Control of Stochastic Processes 048913

Winter 2006 Supplement: interchange arguments

11/12/2006. For the latest see http://www.ee.technion.ac.il/~adam/GRADUATES/048913

Consider the following control problem.

There are K queues, and a vector arrival process: $A_k(n)$ new arrivals to queue k at time n. The single server may serve one queue at a time. Assume that after we start serving a customer we may switch to another customer after any amount of time (preemptive-resume).

Service times are geometric with mean μ_k, i.i.d. for each queue and independent across queues, as well as independent of the arrival process. Note that since arrival times are general, that this is not Markov even if we use a Markov policy.

Assume discounted cost, with linear holding costs. That is, if $X(n)$ denotes the (vector) queue length at time n, we have (a vector of) positive constants C so that the immediate cost is $C \cdot X$. We want to minimize

$$V(X,\pi) = \mathbb{E}_x^\pi \sum_{n=0}^{\infty} \beta^n C \cdot X(n) \ .$$

To find the optimal policy we now use an interchange argument. Fix any policy π. The behavior below is under π: Suppose that we serve queue i at time N, and that at that time queue j is not empty. Let τ denote the (random) first time after N that we serve j (under π). Now interchange service order between the two queues, that is, serve j at time N and i at time τ, keeping everything else exactly as before: note that this means that the policy π behaves as if the change did not take place.

This leads to a change in immediate cost as follows. Serving j means that with probability μ_j a type-j customer will leave at time N (instead of at τ), leading to a reduction of cost of

$$\sum_{n=N+1}^{\tau} C_j \beta^n \ .$$

However, since we now do not serve queue i at time N, with probability μ_i a type-i customer will not leave at time N, leading to an increase in cost of

$$\sum_{n=N+1}^{\tau} C_i \beta^n \ .$$

Thus. if we compute the net gain, it is

$$\mathbb{E}\, \mu_j \sum_{n=N+1}^{\tau} C_j \beta^n - \mathbb{E}\, \mu_i \sum_{n=N+1}^{\tau} C_i \beta^n = (\mu_j C_j - \mu_i C_i)\, \mathbb{E} \sum_{n=N+1}^{\tau} \beta^n \ .$$

We conclude that we should indeed make this change whenever $\mu_j C_j > \mu_i C_i$, that is, we should always serve the non-empty queue with the largest value of $\mu_k C_k$. This is the famous $\mu - C$ rule.

This result can be obtained using dynamic programming arguments, when the arrival process is i.i.d. (so that this is a MDP).

Interchange arguments are a powerful tool—when they apply. In order to make this precise, we need the following.

Definition 0.1 *Let X_i be a RVs defined on probability spaces $(\Omega_i, \mathcal{F}_i, \mathbb{P}_i)$. A coupling is a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and RVs $\tilde{X}_i$ so that for each i*

$$\mathbb{P}(\tilde{X}_i \in A) = \mathbb{P}_i(X_i \in A) \ .$$

The RVs could be sequences, random processes etc. The main idea here is that it may be easier to compare on a sample path basis, so we need to construct all RVS on the same probability space.

In the example above this is done by choosing the RVs that describe completion of service in the right we, so that the two policies can be compared.

Coupling is a powerful method and is a well developed theory. In the context of MDP it became very popular in the 80's, and led to some beautiful and elegant results: however, it was soon discovered that it is rarely applicable.