

COS 140: Foundations of Computer Science

Process Synchronization: Semaphores

Fall 2018

Process Synchronization	3
What is it?	3
Race conditions.	5
Example	6
Mutual exclusion.	7
Semaphores	9
One solution: Semaphores	9
Using semaphores	10
Example	12
Implementation	13
Data structure	13
Counting Semaphores	16
Counting semaphores	16
Example	17
Producer–Consumer Problem	18
Discussion	19
Problems with semaphores	19
Other approaches	20

Homework, etc.

- Reading: Chapter 20
- Homework: Exercises at end of Chapter 20
- Due: Friday, 11/9
- NOTE:** Remember that Prelim II is Wednesday, 11/14!

Copyright © 2002–2018 UMaine Computer Science Department – 2 / 20

Process Synchronization

3 / 20

Process synchronization

- Problems \Leftarrow processes share resource – e.g., memory or device
- Possible: both want to use the resource simultaneously and incompatibly

Copyright © 2002–2018 UMaine Computer Science Department – 3 / 20

Process synchronization

- Example: Two movers, Bob and Jill, are putting furniture into house
 - Both approach doorway (shared resource) simultaneously with a chair they can't see around
 - What happens? Bob goes first; Jill goes first; both get stuck!
- Example: Two tellers, two customers, Sue and Jim, with joint account
 - Joint account = shared resource
 - Both want to put \$100 in at same time, tellers update accounts
 - What can happen?
 - ▷ Sue or Jim go first → \$200 in account
 - ▷ Both at same time...? Maybe only \$100 in account!

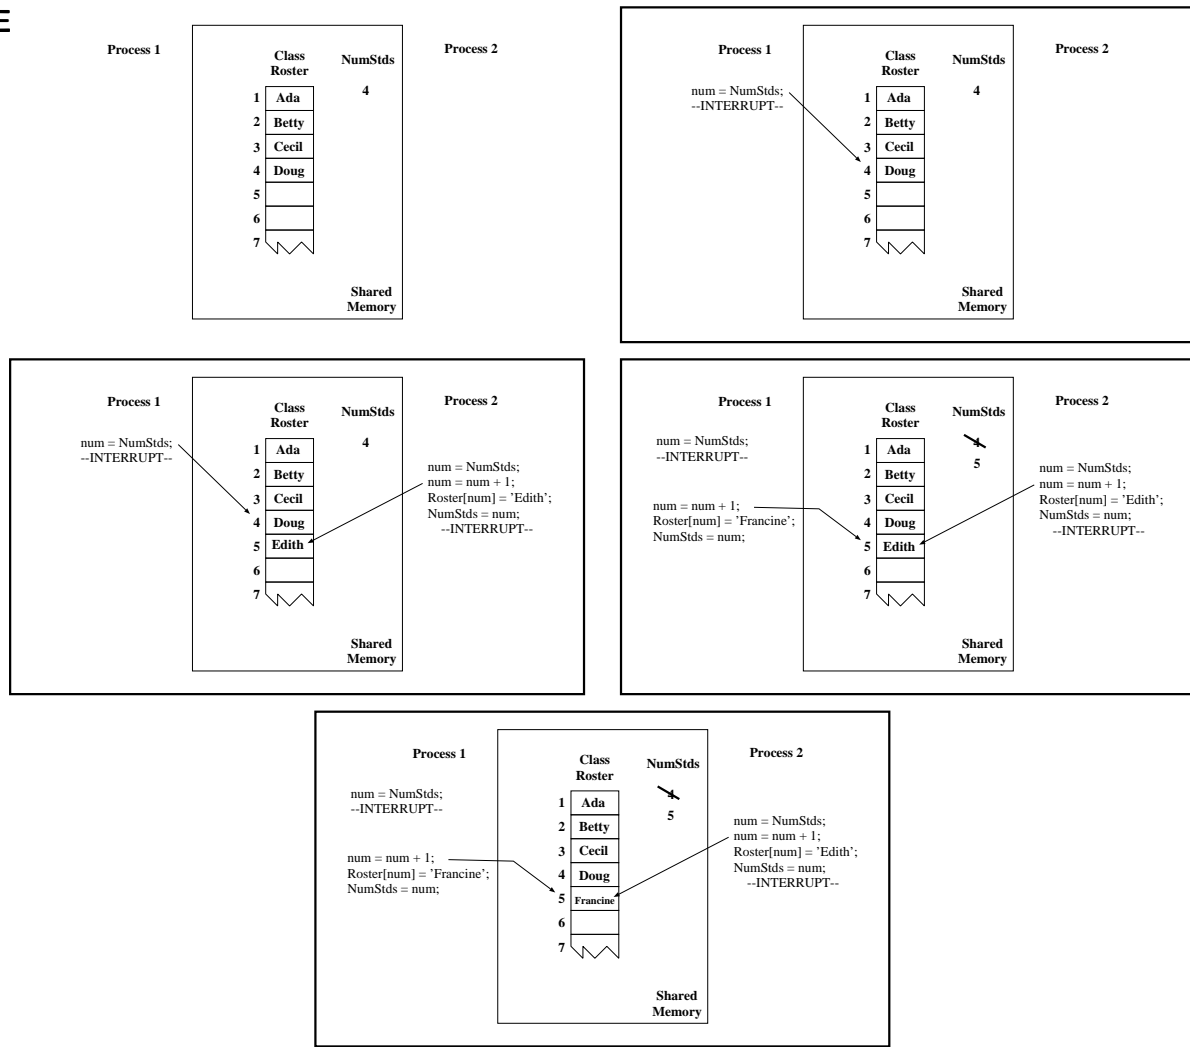
Copyright © 2002–2018 UMaine Computer Science Department – 4 / 20

Race conditions

- Examples of *race conditions*: outcome depends on timing/speed of participants
- Computer: processes are the “racers”
- Resource: any non-simultaneously-sharable thing

Copyright © 2002–2018 UMaine Computer Science Department – 5 / 20

E



Copyright © 2002–2018 UMaine Computer Science Department – 6 / 20

Approach: Mutual exclusion

- Identify *critical regions* where shared resource is being accessed – in teach process
- If we only allow one process into its critical region at a time → no conflict, no race condition

Copyright © 2002–2018 UMaine Computer Science Department – 7 / 20

Solution requirements

- Result is predictable
- Solution does not depend on speed of processes
- Solution does not depend on number of CPUs


Copyright © 2002–2018 UMaine Computer Science Department – 8 / 20

Semaphores

9 / 20

One solution: Semaphores


- Semaphores*:
 - *synchronization primitive*
 - guarantee mutual exclusion when used correctly
- Semaphore: data structure (and associated code) that:
 - Keeps track of whether or not the resource is in use...
 - ...or alternatively, keeps track of whether anyone is in their critical region
 - *Blocks* process if it tries to enter its critical region when someone else is in theirs
- Usually includes *queue* of processes that are blocked



Copyright © 2002–2018 UMaine Computer Science Department – 9 / 20

Using semaphores

- A semaphore S has two major procedures associated with it:
 - $P(S)$, or $\text{Down}(S)$
 - $V(S)$, or $\text{Up}(S)$
 - (V for Dutch *verhoog* (increase) and P for *prolaag* (try to decrease) [Dijkstra])
- Think of semaphore signaling “okay” when up, “stop!” when down




Copyright © 2002–2018 UMaine Computer Science Department – 10 / 20

Using semaphores

- Initially: semaphore is “up”
- Key: wrap calls to Down and Up around each process' critical region
 - Suppose we have semaphore S
 - Process wants to enter critical region, calls $\text{Down}(S)$
 - If semaphore “up” then semaphore \Rightarrow “down”, process continues
 - If semaphore “down” then process blocks
 - When process leaves critical region: call $\text{Up}(S)$

Copyright © 2002–2018 UMaine Computer Science Department – 11 / 20

Example



Copyright © 2002–2018 UMaine Computer Science Department – 12 / 20

Implementation

13 / 20

Data structure

- Semaphores are *data structures*
- Two parts: count and queue
 - Count is an integer
 - Queue contains process IDs
- This type of semaphore = *mutual exclusion*, or *mutex*, semaphore
- Count: set to 1 for mutex semaphores

Copyright © 2002–2018 UMaine Computer Science Department – 13 / 20

Procedures

- Can implement as separate procedures or *methods* of a semaphore class
- Down() and Up() have to be *atomic*
- Only OS can do this!
- Usually have to ask OS for semaphores, then Down() and Up() are accessed via system calls

Copyright © 2002–2018 UMaine Computer Science Department – 14 / 20

Procedures

- Down(mutex):
 - Decrement count
 - If < 0 , then block current process
 - Otherwise, continue (into critical region)
- Up(mutex):
 - Increment count
 - If there are blocked processes, allow one to continue
 - Note: semaphore is still “down” until count is positive
- Example

Copyright © 2002–2018 UMaine Computer Science Department – 15 / 20

Counting semaphores

- Mutex semaphores: just one kind
- Semaphores really *count* number of free resources
- Mutex: only 1 unit of resource = critical region
- Other cases: may have $\geq 1 \Rightarrow$ *counting semaphores*

Copyright © 2002–2018 UMaine Computer Science Department – 16 / 20

Example

- Producer–consumer problem
- Two processes, shared resource of finite size
- Producer puts things into the resource
- Consumer removes them
- Producer must block when full, consumer must block when empty

Copyright © 2002–2018 UMaine Computer Science Department – 17 / 20

Producer–Consumer Problem

```
Semaphores: mutex, full, empty;  
Set count of mutex=1, full=0,  
empty=size of resource;
```

```
Producer:                Consumer:  
loop forever:           loop forever:  
  Get thing to put in;   Down(full);  
  Down(empty);           Down(mutex);  
  Down(mutex);           Take thing out;  
  Put thing in resource; Up(mutex);  
  Up(mutex);             Up(empty);  
  Up(full);              Use thing;  
end loop.                end loop.
```

Copyright © 2002–2018 UMaine Computer Science Department – 18 / 20

Discussion

19 / 20

Problems with semaphores

- Easy to make mistakes:
 - Forget to do Up(S)
 - Too many Down's
 - Crossed semaphores
- Cheating
- Too low-level

Copyright © 2002–2018 UMaine Computer Science Department – 19 / 20

Other approaches

- Event counters
- Monitors
- Message passing
- Equivalence of primitives

Copyright © 2002–2018 UMaine Computer Science Department – 20 / 20