Questions:

1. (30 points) Consider a queueing system where customers arrive according to a Poisson process with rate λ and join a single queue. There are two servers serving customers from this single queue. The service time of a customer is exponentially distributed with rate μ_i if it is served by server i, where $\mu_1 > \mu_2$. Assume that the service is non-preemptive, i.e., once the service of a customer is initiated, the decision cannot be reversed. The problem is to find the server allocation policy that minimizes the infinite-horizon total discounted expected number of customers in the system where the continuous discounting factor is α. Note that in this problem, the tradeoff is between waiting for the fast server (server 1) to become available and committing a customer to the slow server. Hence, it may be optimal to keep the second server idle even if there are customers waiting in the queue. (You can assume that the fast server never idles whenever there is a customer in the system.)

   (a) Formulate this sequential decision making problem as a Markov decision process. (Be specific about the state definition, action space, etc.)

   (b) Write down the optimality equations that need to be solved to compute the optimal policy for the infinite-horizon problem with discounting.

   (c) Show that the optimal policy is a threshold-type policy, i.e., there exists some number N such that the optimal policy is to allocate a customer to the slow server if and only if the number of customers in the queue is at least N. Hint: First obtain a finite-horizon version of the uniformized problem, and then prove that the optimal value function possesses some monotonicity properties. Finally, carry-over your results to the infinite-horizon model by using a rigorous argument that you learned in class.

2. (30 points) A collection of n jobs is to be processed on two parallel machines. The processing time of job i is exponentially distributed with rate μ_i, for $i = 1, \ldots, n$ and is independent of the processing time of all other jobs and the order of processing. Assume that the processing is preemptive, i.e., a job that is already in process can be replaced by another job. Let C_i be the completion time of job i. Define $\max_{i=1,\ldots,n}\{C_i\}$ to be the *makespan*, i.e., the time when all jobs are complete. The objective is to determine the processing order of these jobs dynamically, i.e., based on system state, to minimize the expected makespan.

   (a) Formulate this sequential decision making problem as a Markov decision process. (Be specific about the state definition, action space, etc.)

   (b) Write down the optimality equations that need to be solved to find the optimal policy.

   (c) Show that the optimal policy is to serve the job with the smallest μ_i first.

3. (10 points) For an average-reward unichain Markov decision process with finite state space $\mathcal{S}$, consider a stationary policy γ with long-run average reward g_γ. One can compute g_γ by using one of the solution algorithms for unichain Markov decision processes because it can be shown that there exists a bounded function $h(i)$, $i \in \mathcal{S}$ such that

$$g_\gamma + h(i) = r(i, a_i) + \sum_{j \in \mathcal{S}} p(j|i, a_i)h(j), \forall i \in \mathcal{S}, \tag{1}$$

   where a_i is the action taken in state i under policy γ. On the other hand, from STOR 641, we also know that

$$g_\gamma = \sum_{i \in \mathcal{S}} r(i, a_i)\pi_i, \tag{2}$$

   where π_i is the limiting probability of state i under policy γ. Prove that these two methods of computing g_γ are consistent, i.e., solving equation (1) or using expression (2) will yield the same value for g_γ.

4. (30 points) Consider the dynamic control problem studied in the paper by Ahn H., I. Duenyas, and M. Lewis, "Optimal Control of a Two-stage Tandem Queueing System with Flexible Servers," *Probability in the Engineering and Informational Sciences* **16** (2002), 453–369. In particular, consider the case where both servers are collaborative. For this case, suppose that the problem is generalized as follows. First, servers are not necessarily identical and may have different service rates. Let $\mu_i(j) > 0$ denote the service rate when server j is the only server working at station i, for $i, j \in \{1, 2\}$. Second, when servers collaborate on a job, the pooled service rate is not necessarily equal to the sum of the individual service rates. Let $\bar{\mu_i} > 0$ denote the service rate when both servers collaborate on a job at station $i = 1, 2$. Assuming all other modeling assumptions remain the same, answer the following questions.

   (a) Assuming that the queueing system is stable under any policy, re-write the optimality equations (2.3) and (2.4) in the paper for the generalized problem.

   (b) Will Proposition 3.1 still hold for this generalized MDP? If yes, prove. Otherwise, provide sufficient conditions on $\mu_i(j)$ and $\bar{\mu_i}$ for $i, j \in 1, 2$ such that Proposition 3.1 still holds. (Of course, an immediate sufficient condition is $\mu_i(j) = \mu_i$ and $\bar{\mu_i} = 2\mu_i$ as proved in the paper. You are expected to provide a weaker or a different set of conditions.)