





Assume new processes go on the HEAD of the ready queue:

	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	2	2	2	2	2	2	2	2	2	2	2
	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	
P1	X	X	X	X	X	X	X	X	X	X	R	R	R	R	R	R	R	R	R	R	X	X	X	X	X	X	X	X	X	X	R
P2		R	R	R	R	R	R	R	R	R	X	X	X	X	X	X	X	X	X	X											
P3																						R	R	R	R	R	R	R	R	R	X
P4																															

	3	3	3	3	3	3	3	3	3	3	4	4	4	4	4	4	4	4	4	4	5	5	5
	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2
P1	R	R	R	R	R	R	R	R	R	R	R	R	R	R	X	X	X	X	X				
P2																							
P3	X	X	X	X	X	X	X	X	X	R	R	R	R	R	R	R	R	R	R	X	X	X	X
P4			R	R	R	R	R	R	R	X	X	X	X										

Job throughput: 4 jobs in 51 time units, or **(4/52) jobs per time unit**

Avg waiting time: if wait time = time on ready queue (textbook definition), then this is  $(23 + 9 + 18 + 7)/4 = 57/4$  **time units**

if wait time == time on wait queue (my definition), then this is 0 time units.

Avg turnaround time: turnaround time = time of completion of process – time of submission of process (page 128 of textbook)

$$\text{avg turnaround} = ((48-0) + (19-1) + (52-20) + (43-32))/4 = 109/4 \text{ time units}$$

b. With non-preemptive priority scheduling (given the above priorities)?

Assuming higher numbers mean higher priority:

	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	2	2	2	2	2	2	2	2	2	2	2
	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9		
<b>P1</b>	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X							
<b>P2</b>		R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R
<b>P3</b>																							R	R	R	R	R	X	X	X	X	X
<b>P4</b>																																

	3	3	3	3	3	3	3	3	3	3	4	4	4	4	4	4	4	4	4	4	4	5	5	5
	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	
<b>P1</b>																								
<b>P2</b>	R	R	R	R	R	R	R	R	R	R	R	R	R	X	X	X	X	X	X	X	X	X	X	
<b>P3</b>	X	X	X	X	X	X	X	X	X															
<b>P4</b>			R	R	R	R	R	R	R	X	X	X	X											

Job throughput: 4 jobs in 51 time units, or **(4/52) jobs per time unit**

Avg waiting time: if wait time = time on ready queue (textbook definition), then this is  $(0 + 42 + 5 + 7)/4 = \mathbf{54/4 \text{ time units}}$

if wait time == time on wait queue (my definition), then this is 0 time units.

Avg turnaround time: turnaround time = time of completion of process – time of submission of process (page 128 of textbook)

$$\text{avg turnaround} = ((25-0) + (52-1) + (39-20) + (43-32))/4 = \mathbf{106/4 \text{ time units}}$$

Assuming lower numbers mean higher priority:

	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	2	2	2	2	2	2	2	2	2	2	2
	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9		
<b>P1</b>	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X							
<b>P2</b>		R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	X	X	X	X	X
<b>P3</b>																						R	R	R	R	R	R	R	R	R	R	
<b>P4</b>																																

	3	3	3	3	3	3	3	3	3	3	4	4	4	4	4	4	4	4	4	4	4	5	5	5
	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	
<b>P1</b>																								
<b>P2</b>	X	X	X	X																				
<b>P3</b>	R	R	R	R	R	R	R	R	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	
<b>P4</b>			R	R	X	X	X	X																

Job throughput: 4 jobs in 51 time units, or **(4/52) jobs per time unit**

Avg waiting time: if wait time = time on ready queue (textbook definition), then this is  $(0 + 24 + 18 + 2)/4 = 44/4 = 11$  time units

if wait time == time on wait queue (my definition), then this is 0 time units.

Avg turnaround time: turnaround time = time of completion of process – time of submission of process (page 128 of textbook)

$$\text{avg turnaround} = ((25-0) + (34 -1) + (52 -20) + (38-32))/4 = 86/4 \text{ time units}$$

2. Consider the Sleeping-Barber Problem (p202, question 6.7 in the textbook,) with the modification that there are k barbers and k barber chairs in the barber room, instead of just one. Write a program to coordinate the barbers and the customers using Java, C, or pseudo-code. You can use either semaphores or monitors.

Here's a solution that uses semaphores:

```
// shared data
Semaphore waiting_room_mutex = 1;
Semaphore barber_room_mutex = 1;

Semaphore barber_chair_free = k;
Semaphore sleepy_barbers = 0;

Semaphore barber_chairs[k] = {0, 0, 0, ...};
int barber_chair_states[k] = {0, 0, 0, ...};

int num_waiting_chairs_free = N;

boolean customer_entry( ) {

    // try to make it into waiting room
    wait(waiting_room_mutex);
    if (num_waiting_chairs_free == 0) {
        signal(waiting_room_mutex);
        return false;
    }
    num_waiting_chairs_free--; // grabbed a chair
    signal(waiting_room_mutex);

    // now, wait until there is a barber chair free
    wait(barber_chair_free);

    // a barber chair is free, so release waiting room chair
    wait(waiting_room_mutex);
    wait(barber_room_mutex);
    num_waiting_chairs_free++;
    signal(waiting_room_mutex);

    // now grab a barber chair
    int mychair;
    for (int I=0; I<k; I++) {
        if (barber_chair_states[I] == 0) { // 0 = empty chair
            mychair = I;
            break;
        }
    }
}
```

```

    }
}
barber_chair_states[mychair] = 1; // 1 = haircut needed
signal(barber_room_mutex);

// now wake up barber, and sleep until haircut done
signal(sleepy_barbers);
wait(barber_chairs[mychair]);

// great! haircut is done, let's leave. barber
// has taken care of the barber_chair_states array.
signal(barber_chair_free);
return true;
}

void barber_enters() {
    while(1) {
        // wait for a customer
        wait(sleepy_barbers);

        // find the customer
        wait(barber_room_mutex);
        int mychair;
        for (int I=0; I<k; I++) {
            if (barber_chair_states[I] == 1) {
                mychair = I;
                break;
            }
        }
        barber_chair_states[mychair] = 2; // 2 = cutting hair
        signal(barber_room_mutex);

        // CUT HAIR HERE
        cut_hair(mychair);

        // now wake up customer
        wait(barber_room_mutex);
        barber_chair_states[mychair] = 0; // 0 = empty chair
        signal(barber_chair[mychair]);
        signal(barber_room_mutex);

        // all done, we'll loop and sleep again
    }
}
}

```

3. Consider the following C++-style pseudo-code. We have a producer thread and a consumer thread running concurrently. Both threads have access to the shared data. Will they function correctly (i.e. each produced item will be consumed)? Why or why not?

No, they will NOT function correctly. If the consumer thread gets context switched out after line 46 and before line 48, then the item on the stack might be overwritten by the producer thread (line 27).

4. Use Java, C, or pseudo-code to implement:

- monitors using semaphores
- semaphores using monitors

Your solution may *\*not\** busy-wait.

monitors using semaphores:

- the answer, for the most part, is in section 6.7 of the text. here's some brief pseudocode to fill in the blanks.

```
Semaphore mutex = 1, next = 0;  
int next_count = 0;
```

For each external procedure F:

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

For each condition x

```
int x_count = 0;  
semaphore x_sem = 0;  
  
////////////////////////////////////  
// x.wait  
x_count = x_count + 1;  
if (next_count > 0)  
    signal(next);  
else  
    signal(mutex);  
wait(x_sem);  
x_count = x_count - 1;
```

```

////////////////////////////////////
// x.signal
if (x_count > 0) {
    next_count = next_count + 1;
    signal(x_sem);
    wait(next);
    next_count = next_count - 1;
}

```

- semaphores using monitors:

```

class Semaphore : public Monitor
{
protected:
    int count;
    condition cond;

public:
    Semaphore(int initial) {
        count = initial;
    }

    void wait() {
        count = count - 1;
        while (count < 0) {
            cond.wait();
        }
    }

    void signal() {
        count = count + 1;
        cond.signal();
    }
};

```

5. Use Java, C, or pseudo-code to implement:

- semaphores using locks
- locks using semaphores

Is it possible to produce a solution that doesn't busywait?

**Answer:** (sort of):

It is trivial to build solutions that busy wait (I won't bother writing any out here).

It is extremely hard to build solutions that don't busy wait. Essentially, you end up having to build a miniature thread scheduler, complete with waiting queues. We graded liberally on this bonus question.