COMPUTER SCIENCE TRIPOS Part IA – 2024 – Paper 2

3 Operating Systems (rmm1002)

- (a) When an operating system needs to change a process' state from *ready* to *running* it does so by executing a *context switch*. What state must always be included in the context switch? What state must be flushed from the hardware before the new process is allowed to start running? [4 marks]
- (b) A single-CPU system has five CPU-bound processes arrive simultaneously. Each has total CPU demand 10 ms, 20 ms, 30 ms, 40 ms, 50 ms respectively. Give the schedule by which processes are run and state how many context switches occur under the following scheduling regimes. Assume all processes have the default *nice* level.
 - (i) Round Robin, with a quantum of $q = 5 \,\mathrm{ms.}$
 - (*ii*) Round Robin, with a quantum of q = 20 ms.
 - (*iii*) Completely Fair Scheduler, with a *target latency* of 60 ms, and a *minimum granularity* of 2 ms.
 - (iv) Completely Fair Scheduler, with a target latency of $20 \,\mathrm{ms}$, and a minimum granularity of $5 \,\mathrm{ms}$.

State which would be preferable for processing batch workloads, and which would be preferable for processing interactive workloads. Give a reason in each case. [10 marks]

(c) A modern-day cloud-provider decides to offer a service using Linux to which customers can deploy single *functions* as stand-alone units of computation. Their intention is that the relatively small size of such functions will give them more predictable behaviour, increasing customer satisfaction, as well as enabling instances to be packed more densely onto each physical machine, reducing the cloud-provider's costs. Suggest one reason why either assumption may fail, what effect that failure will have, and how the cloud-provider might try to ensure the assumption holds.