

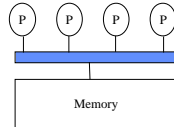
Outline Part 1

- Objectives:
 - To define the **process** and **thread abstractions**.
 - To briefly introduce **mechanisms** for implementing processes (threads).
 - To introduce the **critical section problem**.
 - To learn how to reason about the correctness of concurrent programs.
- Administrative details:
 - **Groups** are listed with “class pics”.
 - Pictures – make sure name-to-face mapping is correct.
 - Password protected
name: cps110 passwd: OntheGo

1

Concurrency

- Multiple things happening simultaneously
 - logically or physically
- Causes
 - Interrupts
 - Voluntary context switch (system call/trap)
 - Shared memory multiprocessor



HW Support for Atomic Operations

- Could provide direct support in HW
 - Atomic increment
 - *Insert node into sorted list??*
- Just provide low level primitives to construct atomic sequences
 - called **synchronization** primitives
 - ```
LOCK(counter->lock); // Wait here until unlocked
counter->value = counter->value + 1;
UNLOCK(counter->lock);
```
- test&set (x) instruction: returns previous value of x and sets x to “1”
  - ```
LOCK(x) => while (test&set(x));
UNLOCK(x) => x = 0;
```

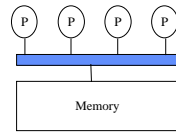
The Basics of Processes

- Processes are the *OS-provided abstraction* of multiple tasks (including user programs) executing concurrently.
- A Process IS: *one instance* of a *program* (which is only a passive set of bits) *executing* (implying an *execution context* – register state, memory resources, etc.)
- OS schedules processes to share CPU.

4

Why Use Processes?

- To capture naturally concurrent activities within the structure of the programmed system.
- To gain speedup by overlapping activities or exploiting parallel hardware.
 - From DMA to multiprocessors



5

Separation of Policy and Mechanism (System Design Principle)

- “Why and What” vs. “How”
- Objectives and strategies vs. data structures, hardware and software implementation issues.
- **Process abstraction vs. Process machinery**

Can you think of examples?

6

Process Abstraction

- Unit of scheduling
- One (or more*) sequential threads of control
 - program counter, register values, call stack
- Unit of resource allocation
 - address space (code and data), open files
 - sometimes called *tasks* or *jobs*
- Operations on processes: fork (clone-style creation), wait (parent on child), exit (self-termination), signal, kill.

Process-related System Calls in Unix.

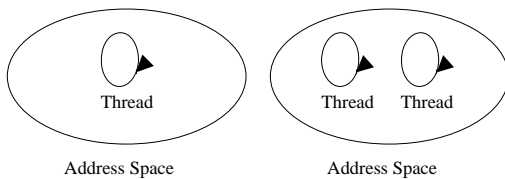
7

Threads and Processes

- Decouple the resource allocation aspect from the control aspect
- Thread abstraction - defines a single sequential instruction stream (PC, stack, register values)
- Process - the resource context serving as a “container” for one or more threads (shared address space)
- Kernel threads - unit of scheduling (kernel-supported thread operations – \rightarrow still slow)

8

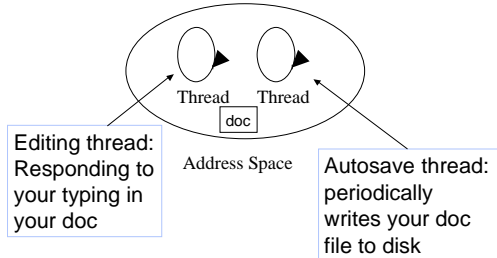
Threads and Processes



9

An Example

Doc formatting process



10

User-Level Threads

- To avoid the performance penalty of kernel-supported threads, implement at user level and manage by a run-time system
 - Contained "within" a single kernel entity (process)
 - Invisible to OS (OS schedules their container, not being aware of the threads themselves or their states). Poor scheduling decisions possible.
- User-level thread operations can be 100x faster than kernel thread operations, but need better integration / cooperation with OS.

11

Process Mechanisms

- PCB data structure in kernel memory represents a process (allocated on process creation, deallocated on termination).
- PCBs reside on various **state queues** (including a different queue for each "cause" of waiting) reflecting the process's state.
- As a process executes, the OS moves its PCB from queue to queue (e.g. from the "waiting on I/O" queue to the "ready to run" queue).

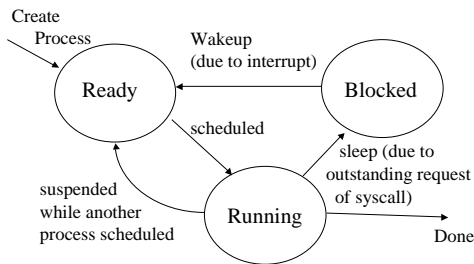
12

Context Switching

- When a process is running, its program counter, register values, stack pointer, etc. are contained in the hardware registers of the CPU. The process has direct control of the CPU hardware for now.
- When a process is not the one currently running, its current register values are saved in a process descriptor data structure (**PCB - process control block**)
- Context switching involves moving state between CPU and various processes' PCBs by the OS.

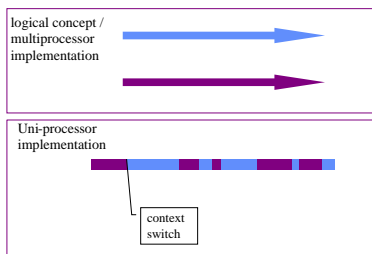
13

Process State Transitions



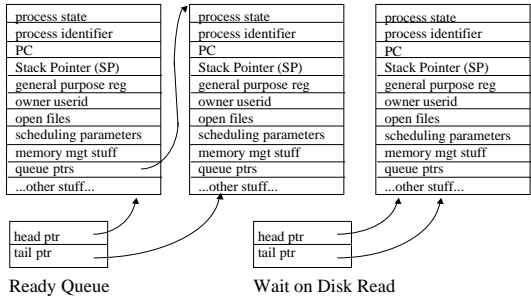
14

Interleaved Schedules

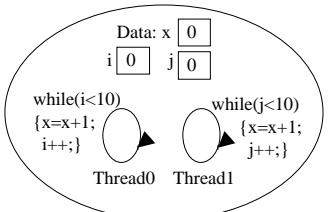


15

PCBs & Queues



The Trouble with Concurrency in Threads...



What is the value of x when both threads leave this while loop?

Nondeterminism

- What unit of work can be performed without interruption? **Indivisible** or **atomic** operations.
- **Interleavings** - possible execution sequences of operations drawn from all threads.
- **Race condition** - final results depend on ordering and may not be "correct".

```
while (i<10) {x=x+1; i++;}
    load value of x into reg
    yield()
    add 1 to reg
    yield()
    store reg value at x
    yield()
```

Interleaving

Thread 0

Load x (value 0)
Incr x (value 1 in reg)

Store x (value 1)

Load x (value 1) for 2nd iteration
Incr x (value 2 in reg)
...

Store x for 10th iteration (value 10)

Thread 1

Load x (value 0)
Incr x (value 1 in reg)
Store x (value 1)

...
Store x for 9th iteration (value 9)

Load x (value 1) for 10th iteration
Incr x (value 2 in reg)

Store x (value 2) for 10th iteration

19

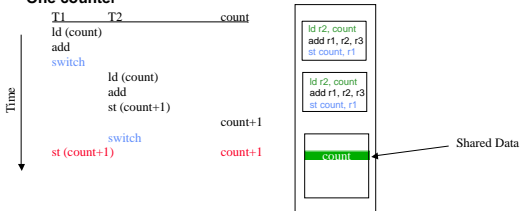
Reasoning about Interleavings

- On a uniprocessor, the possible execution sequences depend on when context switches can occur
 - Voluntary context switch - the process or thread explicitly yields the CPU (blocking on a system call it makes, invoking a Yield operation).
 - Interrupts or exceptions occurring - an asynchronous handler activated that disrupts the execution flow.
 - Preemptive scheduling - a timer interrupt may cause an involuntary context switch at any point in the code.
- On multiprocessors, the ordering of operations on shared memory locations is the important factor.

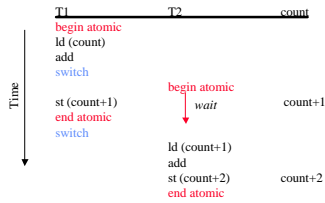
20

The Trouble with Concurrency

- Two threads (T1,T2) in one address space or two processes in the kernel
- One counter



Solution: Atomic Sequence of Instructions



- Atomic Sequence
 - Appears to execute to completion without any intervening operations

Critical Sections

- If a sequence of non-atomic operations must be executed as *if* it were atomic in order to be correct, then we need to provide a way to constrain the possible interleavings in this **critical section** of our code.
 - Critical sections are code sequences that contribute to “bad” race conditions.
 - Synchronization is needed around such critical sections.
- **Mutual Exclusion** - goal is to ensure that critical sections execute atomically w.r.t. related critical sections in other threads or processes.
 - How?

23

The Critical Section Problem

Each process follows this template:

```
while (1)
{ ...other stuff... //processes in here shouldn't stop others
  enter_region( );
  critical_section
  exit_region( );
}
```

The problem is to define enter_region and exit_region to ensure mutual exclusion with some degree of fairness.

24

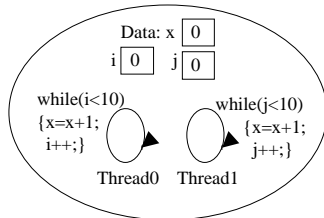
Implementation Options for Mutual Exclusion

- Disable Interrupts
- Busywaiting solutions - spinlocks
 - execute a tight loop if critical section is busy
 - benefits from specialized atomic (read-mod-write) instructions
- Blocking synchronization
 - sleep (enqueued on wait queue) while C.S. is busy

Synchronization primitives (abstractions, such as locks) which are provided by a system may be implemented with some combination of these techniques.

25

The Trouble with Concurrency in Threads...



What is the value of x when both threads leave this while loop?

26

Range of Answers

Process 0

LD x // x currently 0

Add 1

ST x // x now 1, stored over 9

Do 9 more full loops // leaving x at 10

Process1

LD x // x currently 0

Add 1

ST x // x now 1

Do 8 more full loops // x = 9

LD x // x now 1

Add 1

ST x // x = 2 stored over 10

27

The Critical Section Problem

```
while (1)
{ ...other stuff...
  
  critical section
  exit_region( );
  
}
```

28

Proposed Algorithm for 2 Process Mutual Exclusion

```
Boolean flag[2];
proc (int i) {
  while (TRUE){
    compute;
    flag[i] = TRUE ;
    while(flag[(i+1) mod 2]) ;
    critical section;
    flag[i] = FALSE;
  }
}
```

flag[0] = flag[1]= FALSE;
fork (proc, 1, 0);
fork (proc, 1,1);
Is it correct?
Assume they go lockstep.
Both set their own flag to TRUE.
Both busywait forever on the other's flag -> deadlock.

29

Proposed Algorithm for 2 Process Mutual Exclusion

- enter_region:
 needin [me] = true;
 turn = you;
 while (needin [you] && turn == you) {no_op};
- exit_region:
 needin [me] = false;

Is it correct?

30

Interleaving of Execution of 2 Threads (blue and green)

```
enter_region:          enter_region:
needin [me] = true;    needin [me] = true;
turn = you;           turn = you;
while (needin [you] &&  while (needin [you] &&
  turn == you) {no_op};   turn == you) {no_op};
Critical Section      Critical Section
exit_region:          exit_region:
needin [me] = false;   needin [me] = false;
```

31

```
needin [blue] = true;
needin [green] = true;
turn = green;
turn = blue;
while (needin [green] && turn == green)
Critical Section
while (needin [blue] && turn == blue){no_op};
while (needin [blue] && turn == blue){no_op};
needin [blue] = false;
while (needin [blue] && turn == blue)
Critical Section
needin [green] = false;
```

32

Greedy Version (turn = me)

```
needin [blue] = true;
needin [green] = true;
turn = blue;
while (needin [green] && turn == green)
Critical Section
turn = green;
while (needin [blue] && turn == blue)
Critical Section
Oooops!
```

33

Synchronization

- We illustrated the dangers of race conditions when multiple threads execute instructions that interfere with each other when interleaved.
- Goal in solving the critical section problem is to build synchronization so that the sequence of instructions that can cause a race condition are executed **AS IF** they were indivisible (just appearances)
- “Other stuff” can be interleaved with critical section code as well as the enter_region and exit_region protocols, but it is deemed OK.

34

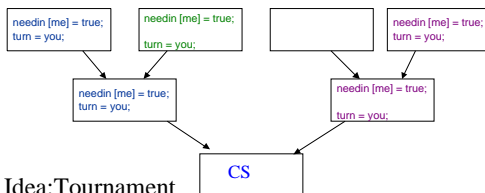
Peterson's Algorithm for 2 Process Mutual Exclusion

- enter_region:
 needin [me] = true;
 turn = you;
 while (needin [you] && turn == you) {no_op};
- exit_region:
 needin [me] = false;

What about more than 2 processes?

35

Can we extend 2-process algorithm to work with n processes?



Idea: Tournament
Details: Bookkeeping (left to the reader)

37

Lamport's Bakery Algorithm

- **enter_region:**
choosing[me] = true;
number[me] = max(number[0:n-1]) + 1;
choosing[me] = false;
for (j=0; j<n; j++) {
 { while (choosing[j] != 0) ;
 while((number[j] != 0) and ((number[j] < number[me])
 or ((number[j] == number[me]) and (j < me)))) ;
 }
}
- **exit_region:**
number[me] = 0;

38

Explanation of Lamport's Bakery Algorithm

```
choosing[me] = true;  
number[me] = max(number[0:n-1]) + 1;  
choosing[me] = false;  
/* choosing[i] is false when number[i] is not changing to non -zero */
```

```
for (j=0; j<n; j++) {  
  { while (choosing[j] != 0) ;  
    while((number[j] != 0) and ((number[j] < number[me])  
      or ((number[j] == number[me]) and (j < me)))) ;  
  }  
}
```

/* While thread i is in this for-loop, number[i] is non-zero; if thread j (j<i) arrives later, while i is examining number[], and sets choosing[j] false, number[j]>number[i].

Even if thread i has examined j already, i can enter CS(i), and j will NOT enter CS(j) until i leaves CS(i), and sets number[i] to 0. */

39

Interleaving / Execution Sequence with Bakery Algorithm

Thread 0

Choosing= False
Number [0]= 0

Thread 1

Choosing= False
Number [1]= 0

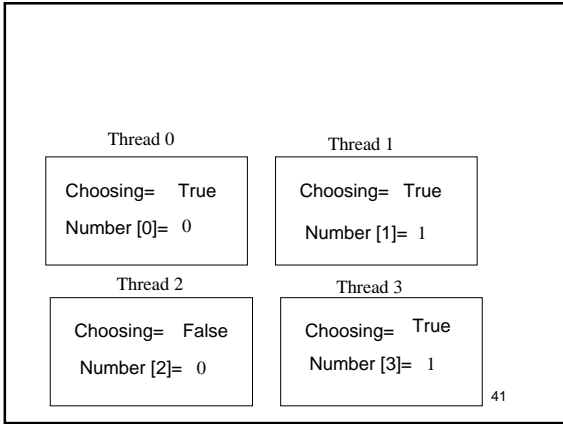
Thread 2

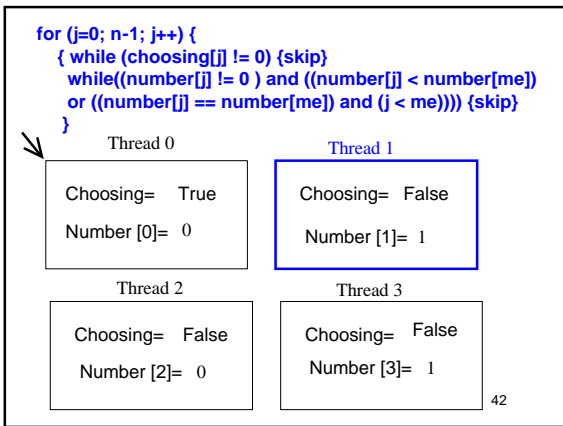
Choosing= False
Number [2]= 0

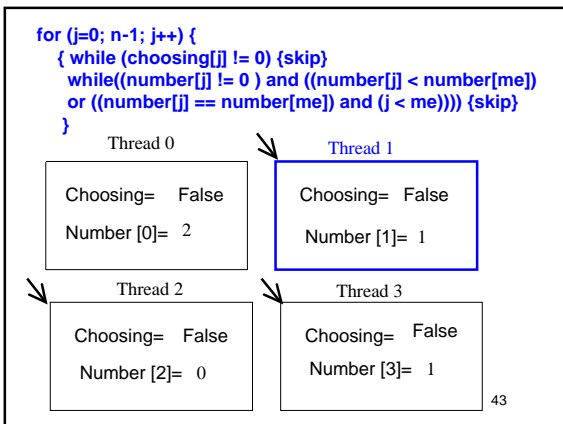
Thread 3

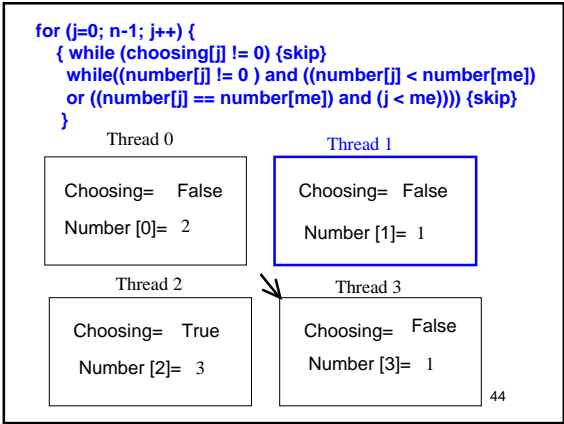
Choosing= False
Number [3]= 0

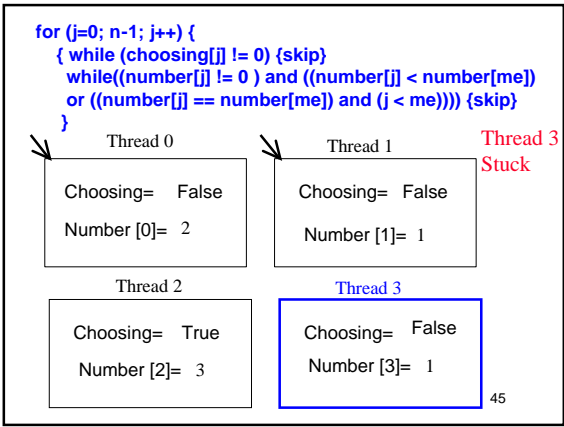
40

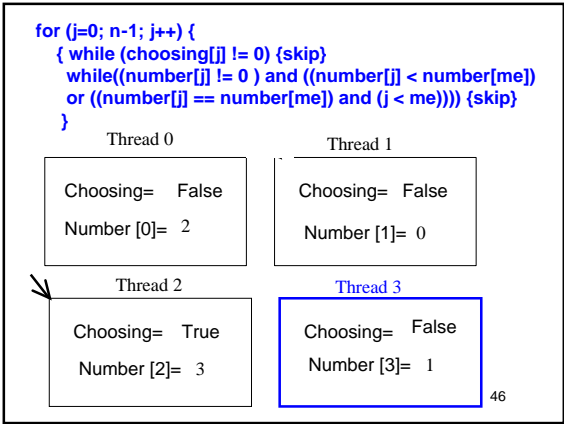








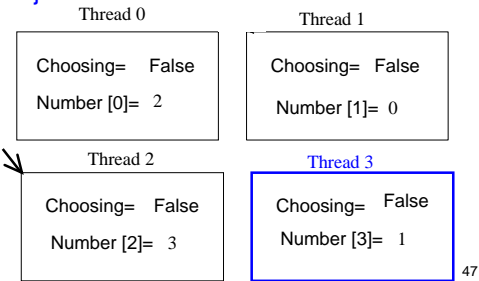




```

for (j=0; n-1; j++) {
  while (choosing[j] != 0) {skip}
  while((number[j] != 0 ) and ((number[j] < number[me])
  or ((number[j] == number[me]) and (j < me)))) {skip}
}

```



Hardware Assistance

- Most modern architectures provide some support for building synchronization: atomic **read-modify-write** instructions.
- Example: **test-and-set (loc, reg)**
 [sets bit to 1 in the new value of loc;
 returns old value of loc in reg]
- Other examples: *compare-and-swap, fetch-and-op*

[] notation means atomic

Busywaiting with Test-and-Set

- Declare a shared memory location to represent a *busyflag* on the critical section we are trying to protect.
- enter_region (or *acquiring* the “lock”):
 waitloop: tsl busyflag, R0 // R0 = busyflag; busyflag = 1
 bnz R0, waitloop // was it already set?
- exit region (or *releasing* the “lock”):
 busyflag = 0

Pros and Cons of Busywaiting

- Key characteristic - the “waiting” process is actively executing instructions in the CPU and using memory cycles.
- Appropriate when:
 - High likelihood of finding the critical section unoccupied (don't take context switch just to find that out) or estimated wait time is very short
- Disadvantages:
 - Wastes resources (CPU, memory, bus bandwidth)

50

Blocking Synchronization

- OS implementation involving changing the state of the “waiting” process from running to blocked.
- Need some synchronization abstraction known to OS - provided by system calls.
 - mutex locks with operations acquire and release
 - semaphores with operations P and V (down, up)
 - condition variables with wait and signal

51

Template for Implementing Blocking Synchronization

- Associated with the lock is a memory location (busy) and a queue for waiting threads/processes.
- *Acquire* syscall:

```
while (busy) {enqueue caller on lock's queue}
/*upon waking to nonbusy lock*/ busy = true;
```
- *Release* syscall:

```
busy = false;
/* wakeup */ move any waiting threads to Ready queue
```

52

Pros and Cons of Blocking

- Waiting processes/threads don't consume CPU cycles
- Appropriate: when the cost of a system call is justified by expected waiting time
 - High likelihood of contention for lock
 - Long critical sections
- Disadvantage: OS involvement
 overhead

53
