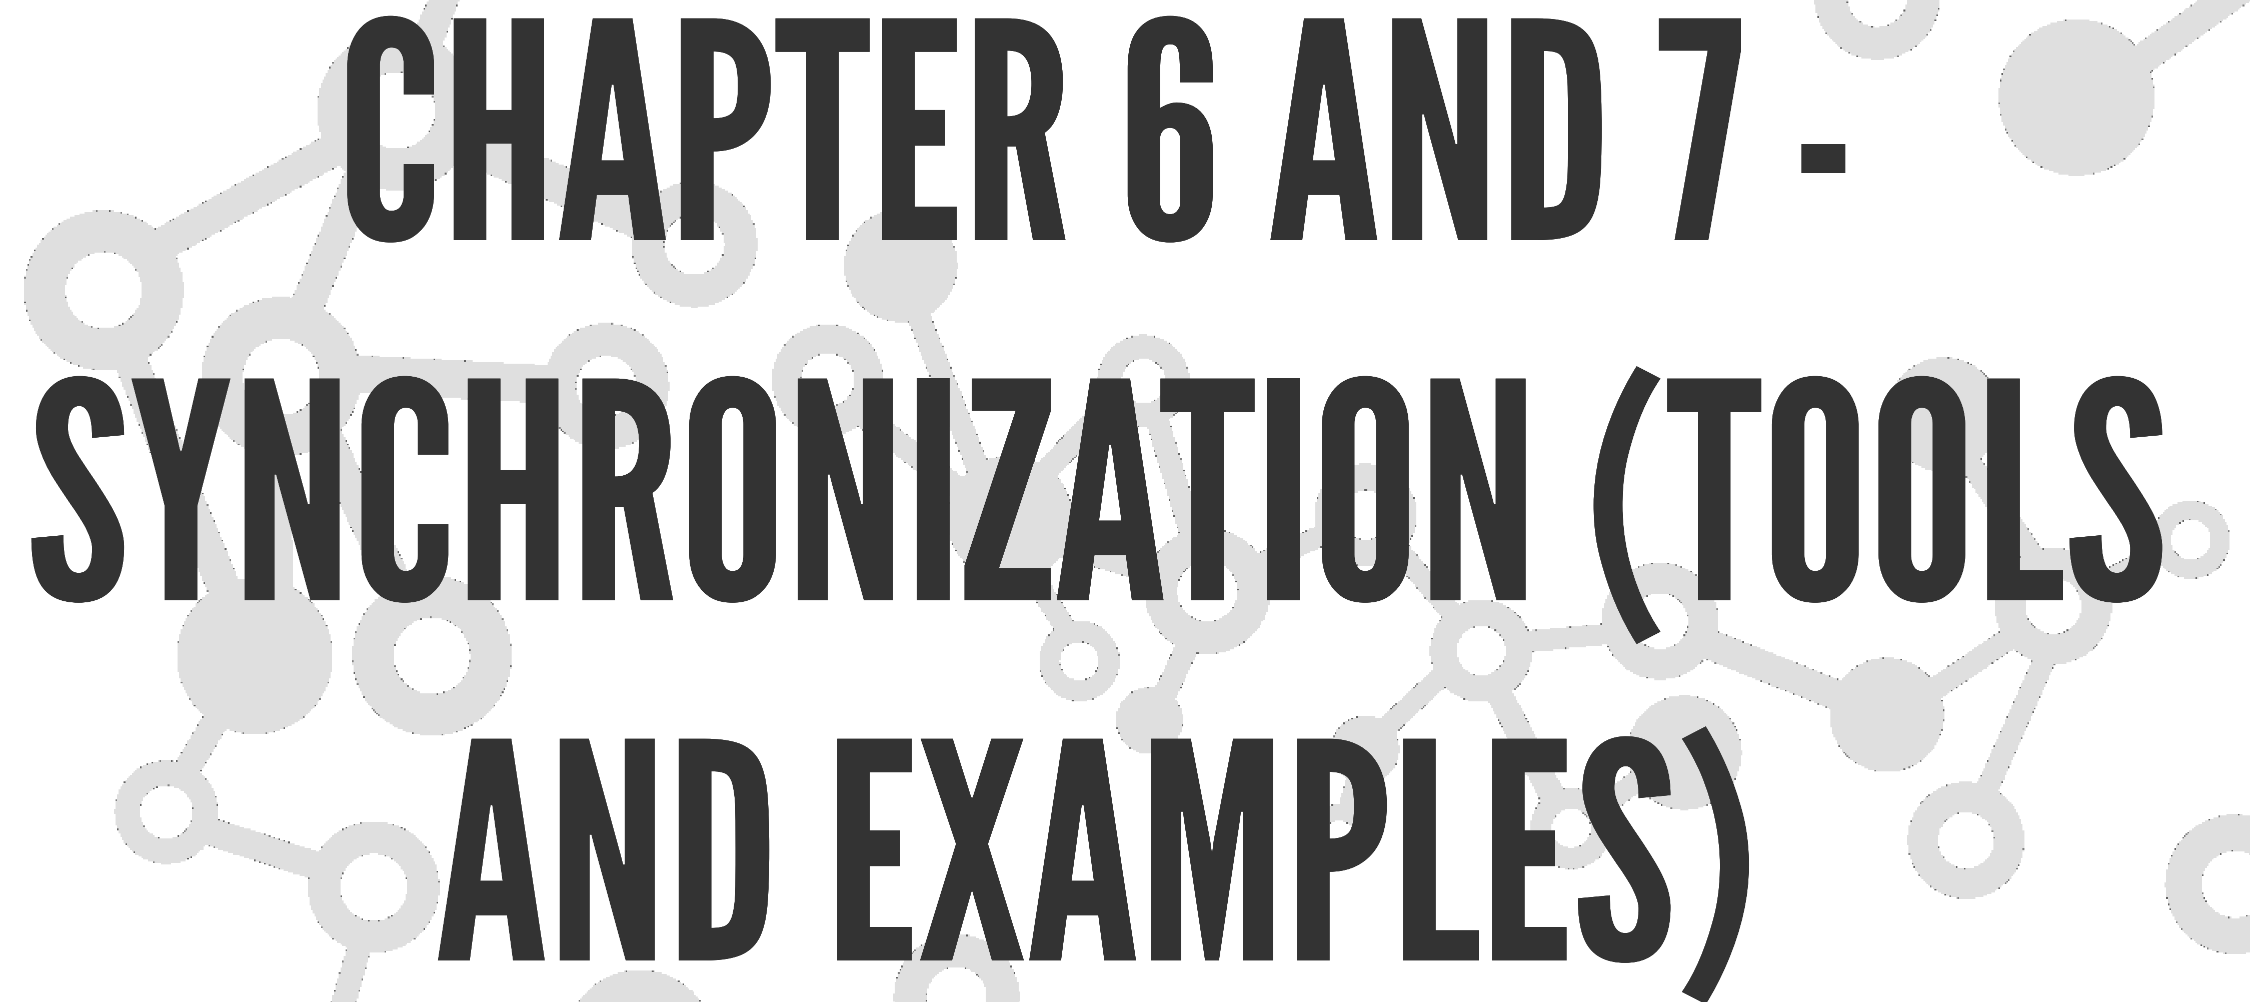

INTRODUCTION

A **cooperating process** can affect or be affected by other processes executing.

- share a logical address space (code and data)
- share data through shared memory or message passing

Concurrent access to shared data may result in data inconsistency

OBJECTIVES

- Introduce the critical-section problem → solutions used to ensure consistency of shared data.
- software and hardware solutions of the critical-section problem.
- examine classical process-synchronization problems.
- explore several tools that are used to solve process synchronization problems.

BACKGROUND

BACKGROUND

Processes can execute concurrently

Concurrent access to shared data may result in data inconsistency

Maintaining data consistency requires mechanisms to ensure the orderly execution of cooperating processes

THE PROBLEM

> Want to provide a solution to the consumer-producer problem that fills all the buffers.

Have integer counter that keeps track of the number of full buffers.

- Initially, counter is set to 0.
- incremented by the producer after it produces a new buffer
- decremented by the consumer after it consumes a buffer.

PRODUCER

```
while (true) {
    /* produce an item in next produced */
    while (counter == BUFFER_SIZE) {
       /* do nothing */
    }
    buffer[in] = next_produced;
    in = (in + 1) % BUFFER_SIZE;
    counter++;
}
```

CONSUMER

```
while (true) {
    while (counter == 0){
        /* do nothing */
    }
    next_consumed = buffer[out];
    out = (out + 1) % BUFFER_SIZE;
    counter--;
    /* consume the item in next consumed */
}
```

CONSUMER-PRODUCER PROBLEM

Works well separately!

May not work correctly when executed concurrently

IMPLEMENTATION

- `counter++` could be implemented as

```
register1 = counter
register1 = register1 + 1
counter = register1
```

- `counter--` could be implemented as

```
register2 = counter
register2 = register2 - 1
counter = register2
```

INITIALLY COUNTER=5

T_0	register1 = counter	register1 = 5
T_1	register1 = register1 + 1	register1 = 6
T_2	register2 = counter	register2 = 5
T_3	register2 = register2 – 1	register2 = 4
T_4	counter = register1	counter = 6
T_5	counter = register2	counter = 4

RACE CONDITION

A **Race Condition** is where several processes access and manipulate the same data concurrently and the outcome of the execution depends on the particular order in which the access takes place.

Need some synchronization to solve this

THE CRITICAL-SECTION PROBLEM

CRITICAL SECTION PROBLEM

Consider system of n processes $\{p_0, p_1, \ldots p_{n-1}\}$

Each process has **critical section** segment of code

- Process may be changing common variables, updating table, writing file, etc
- When one process in critical section, no other may be in its critical section

Th **critical section problem** is to design protocol to solve this

CRITICAL SECTION

General structure of process p_i is

```
do {
        entry section

            critical section

        exit section

            remainder section

} while (true);
```

SOLUTION TO CRITICAL SECTION PROBLEM

A solution must satisfy 3 requirements

- Mutual Exclusion
- Progress
- Bounded Waiting

MUTUAL EXCLUSION

If process P_i is executing in its critical section, then no other processes can be executing in their critical sections

PROGRESS

If no process is executing in its critical section and there exist some processes that wish to enter their critical section, then the selection of the processes that will enter the critical section next cannot be postponed indefinitely

BOUNDED WAITING

A bound must exist on the number of times that other processes are allowed to enter their critical sections after a process has made a request to enter its critical section and before that request is granted

- Assume that each process executes at a nonzero speed
- No assumption concerning relative speed of the n processes

KERNEL MODE EXAMPLE

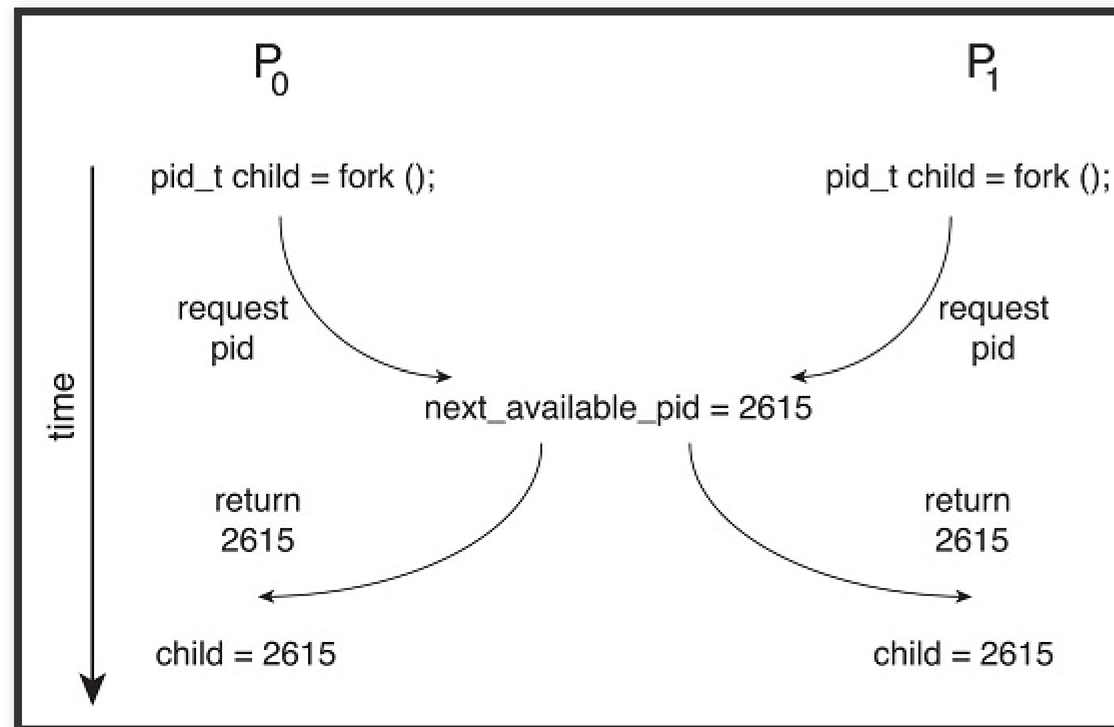

TWO APPROACHES

Two approaches depending on if kernel is preemptive or non-preemptive

- **Preemptive** – allows preemption of process when running in kernel mode
- **Non-preemptive** – runs until exits kernel mode, blocks, or voluntarily yields CPU
  - Essentially free of race conditions in kernel mode

PETERSON'S SOLUTION

A classic software-based solution to the critical-section

PETERSON'S SOLUTION

Good algorithmic description of solving the problem

Two process solution

> ❗ Assume that the load and store instructions are atomic; that is, cannot be interrupted

PETERSON'S SOLUTION

The two processes share two variables:

```
int turn;
Boolean flag[2]
```

The variable turn indicates whose turn it is to enter the critical section

The flag array is used to indicate if a process is ready to enter the critical section. `flag[i] = true` implies that process P_i is ready!

ALGORITHM FOR PROCESS P_I

```
while (true) {
        flag[i] = true;
        turn = j;
        while (flag[j] && turn == j)
                ;

        /* critical section */

        flag[i] = false;

        /*remainder section */
}
```

Provable that 1), 2) and 3) holds

PETERSONS SOLUTION ON MODERN ARCHITECTURE

Modern Architechtures: To improve system performance, processors and/or compilers may reorder read and write operations that have no dependencies

For a multithreaded application with shared data, the reordering of instructions may render inconsistent or unexpected results.

MODERN ARCHITECTURE

Shared between two threads

```
boolean flag = false;
int x = 0;
```

Thread 1

```
while (!flag)
   ;
print x;
```

Thread 2

```
x = 100;
flag = true;
```

PETERSONS SOLUTION ON MODERN ARCHITECTURE

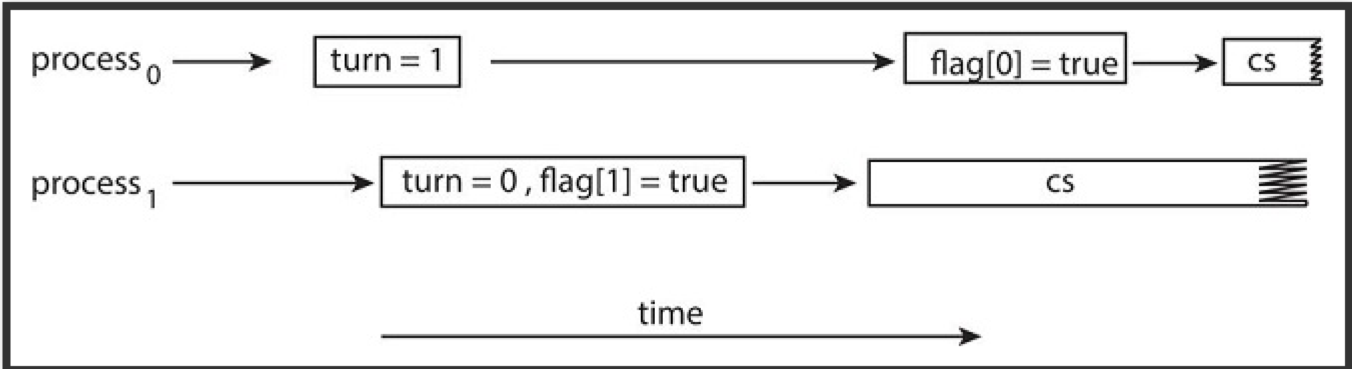

HARDWARE SUPPORT FOR SYNCHRONIZATION

MEMORY BARRIERS

Computer architecture memory guarantees it will provide to an application program:

1. **Strongly ordered:** Memory modification on one processor is immediately visible to all other processors.
2. **Weakly ordered:**, Modifications to memory on one processor may not be immediately visible to other processors.

Vary by processor type

MEMORY BARRIERS

Computer architectures provide instructions that can **force** any changes in memory to be propagated to all other processors ⇒ ensuring that memory modifications are visible to threads running on other processors. Such instructions are known as **memory barriers** or **memory fences**.

System ensures that all loads and stores are completed before any subsequent load or store operations are performed.

MEMORY BARRIERS

Thread 1

```
while (!flag)
   ;
memory_barrier();
print x;
```

Thread 2

```
x = 100;
memory_barrier();
flag = true;
```

HARDWARE INSTRUCTIONS

Modern machines provide special atomic hardware instructions

- Atomic = non-interruptible
- Either test memory word and set value: `test_and_set()`
- Or swap contents of two memory words: `compare_and_swap()`

TEST AND SET INSTRUCTION

Definition

```
boolean test_and_set(boolean *target) {
    boolean read_value = *target;
    *target = true;

    return read_value;
}
```

SOLUTION USING TEST AND SET

Shared boolean variable lock, initialized to FALSE

```
do {
    while( test_and_set(&lock) ) {
        ; /* DO NOTHING */
    }
    /* Chritical section */

    lock = false;

    /* remainder section */
} while(true);
```

COMPARE AND SWAP INSTRUCTION

Definition

```
int compare_and_swap(int *value, int expected, int new_value) {

    int temp = *value;

    if( *value == expected) {
        *value = new_value;
    }
    return temp;
}
```

SOLUTION USING COMPARE_AND_SWAP

Shared Boolean variable lock initialized to FALSE; Each process has a local Boolean variable key

```
do {
    while( compare_and_swap(&lock, 0, 1) != 0 ) {
        ; /* DO NOTHING */
    }

    /* Chritical section */

    lock = 0;

    /* remainder section */
} while(true);
```

BOUNDED-WAITING MUTUAL EXCLUSION WITH TEST AND SET

```
do {
    waiting[i] = true;
    key = true;
    while (waiting[i] && key)
        key = test and set(&lock);
    waiting[i] = false;

    /* critical section */

    j = (i + 1) % n;
    while ((j != i) && !waiting[j])
        j = (j + 1) % n;
    if (j == i)
        lock = false;
    else
    waiting[j] = false;

    /* remainder section */
} while (true);
```

ATOMIC VARIABLES

Typically, the `compare_and_swap()` instruction is not used directly to provide mutual exclusion.

It is used as a basic building block for constructing other tools that solve the critical-section problem.

One such tool is an **atomic variable**, which provides atomic operations on basic data types such as integers and booleans.

ATOMIC VARIABLES

For atomic integer `sequence`

```
increment(&sequence);
```

Can be implemented like

```
void increment(atomic_int *v) {

  int temp;
  do {
    temp = *v;
  }
  while (temp != compare_and_swap(v, temp, temp+1));
}
```

MUTEX LOCKS

MUTEX LOCKS

Previous solutions are complicated and generally inaccessible to application programmers

OS designers build software tools to solve critical section problem

Simplest is mutex lock

Product critical regions with it by first `acquire()` a lock then `release()` it

SOLUTION TO THE CRITICAL-SECTION PROBLEM USING MUTEX LOCKS

```
while (true) {

    // acquire lock

       // critical section

    // release lock

       //  remainder section
}
```

MUTEX LOCKS

Calls to `acquire()` and `release()` must be atomic

Boolean variable indicating if lock is available or not → Usually implemented via hardware atomic instructions

But this solution requires **busy waiting** → This lock therefore called a **spinlock**

ACQUIRE()

```
acquire() {
    while (!available) {
        ; /* busy wait */
    }
    available = false;;
}
```

RELEASE()

```
release() {
    available = true;
}
```

SEMAPHORES

SEMAPHORE

Synchronization tool that does not require busy waiting

Semaphore S – integer variable

Two standard operations modify S

```
wait()
signal()
```

Originally called P ( ) and V ( )

SEMAPHORE

Can only be accessed via two indivisible (atomic) operations

```
wait (S) {
    while (S <= 0) {
        ; // busy wait
    }
    S--;
}
```

```
signal (S) {
    S++;
}
```

SEMAPHORE USAGE

- **Counting semaphore** – integer value can range over an unrestricted domain
- **Binary semaphore** – integer value can range only between 0 and 1
  - Then a mutex lock
- Can implement a counting semaphore S as a binary semaphore
- Can solve various synchronization problems

SEMAPHORE USAGE

Consider P1 and P2 that require S1 to happen before S2

```
P1:
  S1;
  signal(synch);
P2:
  wait(synch);
  S2;
```

SEMAPHORE IMPLEMENTATION

Must guarantee that no two processes can execute `wait()` and `signal()` on the same semaphore at the same time

Thus, implementation becomes the critical section problem where the wait and signal code are placed in the critical section

SEMAPHORE IMPLEMENTATION

- Could now have busy waiting in critical section implementation
  - But implementation code is short
  - Little busy waiting if critical section rarely occupied

SEMAPHORE IMPLEMENTATION - NO BUSY WAITING

With each semaphore there is an associated waiting queue

Each entry in a waiting queue has two data items:

- value (of type integer)
- pointer to next record in the list

SEMAPHORE IMPLEMENTATION - NO BUSY WAITING

Two operations:

- `sleep` – place the process invoking the operation on the appropriate waiting queue
- `wakeup` – remove one of processes in the waiting queue and place it in the ready queue

SEMAPHORE IMPLEMENTATION - NO BUSY WAITING

```
typedef struct{
    int value;
    struct process *list;
} semaphore;
```

SEMAPHORE IMPLEMENTATION - NO BUSY WAITING

```
wait(semaphore *S) {
    S->value--;
    if (S->value < 0) {
        add this process to S->list;
        sleep();
    }
}
```

SEMAPHORE IMPLEMENTATION - NO BUSY WAITING

```
signal(semaphore *S) {
    S->value++;
    if (S->value <= 0) {
        remove a process P from S->list;
        wakeup(P);
    }
}
```

MONITORS

PROBLEMS WITH SEMAPHORES

- Incorrect use of semaphore operations:
  - `signal(mutex) …. wait(mutex)`
  - `wait(mutex) … wait(mutex)`
  - Omitting of wait(mutex) or signal(mutex) (or both)
- Deadlock and starvation

MONITORS

A high-level abstraction that provides a convenient and effective mechanism for process synchronization

Abstract data type, internal variables only accessible by code within the procedure

Only one process may be active within the monitor at a time

MONITORS

```
monitor monitor name
{

   /* shared variable declarations */

   function P1 ( . . . ) {
       . . .
   }

   function P2 ( . . . ) {
       . . .
   }

               .
               .
               .

   function Pn ( . . . ) {
       . . .
   }

   initialization_code ( . . . ) {
       . . .
   }
}
```

SCHEMATIC VIEW OF A MONITOR

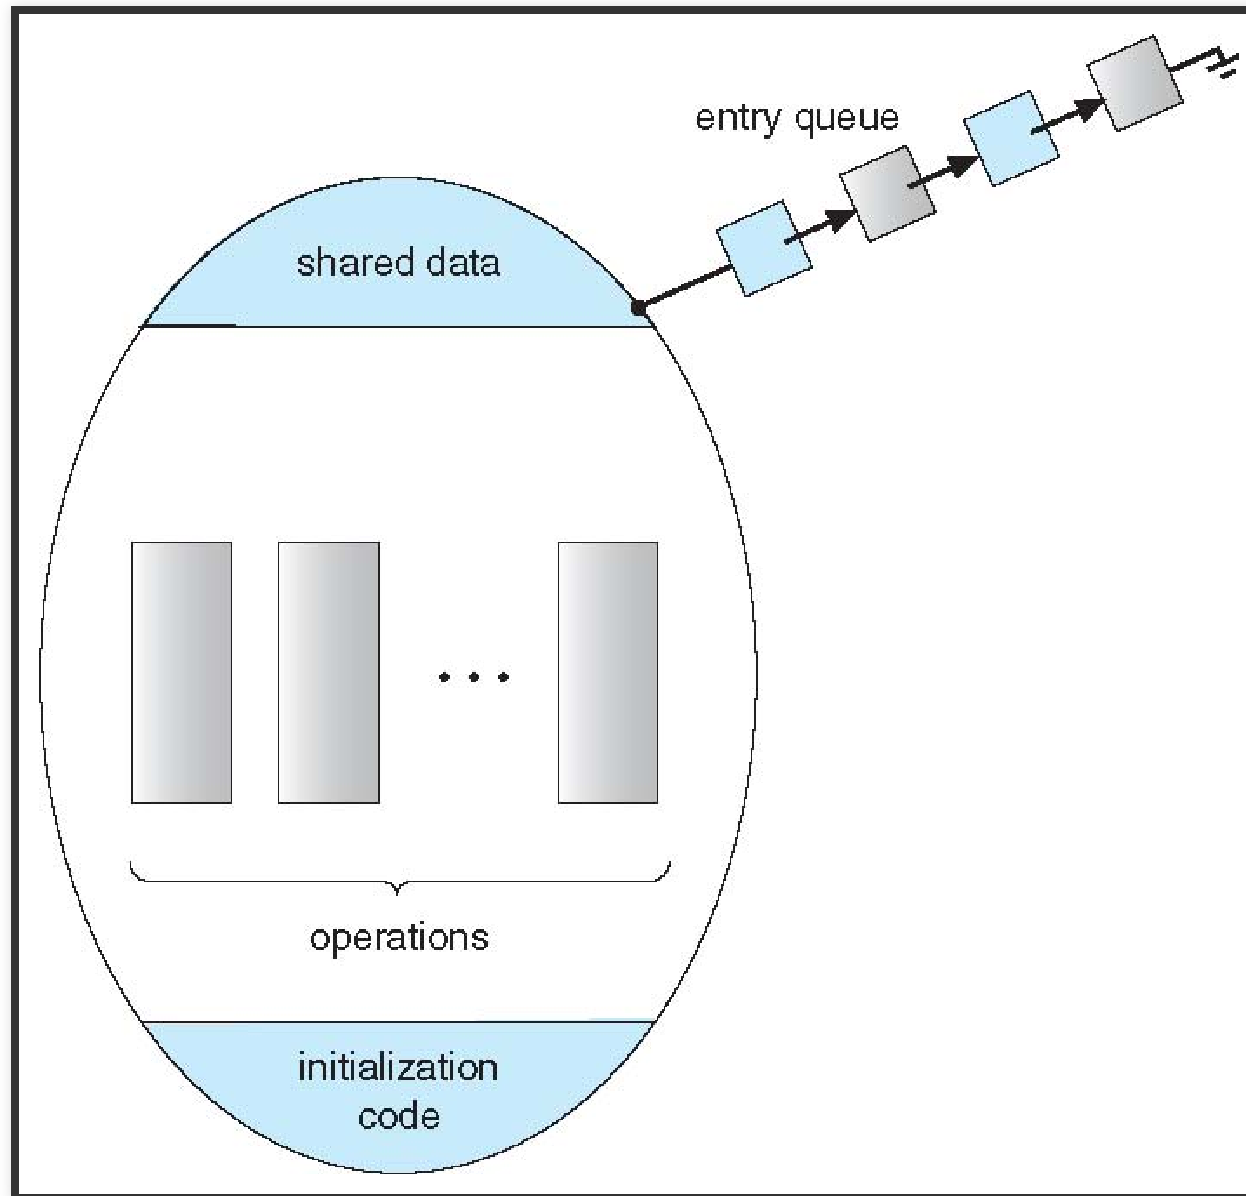

MONITORS

Not powerful enough to model some synchronization schemes

To be more powerfull → `condition` construct

CONDITION VARIABLES

```
condition x, y;
```

- Two operations on a condition variable:
- `x.wait()` – a process that invokes the operation is suspended until `x.signal()`
- `x.signal()` – resumes one of processes (if any) that invoked `x.wait()`
    - If no `x.wait()` on the variable, then it has no effect on the variable

MONITOR WITH CONDITION VARIABLES

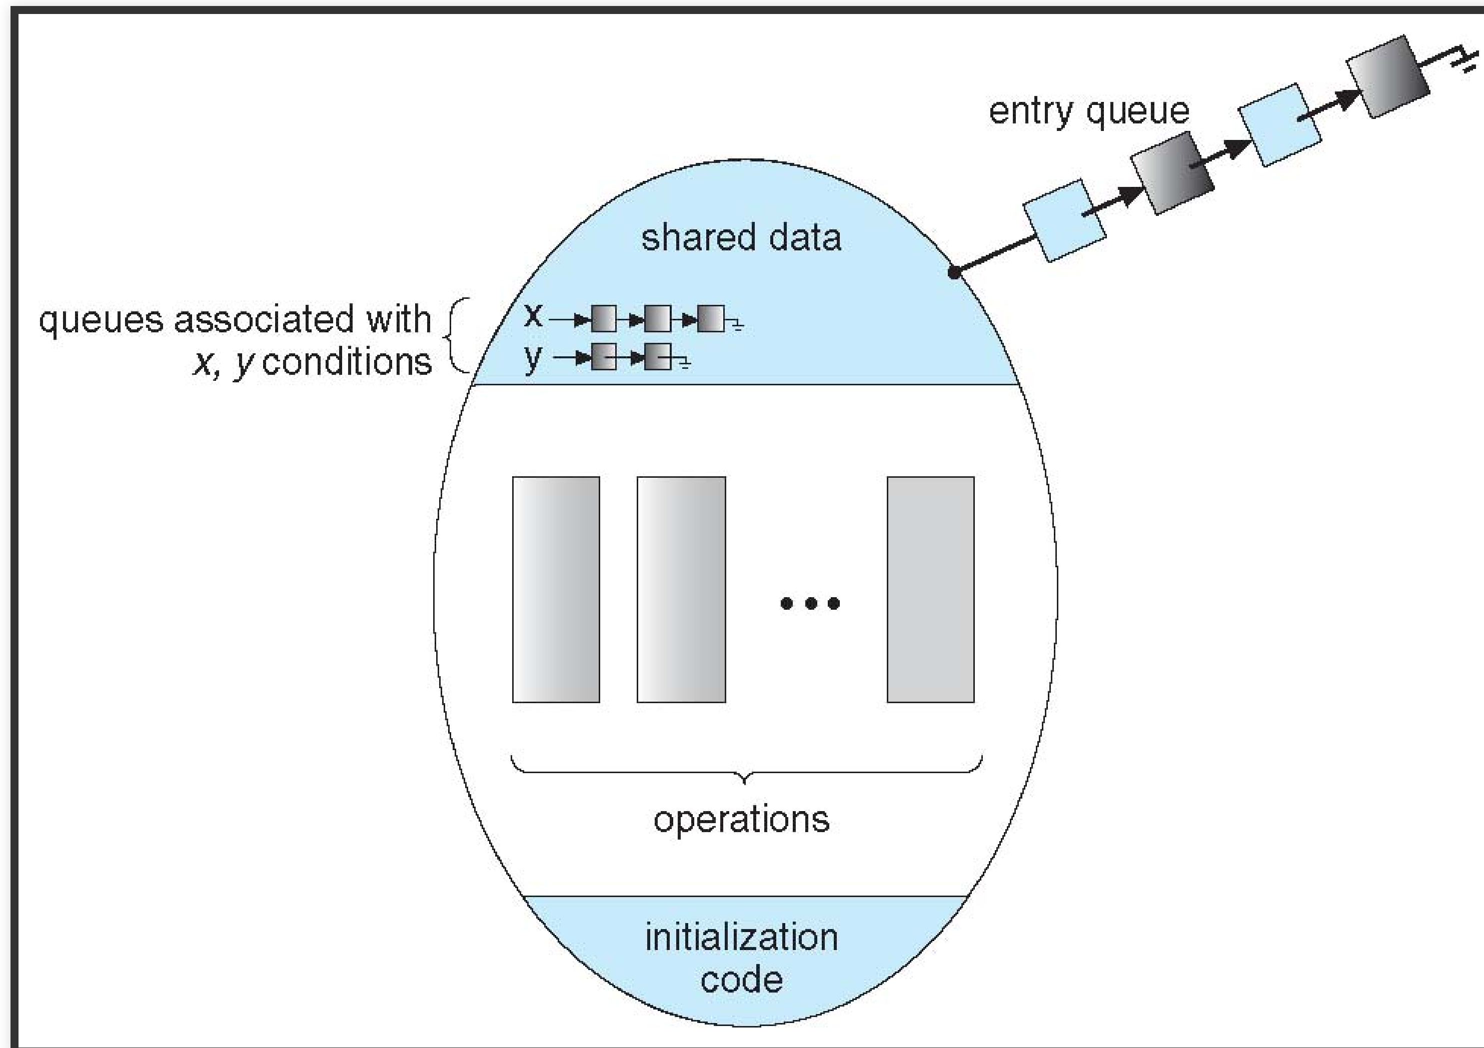

CONDITION VARIABLES CHOICES

- If process P invokes `x.signal()`, with Q in `x.wait()` state, what should happen next?
  - If Q is resumed, then P must wait
- Options include
  - **Signal and wait** – P waits until Q leaves monitor or waits for another condition
  - **Signal and continue** – Q waits until P leaves the monitor or waits for another condition

CONDITION VARIABLES CHOICES

- Both have pros and cons – language implementer can decide
- Monitors implemented in Concurrent Pascal compromise → P executing signal immediately leaves the monitor, Q is resumed
- Implemented in other languages including Mesa, C#, Java

MONITOR IMPLEMENTATION USING SEMAPHORES

```
semaphore mutex; // (initially = 1)
semaphore next;  // (initially = 0)
int next_count = 0;
```

MONITOR IMPLEMENTATION USING SEMAPHORES

Each procedure F will be replaced by

```
wait(mutex);
    ...
  body of F;
    ...
if (next_count > 0) {
    signal(next)
} else {
    signal(mutex);
}
```

Mutual exclusion within a monitor is ensured

MONITOR IMPLEMENTATION – CONDITION VARIABLES

For each condition variable x, we have:

```
semaphore x_sem; // (initially = 0)
int x_count = 0;
```

MONITOR IMPLEMENTATION – CONDITION VARIABLES

The operation `x.wait` can be implemented as:

```
x_count++;
if (next_count > 0) {
    signal(next);
} else {
    signal(mutex);
}
wait(x_sem);
x_count--;
```

MONITOR IMPLEMENTATION – CONDITION VARIABLES

The operation `x.signal` can be implemented as:

```
if (x_count > 0) {
    next_count++;
    signal(x_sem);
    wait(next);
    next_count--;
}
```

RESUMING PROCESSES WITHIN A MONITOR

If several processes queued on condition x, and `x.signal()` executed, which should be resumed?

- FCFS frequently not adequate
- conditional-wait construct of the form `x.wait(c)`
  - Where c is priority number
  - Process with lowest number (highest priority) is scheduled next

A MONITOR TO ALLOCATE SINGLE RESOURCE

```
monitor ResourceAllocator {
    boolean busy;
    condition x;

    void acquire(int time) {
        if (busy)
            x.wait(time);
        busy = true;
    }

    void release() {
        busy = false;
        x.signal();
    }

    initialization code() {
        busy = false;
    }
}
```

MONITOR SHORTCOMMINGS

Monitor concept cannot guarantee that the preceding access sequence will be observed

- A process might access a resource without first gaining access permission to the resource.
- A process might never release a resource once it has been granted access to the resource.
- A process might attempt to release a resource that it never requested.
- A process might request the same resource twice (without first releasing the resource).

LIVENESS

Liveness refers to a set of properties that a system must satisfy to ensure that processes make progress during their execution life cycle.

A process waiting indefinitely under the circumstances just described is an example of a *"liveness failure."*

DEADLOCK

Deadlock – two or more processes are waiting indefinitely for an event that can be caused by only one of the waiting processes

DEADLOCK

Let S and Q be two semaphores initialized to 1

```
P0                  P1
wait(S);            wait(Q);
wait(Q);            wait(S);
...                 ...
signal(S);          signal(Q);
signal(Q);          signal(S);
```

PRIORITY INVERSION

- Assume 3 processes with priority `L < M < H`
- Resource R is being held by `L`
- Proces `H` require resource R, waits for `L`
- Proces `M` becomes runnable, preempting `L`
- Result: `M` affects how long process `H` must wait for `L`

PRIORITY INVERSION

Priority Inversion – Scheduling problem when lower-priority process holds a lock needed by higher-priority process

Solved via **priority-inheritance protocol**

EVALUATION

When should which tool be used?

LOW LEVEL TOOLS

Low level

- test and set
- compare and swap (CAS)

Typically used to implement other tools

CAS-based approaches are considered an optimistic approach—you optimistically first update a variable and then use collision detection to see if another thread is updating the variable concurrently.

CAS-BASED OR TRADITIONAL SYNC.

- **Uncontended:** Both options are generally fast, CAS protection will be somewhat faster than traditional synchronization.
- **Moderate contention:** CAS protection will be faster—possibly much faster—than traditional synchronization.
- **High contention:** Under very highly contended loads, traditional synchronization will ultimately be faster than CAS-based synchronization.

MODERATE CONTENTION

CAS operation succeeds most of the time, and when it fails, it will iterate through the loop a feww times

Mutual-exclusion locking, any attempt to acquire a contended lock will result in a more complicated—and time-intensive—code path that suspends a thread and places it on a wait queue, requiring a context switch to another thread.

RULES OF THUMB

- Atomic integers are much lighter weight than traditional locks
- Spinlocks are used on multiprocessor systems when locks are held for short durations
- Controlling access to a finite number of resources—a counting semaphore is generally more appropriate than a mutex lock

ONGOING RESEARCH

- Designing compilers that generate more efficient code.
- Developing languages that provide support for concurrent programming.
- Improving the performance of existing libraries and APIs.

CLASSIC PROBLEMS OF SYNCHRONIZATION

CLASSIC PROBLEMS OF SYNCHRONIZATION

Classical problems used to test newly-proposed synchronization schemes

- Bounded-Buffer Problem
- Readers and Writers Problem
- Dining-Philosophers Problem

BOUNDED-BUFFER PROBLEM

```
int n;
semaphore mutex = 1;
semaphore empty = n;
semaphore full = 0;
```

PRODUCER PROCESS

```
do {
        . . .
    /* produce an item in next produced */
        . . .
    wait(empty);
    wait(mutex);
        . . .
    /* add next produced to the buffer */
        . . .
    signal(mutex);
    signal(full);
} while (true);
```

CONSUMER PROCESS

```
do {
    wait(full);
    wait(mutex);
        . . .
    /* remove an item from buffer to next consumed */
        . . .
    signal(mutex);
    signal(empty);
        . . .
    /* consume the item in next consumed */
        . . .
} while (true);
```

READERS-WRITERS PROBLEM

A data set is shared among a number of concurrent processes

- **Readers** – only read the data set; they do not perform any updates
- **Writers** – can both read and write

READERS-WRITERS PROBLEM

- **Problem** – allow multiple readers to read at the same time
  - Only one single writer can access the shared data at the same time
- Several variations of how readers and writers are treated
  - all involve priorities

READERS-WRITERS PROBLEM

Shared Data

- Data set

and

```
int read_count = 0;
semaphore rw_mutex = 1;
semaphore mutex = 1;
```

WRITERS PROCESS

```
do {
    wait(rw_mutex);
        . . .
    /* writing is performed */
        . . .
    signal(rw_mutex);
} while (true);
```

READER PROCESS

```
do {
    wait(mutex);
    read count++;
    if (read_count == 1)
        wait(rw_mutex);
    signal(mutex);
        . . .
    /* reading is performed */
        . . .
    wait(mutex);
    read count--;
    if (read_count == 0)
        signal(rw_mutex);
    signal(mutex);
} while (true);
```

PROBLEM VARIATIONS

- First variation – no reader kept waiting unless writer has permission to use shared object
- Second variation – once writer is ready, it performs write asap
- Both may have starvation leading to even more variations
- Problem is solved on some systems by kernel providing reader-writer locks

USEFULL WHEN

- In applications where it is easy to identify which processes only read shared data and which processes only write shared data.
- In applications that have more readers than writers.

DINING-PHILOSOPHERS PROBLEM

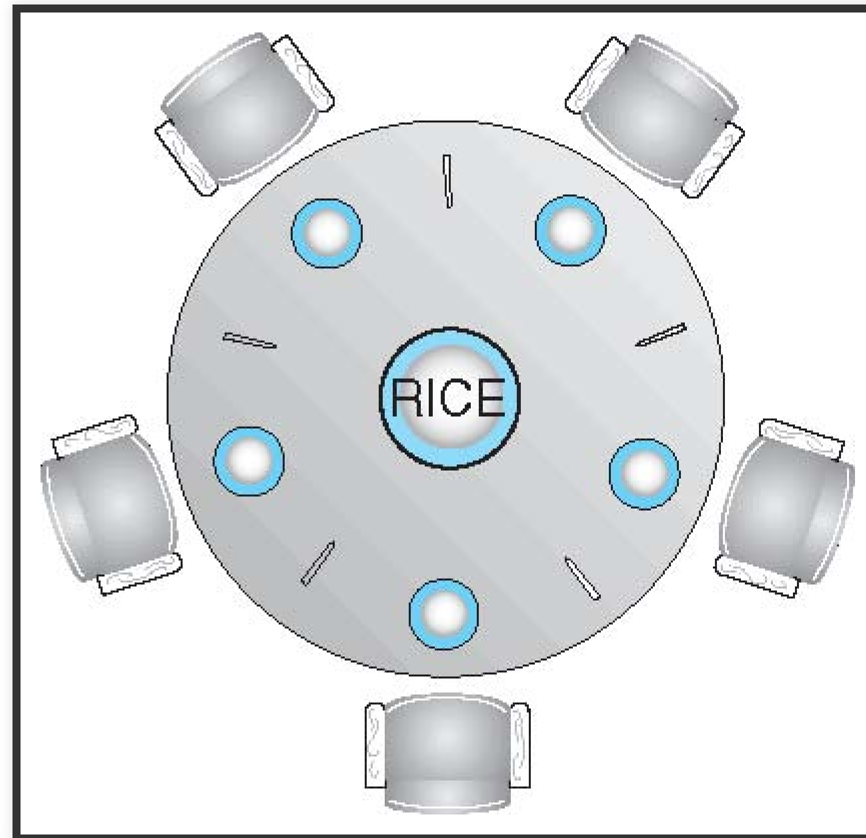

DINING-PHILOSOPHERS PROBLEM

Philosophers spend their lives thinking and eating

Don’t interact with their neighbors, occasionally try to pick up 2 chopsticks (one at a time) to eat from bowl

Need both to eat, then release both when done

DINING-PHILOSOPHERS PROBLEM

Shared data

- Bowl of rice (data set)
- Semaphore `chopstick[5]` initialized to `1`

DINING-PHILOSOPHERS PROBLEM

```
do {
    wait(chopstick[i]);
    wait(chopstick[(i+1) % 5]);
        . . .
    /* eat for awhile */
        . . .
    signal(chopstick[i]);
    signal(chopstick[(i+1) % 5]);
        . . .
    /* think for awhile */
        . . .
} while (true);
```

MONITOR TO DINING PHILOSOPHERS

```
monitor DiningPhilosophers {
    enum { THINKING, HUNGRY, EATING } state[5];
    condition self[5];

    void pickup(int i) {
        state[i] = HUNGRY;
        test(i);
        if (state[i] != EATING)
            self[i].wait();
    }

    void putdown(int i) {
        state[i] = THINKING;
        test((i + 4) % 5);
        test((i + 1) % 5);
    }

    ....
```

SOLUTION TO DINING PHILOSOPHERS

```
    ...

    void test(int i) {
        if ((state[(i + 4) % 5] != EATING) &&
            (state[i] == HUNGRY) &&
            (state[(i + 1) % 5] != EATING)) {
                state[i] = EATING;
                self[i].signal();
        }
    }

    initialization code() {
        for (int i = 0; i < 5; i++)
        state[i] = THINKING;
    }
}
```

SOLUTION TO DINING PHILOSOPHERS

Each philosopher `i` invokes the operations `pickup()` and `putdown()` in the following sequence:

```
DiningPhilosophers.pickup(i);
    EAT
DiningPhilosophers.putdown(i);
```

No deadlock, but starvation is possible

SYNCHRONIZATION WITHIN THE KERNEL

SYNCHRONIZATION IN WINDOWS

- Uses interrupt masks to protect access to global resources on single-processor systems
- Uses spinlocks on multiprocessor systems
  - Spinlocking-thread will never be preempted
- Also provides **dispatcher objects** in user-land which may act mutexes, semaphores, events, and timers

DISPATCHER OBJECTS

Using a dispatcher object, threads synchronize according to several different mechanisms, including mutex locks, semaphores, events, and timers.

- **Events**: An event acts much like a condition variable
- **Timers** notify one or more thread when time expired
- Dispatcher objects either signaled-state (object available) or non-signaled state (thread will block)

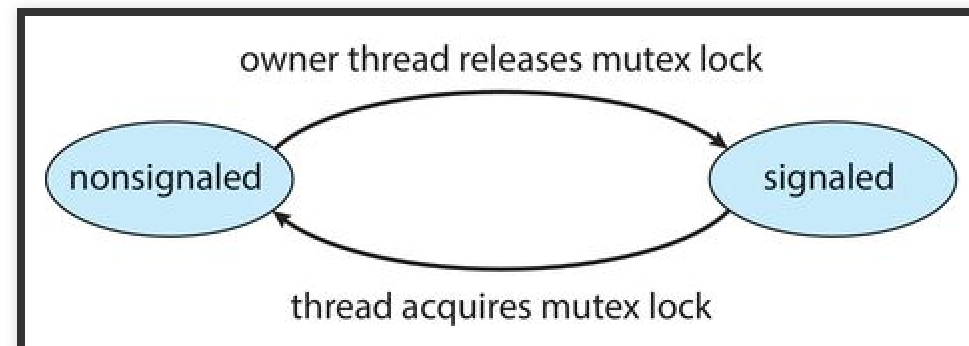

LINUX SYNCHRONIZATION

- Prior to kernel Version 2.6, disables interrupts to implement short critical sections
- Version 2.6 and later, fully preemptive
- Linux provides:
  - atomic integer (*atomic_t*)
  - semaphores
  - spinlocks
  - reader-writer versions of both

LINUX SYNCHRONIZATION

- On single-cpu system, spinlocks replaced by enabling and disabling kernel preemption

Single Processor	Multiple Processors
Disable kernel preemption	Acquire spin lock
Enable kernel preemption	Release spin lock

When a lock must be held for a longer period, semaphores or mutex locks are appropriate for use.

POSIX SYNCHRONIZATION

PTHREADS SYNCHRONIZATION

Pthreads API is OS-independent

It provides:

- mutex locks
- condition variables
  - Non-portable extensions include:
  - read-write locks
  - spinlocks

PTHREADS SYNCHRONIZATION

```
#include <pthread.h>

pthread_mutex_t mutex;

/* create and initialize  the mutex lock */
pthread_mutex_init(&mutex,NULL);

/* acquire the mutex lock */
pthread_mutex_lock(&mutex);

/* critical section */

/* release the mutex lock */
pthread_mutex_unlock(&mutex);
```

POSIX NAMED SEMAPHORES

The function `sem_open()` is used to create and open a POSIX named sempahore:

```
#include <semaphore.h>

sem_t *sem;

/* Create the semaphore and initialize it to 1 */
sem = sem_open(“SEM”, O_CREAT, 0666, 1);

/* acquire the semaphore */
sem_wait(sem);

/* critical section */

/* release the semaphore */
sem_post(sem);
```

POSIX UNNAMED SEMAPHORES

An unnamed semaphore is created and initialized using the sem_init() function, which is passed three parameters:

1. A pointer to the semaphore
2. A flag indicating the level of sharing
3. The semaphore's initial value

POSIX UNNAMED SEMAPHORES

```
#include <semaphore.h>

sem_t sem;

/* Create the semaphore and initialize it to 1 */
sem_init(&sem, 0, 1);

/* acquire the semaphore */
sem_wait(&sem);

/* critical section */

/* release the semaphore */
sem_post(&sem);
```

SYNCHRONIZATION IN JAVA

LOCKS

Every object in Java has associated with it a single lock.

When a method is declared to be synchronized, calling the method requires owning the lock for the object. We declare a synchronized method by placing the synchronized keyword in the method definition

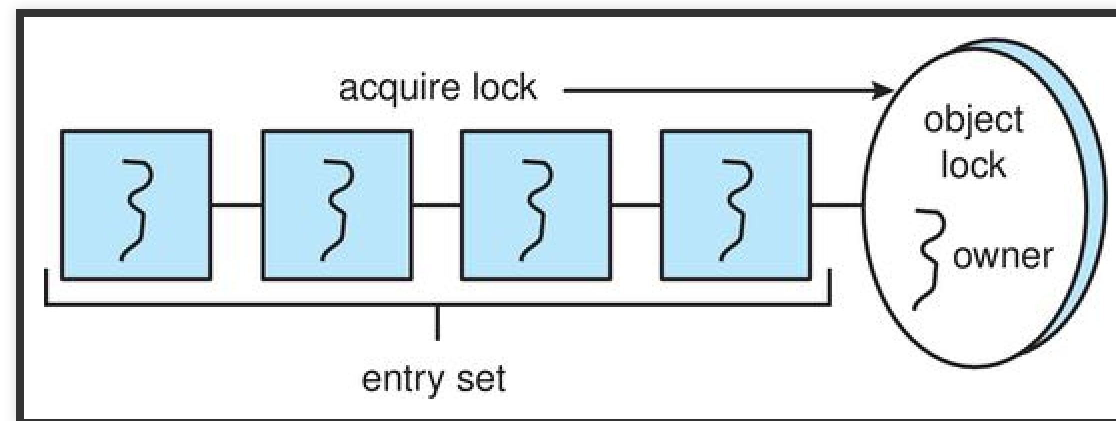

ENTRY AND WAIT SETS

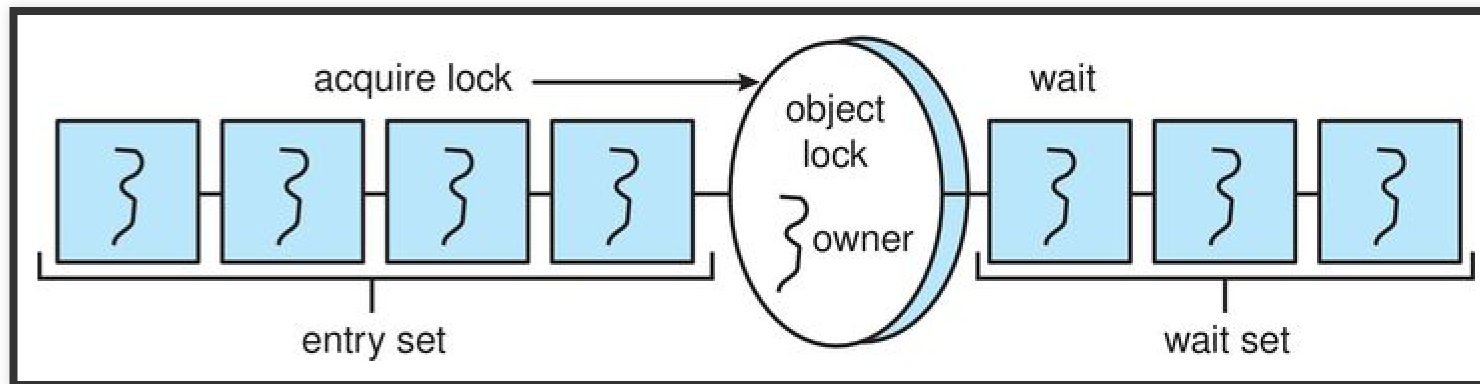

JAVA MONITORS

```
public class BoundedBuffer<E> {

   private static final int BUFFER_SIZE = 5;
   private int count, in, out;
   private E[] buffer;

   public BoundedBuffer() {
      count = 0;
      in = 0;
      out = 0;
      buffer = (E[]) new Object[BUFFER_SIZE];
   }

   /* Producers call this method */
   public synchronized void insert(E item) {
      /* See Next slide */
   }

   /* Consumers call this method */
   public synchronized E remove() {
      /* See Figure 7.11 */
   }
}
```

JAVA MONITORS

```
/* Producers call this method */
public synchronized void insert(E item) {
       while (count == BUFFER_SIZE) {
           try {
               wait();
           }  catch (InterruptedException ie) { }
       }

       buffer[in] = item;
       in = (in + 1) % BUFFER_SIZE;

       count++;
       notify();
}
```

JAVA MONITORS

```
/* Consumers call this method */
public synchronized E remove() {

       E item;
       while (count == 0) {
          try {
             wait();
          }  catch (InterruptedException ie) { }
       }

       item = buffer[out];
       out = (out + 1) % BUFFER_SIZE;

       count--;
       notify();

       return item;
}
```

REENTRANT LOCKS

```
Lock key = new ReentrantLock();

key.lock();

try {
  /* critical section */
} finally {
  key.unlock();
}
```

SEMAPHORES

```
Semaphore sem = new Semaphore(1);

try {
  sem.acquire();

  /* critical section */

} catch (InterruptedException ie) {
} finally {
  sem.release();
}
```

ALTERNATIVE APPROACHES

TRANSACTIONAL MEMORY

originated in database theory

A memory transaction is a sequence of memory read–write operations that are atomic. If all operations in a transaction are completed, the memory transaction is committed. Otherwise, the operations must be aborted and rolled back.

EXAMPLE USING LOCKS

```
void update() {
    acquire();
    /* Modify Shared data */
    release();
}
```

TRANSACTIONAL MEMORY

```
void update() {
    atomic {
      /* Modify Shared data */
    }
}
```

TRANSACTIONAL MEMORY

- transactional memory system—not the developer is responsible for guaranteeing atomicity.
- no locks are involved, deadlock is not possible.
- can identify which statements in atomic blocks can be executed concurrently
  - concurrent read access to a shared variable.
- Software or hardware solution

OPENMP

```
void update(int value) {

    #pragma omp critical
    {
        counter += value;
    }
}
```

FUNCTIONAL PROGRAMMING LANGUAGES

- C, C++, Java, and C# → imperative or procedural languages.
  - program state is mutable, as variables may be assigned different values over time.

FUNCTIONAL PROGRAMMING LANGUAGES

- Erlang, Haskell, Frege and Scala → functional programming languages
  - do not maintain state.
  - once a variable has been defined and assigned a value, its value is immutable
  - no issues like race conditions and deadlocks.

QUESTIONS

BONUS