

## Outline for Today

- Objective:
  - To continue talking about the [critical section problem](#) and get more practice thinking about possible [interleavings](#).
  - Start talking about [synchronization primitives](#).
  - Introduce other “classic” concurrency problems
- Administrative details:
  - Look on the web for TAs’ office hours or check newsgroup for UTAs’ office hours for assignment 1.

1

## Semaphores

- Well-known synchronization abstraction
- Defined as a non-negative integer with two atomic operations
  - P(s) - [wait until  $s > 0$ ;  $s--$ ]
  - V(s) - [ $s++$ ]
- The atomicity and the waiting can be implemented by either busywaiting or blocking solutions.

4

## Semaphore Usage

- Binary semaphores can provide mutual exclusion (solution of critical section problem)
- Counting semaphores can represent a resource with multiple instances (e.g. solving producer/consumer problem)
- Signaling events (persistent events that stay relevant even if nobody listening right now)

5

## The Critical Section Problem

```
while (1)
{ ...other stuff...
  P(mutex)
  critical section
  V(mutex)
}
```

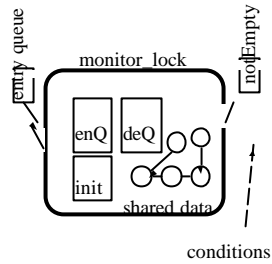
Semaphore:  
mutex initially 1

Fill in the boxes

6

## Monitor Abstraction

- Encapsulates shared data and operations with mutual exclusive use of the object (an associated *lock*).
- Associated *Condition Variables* with operations of *Wait* and *Signal*.



7

## Condition Variables

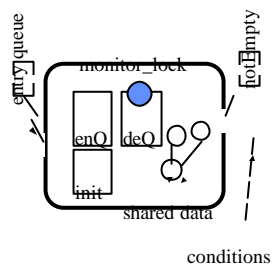
- We build the monitor abstraction out of a lock (for the mutual exclusion) and a set of associated condition variables.
- *Wait on condition*: releases lock held by caller, caller goes to sleep on condition's queue. When awakened, it must reacquire lock.
- *Signal condition*: wakes up one waiting thread.
- *Broadcast*: wakes up all threads waiting on this condition.

8

## Monitor Abstraction

```

EnQ: {acquire (lock);
      if (head == null)
        {head = item;
         signal (lock, notEmpty);}
      else tail->next = item;
      tail = item;
      release(lock);}
deQ: {acquire (lock);
      if (head == null)
        wait (lock, notEmpty);
      item = head;
      if (tail == head) tail = null;
      head=item->next;
      release(lock);}
    
```

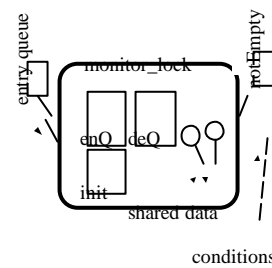


9

## Monitor Abstraction

```

EnQ: {acquire (lock);
      if (head == null)
        {head = item;
         signal (lock, notEmpty);}
      else tail->next = item;
      tail = item;
      release(lock);}
deQ: {acquire (lock);
      if (head == null)
        wait (lock, notEmpty);
      item = head;
      if (tail == head) tail = null;
      head=item->next;
      release(lock);}
    
```



10

## Monitor Abstraction

```

EnQ: {acquire (lock);
      if (head == null)
        {head = item;
         signal (lock, notEmpty);}
      else tail->next = item;
      tail = item;
      release(lock);}
deQ: {acquire (lock);
      if (head == null)
        wait (lock, notEmpty);
      item = head;
      if (tail == head) tail = null;
      head=item->next;
      release(lock);}
  
```

11

## Monitor Abstraction

```

EnQ: {acquire (lock);
      if (head == null)
        {head = item;
         signal (lock, notEmpty);}
      else tail->next = item;
      tail = item;
      release(lock);}
deQ: {acquire (lock);
      if (head == null)
        wait (lock, notEmpty);
      item = head;
      if (tail == head) tail = null;
      head=item->next;
      release(lock);}
  
```

12

## Monitor Abstraction

```

EnQ: {acquire (lock);
      if (head == null)
        {head = item;
         signal (lock, notEmpty);}
      else tail->next = item;
      tail = item;
      release(lock);}
deQ: {acquire (lock);
      if (head == null)
        wait (lock, notEmpty);
      item = head;
      if (tail == head) tail = null;
      head=item->next;
      release(lock);}
  
```

13

## Monitor Abstraction

```

EnQ: {acquire (lock);
      if (head == null)
        {head = item;
         signal (lock, notEmpty);}
      else tail->next = item;
      tail = item;
      release(lock);}
deQ: {acquire (lock);
      while (head == null)
        wait (lock, notEmpty);
      item = head;
      if (tail == head) tail = null;
      head=item->next;
      release(lock);}
  
```

14

## Nachos-style Synchronization

synch.h, cc

- Semaphores
  - Semaphore::P
  - Semaphore::V
- Locks and condition variables
  - Lock::Acquire
  - Lock::Release
  - Condition::Wait (conditionLock)
  - Condition::Signal (conditionLock)
  - Condition::Broadcast (conditionLock)

15

## Tweedledum and Tweedledee

- Separate threads executing their respective procedures. The code below is intended to cause them to forever take turns exchanging insults through the shared variable X in strict alternation.
- The Sleep() and Wakeup() routines operate as follows:
  - Sleep blocks the calling thread,
  - Wakeup unblocks a specific thread if that thread is blocked, otherwise its behavior is unpredictable

16

a) The code shown above exhibits a well-known synchronization flaw. Outline a scenario in which this code would fail, and the outcome of that scenario

```

void Tweedledum()                void Tweedledee()
{
    while(1) {
        Sleep();
        x = Quarrel(x);
        Wakeup(Tweedledee);
    }
}

{
    while(1) {
        x = Quarrel(x);
        Wakeup(Tweedledum);
        Sleep();
    }
}
    
```

If dee goes first to sleep, the wakeup is lost (since dum isn't sleeping yet). Both sleep forever.

17

b) Show how to fix the problem by replacing the Sleep and Wakeup calls with semaphore P (down) and V (up) operations.

```

void Tweedledum()                void Tweedledee()
{
    while(1) {
        P(dum);
        Sleep();
        x = Quarrel(x);
        Wakeup(Tweedledee);
    }
}

{
    while(1) {
        x = Quarrel(x);
        Wakeup(Tweedledum);
        Sleep();
        V(dum);
        P(dee);
    }
}
    
```

semaphore dee = 0;  
semaphore dum = 0;

18

## Classic Problems

There are a number of "classic" problems that represent a class of synchronization situations

- Critical Section problem
- Producer/Consumer problem
- Reader/Writer problem
- 5 Dining Philosophers

19

## The Critical Section Problem

```
while (1)
```

```
{ ...other stuff...
```

```
    [ ]
```

```
    critical section
```

```
    [ ]
```

```
}
```

**DONE**

20

## Producer / Consumer

Producer:

```
while(whatever)
{
  locally generate item
  [ ]
  fill empty buffer with item
  [ ]
}
```

Consumer:

```
while(whatever)
{
  [ ]
  get item from full buffer
  [ ]
  use item
}
```

21

## Producer / Consumer

Producer:

```
while(whatever)
{
  locally generate item
  [ ]
  P(emptybuf);
  [ ]
  fill empty buffer with item
  [ ]
  V(fullbuf);
  [ ]
}
```

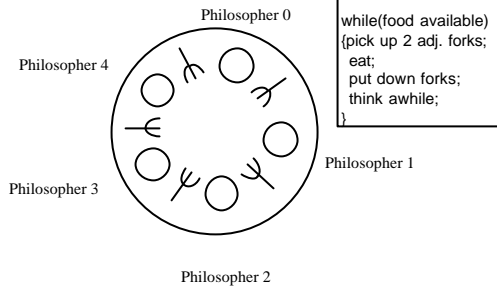
Consumer:

```
while(whatever)
{
  [ ]
  P(fullbuf);
  [ ]
  get item from full buffer
  [ ]
  V(emptybuf);
  [ ]
  use item
}
```

Semaphores: emptybuf initially N; fullbuf initially 0;

22

## 5 Dining Philosophers



23

## Template for Philosopher

```

while (food available)
{
     /*pick up forks*/
    eat;
     /*put down forks*/
    think awhile;
}
    
```

24

## Naive Solution

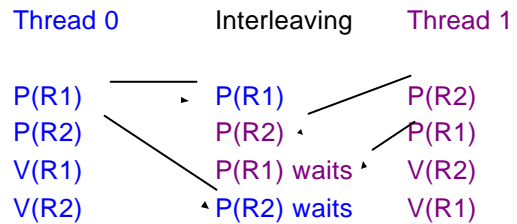
```

while (food available)
{
     /*pick up forks*/
    P(fork[left(me)]);
    P(fork[right(me)]);
    eat;
     /*put down forks*/
    V(fork[left(me)]);
    V(fork[right(me)]);
    think awhile;
}
    
```

Does this work?

25

## Simplest Example of Deadlock



R1 and R2 initially 1 (binary semaphore)

26

## Conditions for Deadlock

- Mutually exclusive use of resources
  - Binary semaphores R1 and R2
- Circular waiting
  - Thread 0 waits for Thread 1 to V(R2) and Thread 1 waits for Thread 0 to V(R1)
- Hold and wait
  - Holding either R1 or R2 while waiting on other
- No pre-emption
  - Neither R1 nor R2 are removed from their respective holding Threads.

27

## Philosophy 101 (or why 5DP is interesting)

- How to eat with your Fellows without causing **Deadlock**.
  - Circular arguments (the circular wait condition)
  - Not giving up on firmly held things (no preemption)
  - Infinite patience with Half-baked schemes (hold some & wait for more)
- Why **Starvation** exists and what we can do about it.

28

## Dealing with Deadlock

It can be **prevented** by breaking one of the prerequisite conditions:

- Mutually exclusive use of resources
  - Example: Allowing shared access to read-only files (readers/writers problem)
- circular waiting
  - Example: Define an **ordering** on resources and acquire them in order
- hold and wait
- no pre-emption

29

## Circular Wait Condition

```
while (food available)
{
  if (me == 0) {P(fork[left(me)]); P(fork[right(me)]);}
  else {(P(fork[right(me)]); P(fork[left(me)]); }
  eat;
  V(fork[left(me)]); V(fork[right(me)]);
  think awhile;
}
```

30

## Hold and Wait Condition

```

while (food available)
{
    P(mutex);
    while (forks [me] != 2)
        {blocking[me] = true; V(mutex); P(sleepy[me]); P(mutex);}
    forks [leftneighbor(me)] --; forks [rightneighbor(me)]--;
    V(mutex);
    eat;
    P(mutex); forks [leftneighbor(me)] ++; forks [rightneighbor(me)]++;
    if (blocking[leftneighbor(me)]) {blocking [leftneighbor(me)] = false;
    V(sleepy[leftneighbor(me)]};
    if (blocking[rightneighbor(me)]) {blocking[rightneighbor(me)] = false;
    V(sleepy[rightneighbor(me)]}; V(mutex);
    think awhile;
}
    
```

31

## Starvation

The difference between deadlock and starvation is subtle:

- Once a set of processes are deadlocked, there is no future execution sequence that can get them out of it.
- In starvation, there does exist some execution sequence that is favorable to the starving process although there is no guarantee it will ever occur.
- Rollback and Retry solutions are prone to starvation.
- Continuous arrival of higher priority processes is another common starvation situation.

32

## 5DP - Monitor Style

```

Boolean eating [5];
Lock forkMutex;
Condition forksAvail;

void PickupForks (int i) {
    forkMutex.Acquire( );
    while ( eating[(i-1)%5]
    || eating[(i+1)%5] )
        forksAvail.Wait(&forkMutex);
    eating[i] = true;
    forkMutex.Release( );
}

void PutdownForks (int i) {
    forkMutex.Acquire( );
    eating[i] = false;
    forksAvail.Broadcast(&forkMute
    x);
    forkMutex.Release( );
}
    
```

33

## What about this?

```

while (food available)
{
    forkMutex.Acquire( );
    while (forks [me] != 2) {blocking[me]=true;
    forkMutex.Release( ); sleep( ); forkMutex.Acquire( );}
    forks [leftneighbor(me)]--; forks [rightneighbor(me)]--;
    forkMutex.Release( );
    eat;
    forkMutex.Acquire( );
    forks[leftneighbor(me)] ++; forks [rightneighbor(me)]++;
    if (blocking[leftneighbor(me)] || blocking[rightneighbor(me)])
        wakeup( ); forkMutex.Release( );
    think awhile;
}
    
```

34

## Readers/Writers Problem

Synchronizing access to a file or data record in a database such that any number of threads requesting read-only access are allowed but only one thread requesting write access is allowed, excluding all readers.

35

## Template for Readers/Writers

```
Reader()                               Writer()
{while (true)                            {while (true)
{
    /*request r access*/                 /*request w access*/
    read                                write
    /*release r access*/                 /*release w access*/
}
}
```

36

## Template for Readers/Writers

```
Reader()                               Writer()
{while (true)                            {while (true)
{
    fd = open(foo, 0);                   fd = open(foo, 1);
    read                                write
    close(fd);                           close(fd);
}
}
```

37

## Template for Readers/Writers

```
Reader()                               Writer()
{while (true)                            {while (true)
{
    startRead();                          startWrite();
    read                                write
    endRead();                            endWrite();
}
}
```

38

### R/W - Monitor Style

```

Boolean busy = false;
int numReaders = 0;
Lock filesMutex;
Condition OKtoWrite, OKtoRead;

void startRead () {
    filesMutex.Acquire( );
    while ( busy )
        OKtoRead.Wait(&filesMutex);
    numReaders ++;
    filesMutex.Release( );
}

void endRead () {
    filesMutex.Acquire( );
    numReaders --;
    if (numReaders == 0)
        OKtoWrite.Signal(&filesMutex);
    filesMutex.Release( );
}

void startWrite() {
    filesMutex.Acquire( );
    while (busy || numReaders != 0)
        OKtoWrite.Wait(&filesMutex);
    busy = true;
    filesMutex.Release( );
}

void endWrite() {
    filesMutex.Acquire( );
    busy = false;
    OKtoRead.Broadcast(&filesMutex);
    OKtoWrite.Signal(&filesMutex);
    filesMutex.Release( );
}

```

39

### Semaphore Solution with Writer Priority

```

int readCount = 0, writeCount = 0;
semaphore mutex1 = 1, mutex2 = 1;
semaphore readBlock = 1;
semaphore writePending = 1;
semaphore writeBlock = 1;

```

40

```

Reader(){
    while (TRUE) {
        other stuff;
        P(writePending);
        P(readBlock);
        P(mutex1);
        readCount = readCount + 1;
        if (readCount == 1)
            P(writeBlock);
        V(mutex1); V(readBlock);
        V(writePending);
        access resource;
        P(mutex1);
        readCount = readCount - 1;
        if (readCount == 0)
            V(writeBlock);
        V(mutex1);
    }
}

Writer(){
    while(TRUE){
        other stuff;
        P(mutex2);
        writeCount = writeCount + 1;
        if ( writeCount == 1)
            P(readBlock);
        V(mutex2);
        P(writeBlock);
        access resource;
        V(writeBlock);
        P(mutex2);
        writeCount - writeCount - 1;
        if ( writeCount == 0)
            V(readBlock);
        V(mutex2);
    }
}

```

41

```

Reader(){
    while (TRUE) {
        other stuff;
        P(writePending);
        P(readBlock);
        P(mutex1);
        readCount = readCount + 1;
        if (readCount == 1)
            P(writeBlock);
        V(mutex1); V(readBlock);
        V(writePending);
        access resource;
        P(mutex1);
        readCount = readCount - 1;
        if (readCount == 0)
            V(writeBlock);
        V(mutex1);
    }
}

Writer(){
    while(TRUE){
        other stuff;
        P(mutex2);
        writeCount = writeCount + 1;
        if ( writeCount == 1)
            P(readBlock);
        V(mutex2);
        P(writeBlock);
        access resource;
        V(writeBlock);
        P(mutex2);
        writeCount - writeCount - 1;
        if ( writeCount == 0)
            V(readBlock);
        V(mutex2);
    }
}

```

42

Assume the writePending semaphore was omitted. What would happen?

```

Reader(){
while (TRUE) {
  other stuff;
  P(writePending);
  P(readBlock);
  P(mutex1);
  readCount = readCount + 1;
  if(readCount == 1)
    P(writeBlock);
  V(mutex1); V(readBlock);
  V(writePending);
  access resource;
  P(mutex1);
  readCount = readCount - 1;
  if(readCount == 0)
    V(writeBlock);
  V(mutex1);
}}

Writer(){
while(TRUE){
  other stuff;
  P(mutex2);
  writeCount = writeCount + 1;
  if ( writeCount == 1)
    P(readBlock);
  V(mutex2);
  P(writeBlock);
  access resource;
  V(writeBlock);
  P(mutex2);
  writeCount = writeCount - 1;
  if ( writeCount == 0)
    V(readBlock);
  V(mutex2);
}}

```

Assume the writePending semaphore was omitted. What would happen?

```

Reader(){
while (TRUE) {
  other stuff;
  P(writePending);
  P(readBlock);
  P(mutex1);
  readCount = readCount + 1;
  if(readCount == 1)
    P(writeBlock);
  V(mutex1); V(readBlock);
  V(writePending);
  access resource;
  P(mutex1);
  readCount = readCount - 1;
  if(readCount == 0)
    V(writeBlock);
  V(mutex1);
}}

Writer(){
while(TRUE){
  other stuff;
  P(mutex2);
  writeCount = writeCount + 1;
  if ( writeCount == 1)
    P(readBlock);
  V(mutex2);
  P(writeBlock);
  access resource;
  V(writeBlock);
  P(mutex2);
  writeCount = writeCount - 1;
  if ( writeCount == 0)
    V(readBlock);
  V(mutex2);
}}

```

Assume the writePending semaphore was omitted. What would happen?

Assume the writePending semaphore was omitted in the solution just given. What would happen?

This is supposed to give writers priority. However, consider the following sequence:

- Reader 1 arrives, executes thro' P(readBlock);
- Reader 1 executes P(mutex1);
- Writer 1 arrives, waits at P(readBlock);
- Reader 2 arrives, waits at P(readBlock);
- Reader 1 executes V(mutex1); then V(readBlock);
- Reader 2 may now proceed...**wrong**

45