

7. Summary

MLFQ Rules:

1. If $\text{Priority}(A) > \text{Priority}(B)$, A runs (B doesn't).
 2. If $\text{Priority}(A) = \text{Priority}(B)$, A & B run in round-robin fashion using the time slice (quantum length) of the given queue.
 3. When a job enters the system, it is placed at the highest priority (the topmost queue).
 4. Once a job uses up its time allotment at a given level (regardless of how many times it has given up the CPU), its priority is reduced (i.e., it moves down one queue).
 5. After some time period S, move all the jobs in the system to the topmost queue.
-

Other Scheduling Policies: (skip)

Chapter 9: fair-share, lottery scheduling (random), stride scheduling (deterministic)

Chapter 10: multiprocessor scheduling