1. Some operating system uses a preemptive SJF (Shortest Job First) algorithm for CPU scheduling. Processes with a lower ID number have priority in case of a tie. Several processes are submitted for execution: their arrival times and burst-intervals are summarized in the following table.

Process	Arrival	Burst Distribution
P1	0	C[a](1), D1[b](1.5), C(1), D2[c](2), C(2)
P2	0.5	C(1.5), D2(2.5), C(0.5), D2(1.5), C(1)
P3	1	C(0.5), D1(0.5), C(0.5), D2(0.5), C(0.5), D1(0.5), C(0.5), D2(0.5), C(0.5), D1(0.5), C(0.5)

a. C stands for the CPU and the number between parantheses for the length of the CPU burst
b. D1 stands for an I/O request on disk #1; the number between parantheses is the length of the I/O burst
c. D2 stands for an I/O request on disk #2; the number between parantheses is the length of the I/O burst

The length of the I/O burst time does not include the waiting time in the disk queue. Disk requests are processed on a FCFS basis.

a) Draw the Gantt chart illustrating the execution of these processes. Use the ruled paper attached to your exam.

b) What is the CPU utilization?

CPU_{util} = (12 - 2)/12 = 83.3%

c) What is the throughput?

Assume that the burst times are given in seconds.

throughput = jobs_completed/time = 3/12 = 0.25 jobs/s

d) What is the average turnaround time?

t_1 = 10 s
t_2 = 11.5 - 0.5 = 11 s
t_3 = 12 - 1 = 11 s

t_{avg} = (10+11+11)/3 = 10.67 s

e) What is the average waiting time? The waiting time is defined as the *total* time spent by

© Virgil Bistriceanu, 1996

a process in the waiting queue.

Note: the waiting time is the CPU waiting time

w_1 = 1 s
w_2 = 2 s
w_3 = 0.5 s

w_{avg} = (1 + 2 +0.5)/3 = 1.17

f) What is the average waiting time spent in disk waiting queues?

wd_1 = 1.5 s
wd_2 = 2 s
wd_3 = 5 s

wd_{avg} = (1.5 + 2 + 5)/3 = 2.83 s

2. Fetch And Add (FAA), *atomically* increments a variable and returns the old value of the variable. `FAA(var, increment)` adds `increment` to the variable `var` and returns the initial value of `var`. We plan to use this for mutual exclusion:

```
shared int S;

// each process executes
while(TRUE){
      enter_section();
      ......                    // here is the critical section
      exit_section();
      ......                    // this is the remainder section
}

// enter_section() returns when it is safe to enter the critical
// section
void enter_section()
{
      while( FAA(S,-1) < 1 ) {
            FAA(S,1);
      }
}
```

a) What should be the initial value of S for this code to work? **$S = 1$**

b) Write the code for `exit_section()`.

```
void exit_section()
{
      FAA(S, 1);
}
```


© Virgil Bistriceanu, 1996

3. Two processes, P1 and P2, execute at the same time the following code:

```
for (i=0, i<2; i++)
      S = S+1;
```

where S is a shared integer with the initial value of 0.

a) What would be the value of S if the two processes would execute to completion in the order P1, P2 (this means that P1 starts and runs to completion, then P2 executes)?

$S = 4$

b) What are possible values for S in reality? Hint: write the assembly code that implements the `for` loop.

```
      li    r1, 0        # i in r1
      li    r2, 2        # loop limit in r2
loop: bge r1, r2, done   # exit if i greater than limit
      ld r3, S           # r3 gets the value of S
      addi r3, 1         # increment the value of S
      st S, r3           # store the value of S back in memory
      addi r1, 1         # i <- i+1
      j loop
done:
```

Depending on the order instructions of the two processes are interleaved, the possible values for S are 2, 3 and 4.

c) Use semaphore(s) to fix the problem: in other words, modify the code in such a way as to always obtain the result described in part a.

```
semaphore mutex=1;

for (i=0; i<2; i++) {
      wait(mutex);       // wait(&mutex) to be more precise
      S = S+1;
      signal(mutex);
}
```


© Virgil Bistriceanu, 1996

4. At same place, the North-South Highway goes over a river. The bridge is in construction and the traffic on the bridge is on a single lane. At most MAX cars can be on the bridge at any time, all going the same direction (either North or South, but never in both directions at the same time).

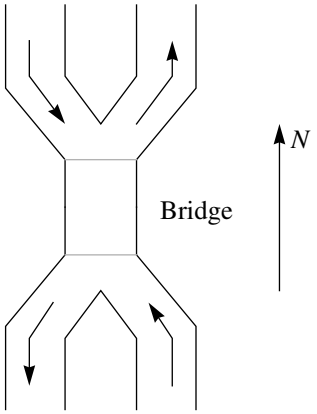

a) Devise a solution with semaphores for this problem.

```
semaphore max_S=MAX,    // counts southbound cars on the bridge
      max_N=MAX,  // counts northbound cars on the bridge
      mutex_S=1,  // mutual exclusion for S variables
      mutex_N=1,  // mutual exclusion for N variables
      bridge=1;   // only one dirsction at a time on the bridge
int cars_S=0, cars_N=0; // needed to decide when to release bridge

// Going_south_car:

      wait(max_S);
      wait(mutex_S);
      cars_S++;
      if (cars_S == 1) wait(bridge);
      signal(mutex_S);

            // on the bridge //

      wait(mutex_S);
      cars_S--;
      signal(max_S);    // since I exit I can let others in
      if (cars_S == 0) signal(bridge);
      signal(mutex_S);
```

b) Is your solution starvation-free? Explain.

No, a continuous flow of southbound cars will starve northbound cars.

© Virgil Bistriceanu, 1996

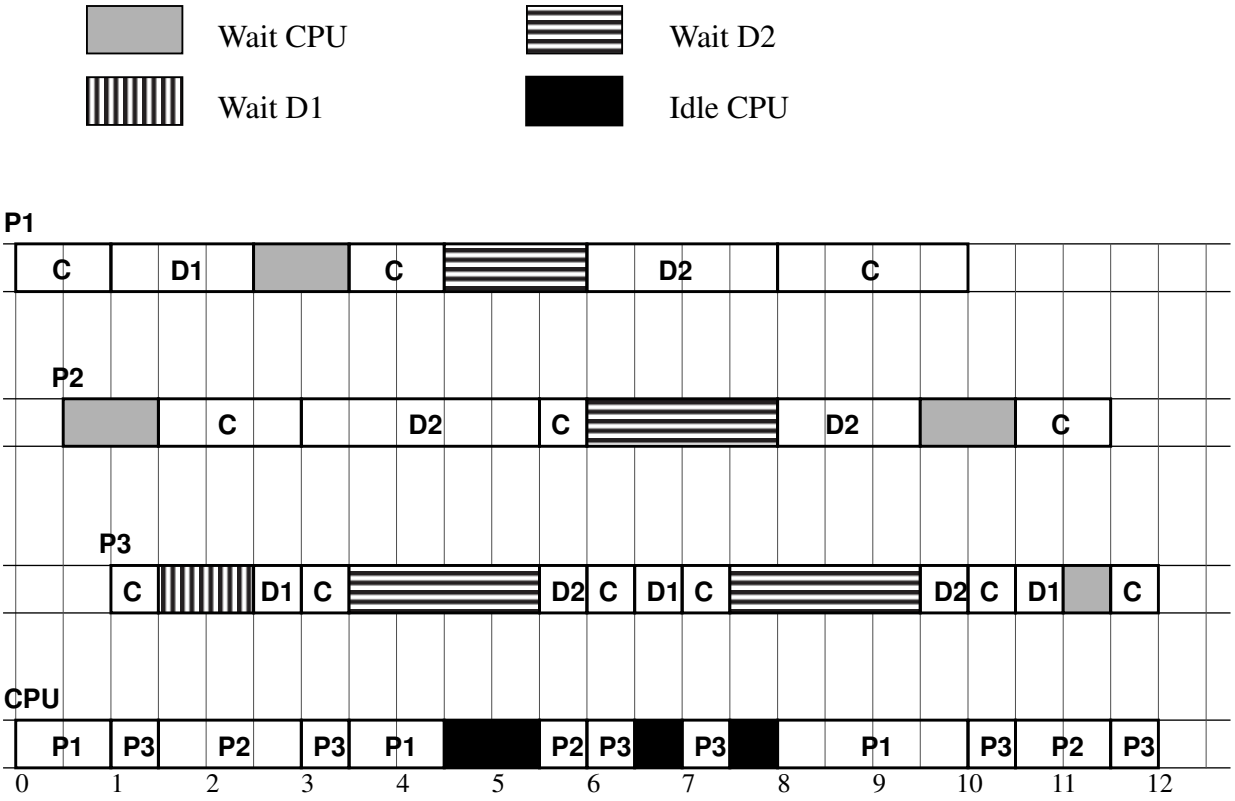

© Virgil Bistriceanu, 1996