14 marks Question 5. CPU Scheduling

12 marks **Part (a)** An OS uses a multi-level feedback scheduler with 2 levels. New and returning processes start at level 0, which uses round robin scheduling with a quantum of 2. A process that uses its entire quantum at level 0 gets moved to level 1, which uses first-come-first-served scheduling. The scheduler always chooses a process from the lowest-numbered non-empty level. Admission of a new process or a process returning from IO does not force a scheduling decision; those processes are placed on the level 0 queue to wait until a scheduling decision is to be made.

Initially, there is one process, P0. New processes P1 and P6 are created at times 1 and 6 respectively. Those are the only three processes you need to schedule. Each process has a CPU burst of 5 time units, an IO burst of 3 time units, and a CPU burst of 1 time unit. During an IO burst, a process is blocked (on a wait queue); it does not require the processor.

Assume context switches take no time. Fill in the timeline below for each process. Note when a process is New, when it has the CPU, when it loses the CPU due to a *Preempt*, what queue it is on (L0, L1, IO), and when it *Exits*. The first 5 time units have been filled in for you.

Time	0	1	2	3	4	5	6
P0	CPU	CPU	Preempt, L1	L1	CPU		
P1		New, L0	CPU	CPU	Preempt, L1		
P6							
Time	7	8	9	10	11	12	13
P0							
P1							
P6							
Time	14	15	16	17	18	19	20
P0							
P1							
P6							

1 mark

Part (b) For your solution above, what is the total wait time (time spent ready but not executing) for all processes?

1 mark **Part (c)** For **your solution above**, what is the average response time (time spent waiting for the processes's first time slice)?