

# Last Class: CPU Scheduling

- **Scheduling Algorithms:**

- FCFS
- Round Robin
- SJF
- Multilevel Feedback Queues
- Lottery Scheduling

- **Review questions:**

- How does each work?
- Advantages? Disadvantages?



## Multilevel Feedback Queues (MLFQ)

- Multilevel feedback queues use past behavior to predict the future and assign job priorities  
=> overcome the prediction problem in SJF
- If a process is I/O bound in the past, it is also likely to be I/O bound in the future (programs turn out not to be random.)
- To exploit this behavior, the scheduler can favor jobs that have used the least amount of CPU time, thus approximating SJF.
- This policy is **adaptive** because it relies on past behavior and changes in behavior result in changes to scheduling decisions.



# Approximating SJF: Multilevel Feedback Queues

- Multiple queues with different priorities.
- Use Round Robin scheduling at each priority level, running the jobs in highest priority queue first.
- Once those finish, run jobs at the next highest priority queue, etc. (Can lead to starvation.)
- Round robin time slice increases exponentially at lower priorities.

	Priority	Time Slice			
<table border="1"><tr><td>G</td><td>F</td><td>A</td></tr></table>	G	F	A	1	1
G	F	A			
<table border="1"><tr><td></td><td></td><td>E</td></tr></table>			E	2	2
		E			
<table border="1"><tr><td></td><td>D</td><td>B</td></tr></table>		D	B	3	4
	D	B			
<table border="1"><tr><td></td><td></td><td>C</td></tr></table>			C	4	8
		C			



## Adjusting Priorities in MLFQ

- Job starts in highest priority queue.
  - If job's time slices expires, drop its priority one level.
  - If job's time slices does not expire (the context switch comes from an I/O request instead), then increase its priority one level, up to the top priority level.
- ⇒ CPU bound jobs drop like a rock in priority and I/O bound jobs stay at a high priority.



# Multilevel Feedback Queues: Example 1

- 3 jobs, of length 30, 20, and 10 seconds each, initial time slice 1 second, context switch time of 0 seconds, all CPU bound (no I/O), 3 queues

Queue	Time Slice	Job
1	1	
2	2	
3	4	

Job	Length	Completion Time		Wait Time	
		RR	MLFQ	RR	MLFQ
1	30				
2	20				
3	10				
Average					



# Multilevel Feedback Queues: Example 1

- 5 jobs, of length 30, 20, and 10 seconds each, initial time slice 1 second, context switch time of 0 seconds, all CPU bound (no I/O), 3 queues

Queue	Time Slice	Job
1	1	$1_1^1, 2_2^1, 3_3^1$
2	2	$1_5^3, 2_7^3, 3_9^3$
3	4	$1_{13}^7, 2_{17}^7, 3_{21}^7$ $1_{25}^{11}, 2_{29}^{11}, 3_{32}^{10} \dots$

Job	Length	Completion Time		Wait Time	
		RR	MLFQ	RR	MLFQ
1	30	60	60	30	30
2	20	50	53	30	33
3	10	30	32	20	22
Average		46 2/3	48 1/3	26 2/3	28 1/3



# Multilevel Feedback Queues: Example 2

• 3 jobs, of length 30, 20, and 10 seconds, the 10 sec job has 1 sec of I/O every other sec, initial time slice 2 sec, context switch time of 0 sec, 2 queues.

Queue	Time Slice	Job
1	1	
2	2	

Job	Length	Completion Time		Wait Time	
		RR	MLFQ	RR	MLFQ
1	30				
2	20				
3	10				
Average					



# Multilevel Feedback Queues: Example 2

• 3 jobs, of length 30, 20, and 10 seconds, the 10 sec job has 1 sec of I/O every other sec, initial time slice 1 sec, context switch time of 0 sec, 2 queues.

Job	Length	Completion Time		Wait Time	
		RR	MLFQ	RR	MLFQ
1	30	60	60	30	30
2	20	50	50	30	30
3	10	30	18	20	8
Average		46 2/3	45	26 2/3	25 1/3

Queue	Time Slice	Job
1	1	1 <sup>1</sup> , 2 <sup>1</sup> , 3 <sup>1</sup> 3 <sup>3</sup> , 3 <sup>5</sup> , 3 <sup>7</sup> , 3 <sup>9</sup> , 3 <sup>11</sup> , 3 <sup>13</sup> 3 <sup>15</sup> , 3 <sup>17</sup>
2	2	1 <sup>3</sup> , 2 <sup>3</sup> , 1 <sup>5</sup> , 2 <sup>5</sup> , 1 <sup>7</sup> , 2 <sup>7</sup> 1 <sup>9</sup> , 2 <sup>9</sup> , 1 <sup>11</sup> , 2 <sup>11</sup> , 1 <sup>13</sup> , 2 <sup>13</sup> 1 <sup>15</sup> , 2 <sup>15</sup> , 1 <sup>17</sup> , 2 <sup>17</sup> , 1 <sup>19</sup> , 2 <sup>19</sup> 1 <sup>21</sup> , 2 <sup>21</sup> , 1 <sup>23</sup> , 2 <sup>23</sup> , 1 <sup>25</sup> , 2 <sup>25</sup> 1 <sup>27</sup> , 2 <sup>27</sup> , 1 <sup>29</sup> , 2 <sup>29</sup>



# Improving Fairness

Since SJF is optimal, but unfair, any increase in fairness by giving long jobs a fraction of the CPU when shorter jobs are available will degrade average waiting time.

Possible solutions:

- Give each queue a fraction of the CPU time. This solution is only fair if there is an even distribution of jobs among queues.
- Adjust the priority of jobs as they do not get serviced (Unix originally did this.) This ad hoc solution avoids starvation but average waiting time suffers when the system is overloaded because all the jobs end up with a high priority,.



# Lottery Scheduling

- Give every job some number of lottery tickets.
- On each time slice, randomly pick a winning ticket.
- On average, CPU time is proportional to the number of tickets given to each job.
- Assign tickets by giving the most to short running jobs, and fewer to long running jobs (approximating SJF). To avoid starvation, every job gets at least one ticket.
- Degrades gracefully as load changes. Adding or deleting a job affects all jobs proportionately, independent of the number of tickets a job has.



# Lottery Scheduling: Example

- Short jobs get 10 tickets, long jobs get 1 ticket each.

# short jobs/ # long jobs	% of CPU each short job gets	% of CPU each long job gets
1/1	91%	9%
0/2		
2/0		
10/1		
1/10		



# Lottery Scheduling Example

- Short jobs get 10 tickets, long jobs get 1 ticket each.

# short jobs/ # long jobs	% of CPU each short job gets	% of CPU each long job gets
1/1	91% (10/11)	9% (1/11)
0/2		50% (1/2)
2/0	50% (10/20)	
10/1	10% (10/101)	< 1% (1/101)
1/10	50% (10/20)	5% (1/20)



# Today: Synchronization

•What kind of knowledge and mechanisms do we need to get independent processes to communicate and get a consistent view of the world (computer state)?

•**Example: Too Much Milk**

<i>Time</i>	You	Your roommate
3:00	Arrive home	
3:05	Look in fridge, no milk	
3:10	Leave for grocery	
3:15		Arrive home
3:20	Arrive at grocery	Look in fridge, no milk
3:25	Buy milk	Leave for grocery
3:35	Arrive home, put milk in fridge	
3:45		Buy milk
3:50		Arrive home, put up mlk
3:50		Oh no!



## Synchronization Terminology

- **Synchronization:** use of atomic operations to ensure cooperation between threads
- **Mutual Exclusion:** ensure that only one thread does a particular activity at a time and *excludes* other threads from doing it at that time
- **Critical Section:** piece of code that only one thread can execute at a time
- **Lock:** mechanism to prevent another process from doing something
  - Lock before entering a critical section, or before accessing shared data.
  - Unlock when leaving a critical section or when access to shared data is complete
  - Wait if locked

=> All synchronization involves waiting.



# Too Much Milk: Solution 1

- What are the correctness properties for this problem?
  - Only one person buys milk at a time.
  - Someone buys milk if you need it.
- Restrict ourselves to atomic loads and stores as building blocks.
  - Leave a note (a version of lock)
  - Remove note (a version of unlock)
  - Do not buy any milk if there is note (wait)

Thread A

```
if (noMilk & NoNote) {  
    leave Note;  
    buy milk;  
    remove note;  
}
```

Does this work?

Thread B

```
if (noMilk & NoNote) {  
    leave Note;  
    buy milk;  
    remove note;  
}
```



# Too Much Milk: Solution 2

How about using labeled notes so we can leave a note before checking the the milk?

Thread A

```
leave note A  
if (noNote B) {  
    if (noMilk){  
        buy milk;  
    }  
}  
remove note;
```

Thread B

```
leave note B  
if (noNote A) {  
    if (noMilk){  
        buy milk;  
    }  
}  
remove note;
```

Does this work?





# Too Much Milk: Solution 3

Thread A	Thread B
leave note A	leave note B
X: while (Note B) {	Y: if (noNote A) {
do nothing;	if (noMilk){
}	buy milk;
if (noMilk){	}
buy milk;	}
}	remove note B;
remove note A;	

Does this work?



## Correctness of Solution 3

- At point Y, either there is a note A or not.
  1. If there is no note A, it is safe for thread B to check and buy milk, if needed. (Thread A has not started yet).
  2. If there is a note A, then thread A is checking and buying milk as needed or is waiting for B to quit, so B quits by removing note B.
- At point X, either there is a note B or not.
  1. If there is not a note B, it is safe for A to buy since B has either not started or quit.
  2. If there is a note B, A waits until there is no longer a note B, and either finds milk that B bought or buys it if needed.
- Thus, thread B buys milk (which thread A finds) or not, but either way it removes note B. Since thread A loops, it waits for B to buy milk or not, and then if B did not buy, it buys the milk.



# Is Solution 3 a good solution?

- It is too complicated - it was hard to convince ourselves this solution works.
- It is asymmetrical - thread A and B are different. Thus, adding more threads would require different code for each new thread and modifications to existing threads.
- A is *busy waiting* - A is consuming CPU resources despite the fact that it is not doing any useful work.

=> This solution relies on loads and stores being atomic.



## Language Support for Synchronization

Have your programming language provide atomic routines for synchronization.

- **Locks:** one process holds a lock at a time, does its critical section releases lock.
- **Semaphores:** more general version of locks.
- **Monitors:** connects shared data to synchronization primitives.

=> All of these require some hardware support, and waiting.



# Locks

- **Locks:** provide mutual exclusion to shared data with two “atomic” routines:
  - **Lock.Acquire** - wait until lock is free, then grab it.
  - **Lock.Release** - unlock, and wake up any thread waiting in Acquire.

Rules for using a lock:

- Always acquire the lock before accessing shared data.
- Always release the lock after finishing with shared data.
- Lock is initially free.



## Implementing Too Much Milk with Locks

### Too Much Milk

Thread A	Thread B
<pre>Lock.Acquire(); if (noMilk){     buy milk; } Lock.Release();</pre>	<pre>Lock.Acquire(); if (noMilk){     buy milk; } Lock.Release();</pre>

- This solution is clean and symmetric.
- How do we make Lock.Acquire and Lock.Release atomic?



# Hardware Support for Synchronization

- Implementing high level primitives requires low-level hardware support
- What we have and what we want

	Concurrent programs
Low-level atomic operations (hardware)	load/store    interrupt disable    test&set
High-level atomic operations (software)	lock            semaphore monitors      send & receive



## Implementing Locks By Disabling Interrupts

- There are two ways the CPU scheduler gets control:
  - **Internal Events:** the thread does something to relinquish control (e.g., I/O).
  - **External Events:** interrupts (e.g., time slice) cause the scheduler to take control away from the running thread.
- On uniprocessors, we can prevent the scheduler from getting control as follows:
  - **Internal Events:** prevent these by not requesting any I/O operations during a critical section.
  - **External Events:** prevent these by disabling interrupts (i.e., tell the hardware to delay handling any external events until after the thread is finished with the critical section)
- Why not have the OS support `Lock::Acquire()` and `Lock::Release` as system calls?



# Implementing Locks by Disabling Interrupts

- For uniprocessors, we can disable interrupts for high-level primitives like locks, whose implementations are private to the kernel.
- The kernel ensures that interrupts are not disabled forever, just like it already does during interrupt handling.

```
class Lock {
public:
    void Acquire();
    void Release();
private:
    int value;
    Queue Q;
}
Lock::Lock {
    // lock is free
    value = 0;
    // queue is empty
    Q = 0;
}
Lock::Acquire(T:Thread){
    disable interrupts;
    if (value == BUSY) {
        add T to Q
        T->Sleep();
    } else {
        value = BUSY;
    }
}
Lock::Release() {
    disable interrupts;
    if queue not empty {
        take thread T off Q
        put T on ready queue
    } else {
        value = FREE
    }
}
enable interrupts; }
```



## Wait Queues

When should Acquire re-enable interrupts when going to sleep?

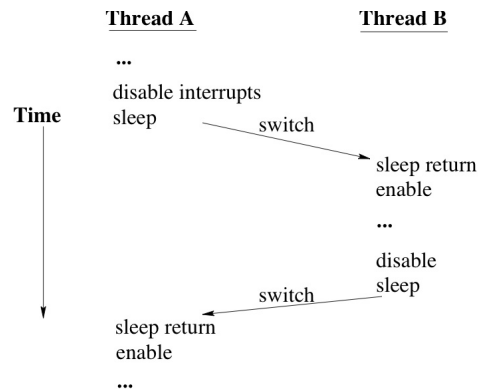
- **Before putting the thread on the wait queue?**
  - No, Release could check the queue, and not wake up the thread.
- **After putting the thread on the wait queue, but before going to sleep?**
  - No, Release could put the thread on the ready queue, but it could already be on the ready queue. When the thread wakes up, it will go to sleep, missing the wakeup from Release.

=>We still have a problem with multiprocessors.



# Example

- When the sleeping thread wakes up, it returns from Sleep back to Acquire.
- Interrupts are still disabled, so its ok to check the lock value, and if it is free, grab the lock and turn on interrupts.



# Summary

- Communication among threads is typically done through shared variables.
- Critical sections identify pieces of code that cannot be executed in parallel by multiple threads, typically code that accesses and/or modifies the values of shared variables.
- Synchronization primitives are required to ensure that only one thread executes in a critical section at a time.
  - Achieving synchronization directly with loads and stores is tricky and error-prone
  - *Solution*: use high-level primitives such as locks, semaphores, monitors

