RTDS408 Tutorial Problems #1 - Real-Time Scheduling Theory

1. Consider the case of three periodic tasks:

   Task t_1: $C_1 = 20$ ms; $T_1 = 100$ ms
   Task t_2: $C_2 = 40$ ms; $T_2 = 150$ ms
   Task t_3: $C_3 = 100$ ms; $T_3 = 350$ ms

   Apply the Utilization Bound Theorem to determine if these tasks are schedulable using a rate monotonic scheduling strategy. Suppose the computation time for task 1 doubles to 40 msec, now determine if the tasks are schedulable, and then apply the less conservative Completion Time Theorem.

2. Suppose we have four tasks: two periodic, one aperiodic, and one interrupt driven aperiodic. The non-interrupt driven tasks require access to a shared data store, and we wish to give the interrrupt-drive task the highest priority:

   periodic task t_1: $C_1 = 30$ ms, $T_1 = 100$ ms
   aperiodic task t_2: $C_2 = 30$ ms, $T_2 = 150$ ms
   interrupt driven aperiodic task t_a: $C_a = 10$ ms, $T_a = 200$ ms
   periodic task t_3: $C_3 = 30$ ms, $T_3 = 300$ ms

   The context switch time is included in the indicated CPU times. Use the Generalized Utilization Bound Theorem to determine if this task set is schedulable.

3. Given two tasks T_1 and T_2 with two shared data structures protected with binary semaphores S_1 and S_2, show how the *priority ceiling protocol* prevents mutual deadlock and guarantees that a high-priority task will be blocked by at most one critical section of any lower priority task.