

Question 5.1: Scheduling Basics

- a. What is the purpose of scheduling?
- b. What is the difference between a long and short-term scheduler?
- c. Consider an operating system that supports the five task states “new”, “running”, “ready”, “waiting”, and “terminated”. Depict the possible state transitions and the events that cause them.
- d. What quantitative metrics/criteria can be used to estimate the quality of a scheduling policy?
- e. What kind of hardware support is required for an operating system that implements a non-cooperative scheduling policy?
- f. What are common values for the length of a timeslice?
- g. Discuss pros and cons of choosing a short timeslice length vs. choosing a longer timeslice length.

Question 5.2: Scheduling Policies

- a. How does SJF scheduling work? What is the difference between preemptive and non-preemptive SJF?
- b. What is the basic idea of priority scheduling?
- c. What is the major problem of priority-based algorithms? What is a possible solution?
- d. Explain the multilevel feedback queue scheduling algorithm. What kind of processes can be found in the high-level queues, what kind of processes sink to the lower-level queues?
- e. Describe the idea of lottery scheduling.
- f. Enumerate possible advantages of lottery scheduling over “classical” priority-based scheduling algorithms.