ECE 555 **Assignment # 3**

Issued: February 11 Due: February 17, 2011

Solutions

Reading: Begin Chapter 9 of CTCN.

Exercises:

7. *Reversibility* The M/M/1 queue in discrete time is defined by $Q(t+1) = [Q(t) - S(t+1) + A(t+1)]_+$, where $\boldsymbol{A}$ is i.i.d. Bernoulli, and $S(t) = 1 - A(t)$ for $t \geq 1$. Denote $\alpha = \mathsf{E}[A(t)]$, $\mu = 1 - \alpha$, and the load is $\rho = \alpha/\mu$. Consider the case of a finite waiting room, of size N,

$$Q(t+1) = [Q(t) - S(t+1) + A(t+1)]_0^N$$

where $[x]_0^N = \max(\min(x, N), 0)$, $x \in \mathbb{R}$. Then $\boldsymbol{Q}$ is a Markov chain on the finite set $\{0, \ldots, N\}$. Let P denote the transition matrix.
Verify that the chain is *reversible*: There is a probability measure π satisfying the *detailed balance equations*,

$$\pi(x)P(x,y) = \pi(y)P(y,x)$$

Note that on summing each side of this equation over x, you obtain invariance $\pi P = \pi$.
Hint: When $N = \infty$ and $\rho < 1$, we have $\pi(x) = (1-\rho)\rho^x$.

***Solution*:** Let's see if $\mu(x) = \rho^x$ satisfies the detailed balance equations: For $1 \leq x \leq N-1$ we have,

$$\begin{aligned}\mu(x)P(x,y) &= \rho^x(\alpha \mathbf{1}\{y = x+1\} + \mu \mathbf{1}\{y = x-1\}\,,\\ \mu(y)P(y,x) &= \rho^y(\alpha \mathbf{1}\{x = y+1\} + \mu \mathbf{1}\{x = y-1\})\,.\end{aligned}$$

We can transform the second identity into the first to obtain the detailed balance equations for this range of x:

$$\begin{aligned}\mu(y)P(y,x) &= \rho^{x-1}\alpha \mathbf{1}\{x = y+1\} + \rho^{x+1}\mu \mathbf{1}\{x = y-1\}\\ &= \rho^x \alpha/\rho \mathbf{1}\{x = y+1\} + \rho^x \mu\rho \mathbf{1}\{x = y-1\}\\ &= \rho^x(\mu \mathbf{1}\{x = y+1\} + \alpha \mathbf{1}\{x = y-1\})\\ &= \mu(x)P(x,y)\end{aligned}$$

For $x = 0$ we have,

$$\begin{aligned}\mu(x)P(x,y) &= \alpha \mathbf{1}\{y = x+1\} + \mu \mathbf{1}\{y = x\}\,,\\ \mu(y)P(y,x) &= \rho^y(\mu \mathbf{1}\{y = 1\} + \mu \mathbf{1}\{y = 0\})\,.\end{aligned}$$

Consequently, for $x = 0$,

$$\begin{aligned}\mu(y)P(y,x) &= \rho^y(\mu \mathbf{1}\{y = 1\} + \mu \mathbf{1}\{y = 0\})\\ &= \alpha \mathbf{1}\{y = 1\} + \mu \mathbf{1}\{y = 0\}\\ &= \alpha \mathbf{1}\{y = x+1\} + \mu \mathbf{1}\{y = x\}\\ &= \mu(x)P(x,y)\end{aligned}$$

For $x = N$ we have,

$$\mu(x)P(x,y) = \rho^x(\alpha\mathbf{1}\{y = x\} + \mu\mathbf{1}\{y = x-1\}),$$
$$\mu(y)P(y,x) = \rho^y(\alpha\mathbf{1}\{y = N\} + \alpha\mathbf{1}\{y = N-1\}).$$

Consequently, in this case

$$\begin{aligned}\mu(y)P(y,x) &= \rho^y(\alpha\mathbf{1}\{y = N\} + \alpha\mathbf{1}\{y = N-1\}) \\ &= \rho^N\alpha\mathbf{1}\{y = N\} + \rho^{N-1}\alpha\mathbf{1}\{y = N-1\} \\ &= \rho^N\alpha\mathbf{1}\{y = N\} + \rho^N\mu\mathbf{1}\{y = N-1\} \\ &= \mu(x)P(x,y)\end{aligned}$$

Note that $\mu(\mathsf{X}) := = \mu(\{1,2,\ldots,N\}) = \sum_{x=0}^{N}\rho^x = \frac{1-\rho^{N+1}}{1-\rho} < \infty$. Thus

$$\pi(x) = \mu(x)/\mu(\mathsf{X}) = \frac{1-\rho}{1-\rho^{N+1}}\rho^x$$

is a probability measure satisfying the detailed balanced equation. □

8. *Rate of convergence in value iteration* In the previous model take $\rho = \alpha/\mu = 0.95$. Work out the following using Matlab.

(i) Compute the first and second largest *eigenvectors* of P for a three values of N (say, $N = 5, 10, 50$).

(ii) For each of these values of N, obtain the solution to Poisson's equation with $c(x) = x$, using the value iteration algorithm. You might experiment with different initial conditions: $V_0(x) = 0$, or $V_0(x) = \frac{1}{2}(\mu-\alpha)^{-1}x^2$ (the fluid value function).

(iii) Estimate the rate of convergence λ, where $\Lambda = \log(\lambda)$ is given by,

$$\Lambda := \lim_{n\to\infty} n^{-1}\log(\|h - h_n\|)$$

How does λ compare with λ_2, the second largest eigenvalue for P?

***Solution*:** For part (i), we use the command $\texttt{eig}(P)$ where P is the transition matrix. The command $\texttt{sort}(\texttt{eig}(P), '\texttt{descend}')$ will sort the eigenvalues in the descending order. The following is a list of the eigenvalues of P for $N = 5, 10, 50$, and Figure 2 is a plot of $\lambda_2(P)$ for $2 \le N \le 50$. Figure 2 also shows all eigenvalues of P when $N = 50$.
For part (ii) we apply the VIA described in lecture and in the notes.
Figure 3 plots our estimate of h after 100 iterations for $N = 10$.
For (iii), recall the "bit of theory":

$$P^n = (P - 1\otimes\pi)^n + 1\otimes\pi$$

The matrix $(P - 1\otimes\pi)^n$ has maximal eigenvalue equal to λ_2^n, which is strictly less than one, and hence the matrix product converges to zero geometrically fast. Therefore, the difference converges to zero at the same rate:

$$P^n(x,z) - P^n(y,z) = (P - 1\otimes\pi)^n(x,z) - (P - 1\otimes\pi)^n(y,z)$$

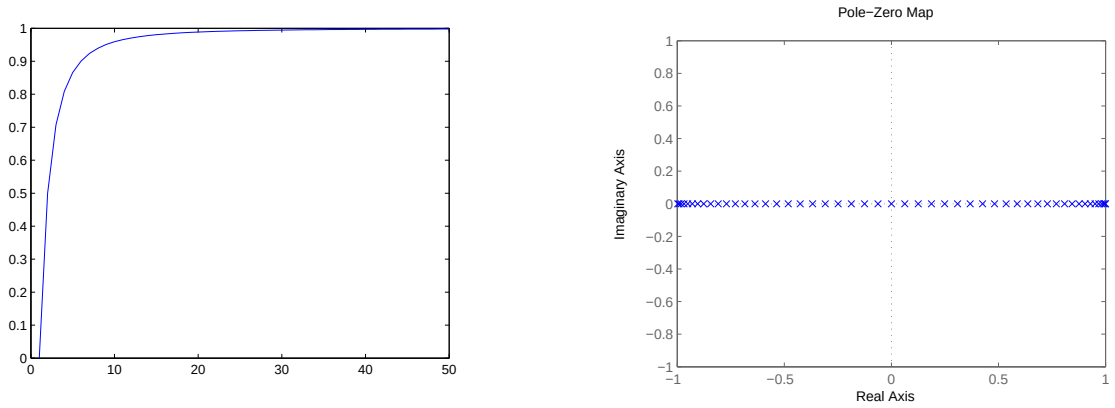

Figure 1: Shown on the left is a plot of $\lambda_2(P)$ vs. N. Plotted on the right are the eigenvalues of P in the complex plane

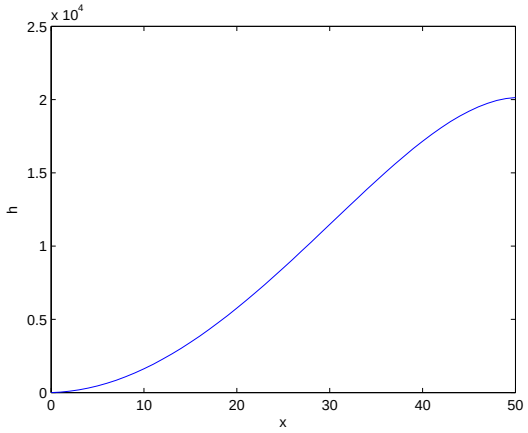

Figure 2: Plot of h for $N = 50$. A quadratic approximation is accurate "where the process lives" — $x \leq 30$.

The VIA computes,

$$V_n(x) = P^n V_0(x) + \sum_{i=0}^{n-1} P^i c(x)$$

and $h_n(x) = V_n(x) - V_n(x^*)$ can thus be expressed,

$$h_n(x) = (P - 1 \otimes \pi)^n V_0(x) - (P - 1 \otimes \pi)^n V_0(x^*) + \sum_{i=0}^{n-1} \Big((P - 1 \otimes \pi)^i c(x) - (P - 1 \otimes \pi)^i c(x^*) \Big)$$

A solution to Poisson's equation is given by,

$$h(x) = \sum_{i=0}^{\infty} \Big((P - 1 \otimes \pi)^i c(x) - (P - 1 \otimes \pi)^i c(x^*) \Big)$$

Hence, the error can be expressed,

$$h(x) - h_n(x) = -(P - 1 \otimes \pi)^n V_0(x) + (P - 1 \otimes \pi)^n V_0(x^*) + \sum_{i=n}^{\infty} \Big((P - 1 \otimes \pi)^i c(x) - (P - 1 \otimes \pi)^i c(x^*) \Big)$$

which converges to zero at rate λ_2^n.

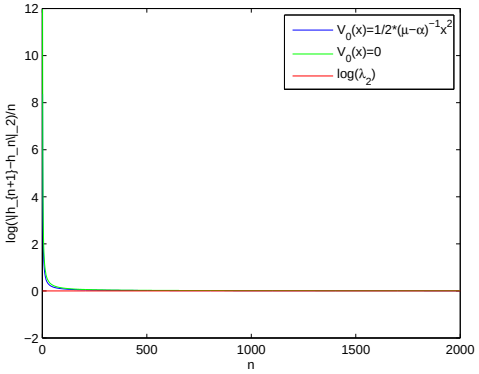

Figure 3: Blue: $\log(h_{n+1} - h_n)/n$ plotted against n and $h_0(x) \equiv \frac{1}{2}(\mu - \alpha)^{-1}x^2$. Green: $\log(h_{n+1} - h_n)/n$ plotted against n and $h_0(x) \equiv 0$.Red: $\log(\lambda_2)$.

However, Figure 3 plots $\log(h_{n+1} - h_n)/n$ for $n \geq 1$. The results are horrible! In theory, this should converge linearly to $-\infty$. Apparently, with this high load, numerical effects mask the geometric rate of convergence.

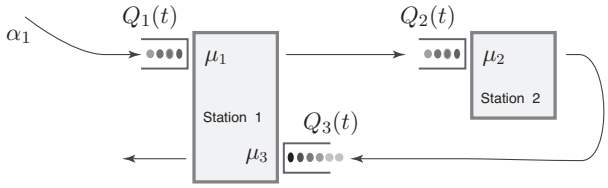

Figure 4: Simple re-entrant line

Numerics are bad because of the high load... As far as numerics go, it seems Prof. Meyn should not have set such a high load! Usually, numerics are much better behaved.
Figure 4 shows a network example from CTCN. Figure 5 provides examples of the convergence of value iteration to compute an *optimal policy*, with $c(x, u) = \|x\|_1 = \sum x_i$. The algorithm was initialized with $V_0 \equiv 0$, or with two different fluid value functions. The vertical axis shows the performance in terms of average cost for the nth policy. The improvement in convergence in the initial stages is remarkable.

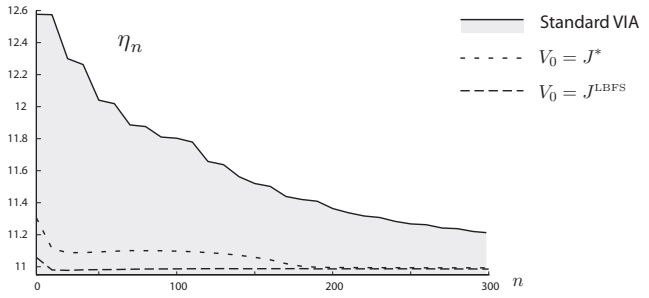

Figure 5: Convergence of the VIA with three different initializations: $V_0 \equiv 0$, or V_0 taken as a value function for the fluid mode for two different policies. The fluid value function was under the optimal fluid policy in one experiment, and in the other the policy was the LBFS (last buffer first served) priority policy.

The queue without truncation Final note: You might think that for the full M/M/1 queue, the eigenvector will be of the form $h(x) = \gamma^x$ for some $\gamma > 1$. For this function we have,

$$Ph\,(x) = \alpha h(x+1) + \mu h(x-u)$$

where $u = \mathbf{1}\{x \geq 1\}$. For $x \geq 1$ this gives, $Ph = \lambda h$, where $\lambda = \alpha\gamma + \mu\gamma^{-1}$. However, the eigenvector equation fails for $x = 0$.
Suppose we choose $\lambda \in (-1, 1)$, with $\lambda^2 > 4\alpha\mu$. Then, there are two solutions γ_+, γ_- to the equation,

$$\alpha\gamma + \mu\gamma^{-1} = \lambda$$

We can then take $h(x) = A_+\gamma_+^x + A_-\gamma_-^x$: For arbitrary constants $\{A_+, A_-\}$, the eigenvector equation holds for $x \geq 1$. We choose $\{A_+, A_-\}$ so that $Ph(0) = \lambda h(0)$ holds for $x = 0$. We have,

$$Ph\,(0) = \alpha h(1) + \mu h(0) = \alpha(A_+\gamma_+ + A_-\gamma_-) + \mu(A_+ + A_-)$$

Setting the right hand side equal to $\lambda h(0) = \lambda[A_+ + A_-]$ gives a linear relationship between A_+ and A_-.
We can attempt to summarize all of this as follows:

(i) For the M/M/1 queue there is no spectral gap: Any λ satisfying $\lambda^2 > 4\alpha\mu$ is an eigenvalue.

(ii) For the truncated model, $\lambda_2(N) \to 1$ as $N \to \infty$.

(iii) However, for the M/M/1 queue without truncation, there is NO spectrum in L_2 except for the trivial eigenvector $h \equiv 1$. You can check that any of the eigenvectors corresponding to $\lambda^2 < 1$, we have $\sum(h(n))^2\pi(n) = \infty$.

9. *Inverse dynamic programming* Consider the controlled Markov chain, evolving on $\mathbb{R}_+$:

$$X(t+1) = X(t) - U(t) + A(t+1)\,,$$

where $\boldsymbol{A}$ is i.i.d. on $\mathbb{R}_+$, with finite variance. The input is constrained: Given $X(t) = x$, we have $U(t) \in \mathsf{U}(x)$, where $\mathsf{U}(x) = \{x : 0 \le u \le x\}$. Let $h(x) = x^2$, and find a function $c(x)$ and constant η^* so that the ACOE holds,

$$\min_{u \in \mathsf{U}(x)} \{c(x) + u^2 + D_u h\,(x)\} = \eta^*, \qquad x \ge 0.$$

***Solution*:** For any given constant η^*, we *define*

$$c(x) = \eta^* - \min_{u \in \mathsf{U}(x)} \{u^2 + D_u h\,(x)\}, \qquad x \ge 0. \tag{1}$$

Without much effort at all we can see that this function is *coercive* — It's sublevel sets are bounded subsets of $\mathbb{R}_+$. To see this, let $\alpha = \mathsf{E}[A(1)]$, and observe that for $x \ge \alpha + 1$, we obtain a bound on c by setting $u \equiv \alpha + 1$:

$$c(x) \ge \eta^* - \{1 + D_1 h\,(x)\} = \eta^* - 1 - \big(\mathsf{E}[(x - 1 + A(1) - \alpha)^2] - x^2\big).$$

Since $A(1) - \alpha$ has zero mean, this simplifies to

$$c(x) \ge \eta^* - 1 - \big([(x-1)^2 + \sigma_A^2] - x^2\big).$$

where σ_A^2 denotes the variance of $A(1)$. The bound can be further simplified to,

$$c(x) \ge \eta^* - 2 - \sigma_A^2 + 2x$$

To compute c we must obtain the minimizer u^*, which requires differentiation: Assuming that the optimizer is in the interior of $\mathsf{U}(x)$ we conclude that,

$$0 = \frac{d}{du}\Big(u^2 - \big(\mathsf{E}[(x - u + A(1))^2]\big)\Big)\Big|_{u=u^*} = 2u^* - 2\mathsf{E}[(x - u^* + A(1))]$$

Hence $u^* = \frac{1}{2}(x + \alpha)$. This is feasible for $x \ge \alpha$. For $x < \alpha$ we conclude that $u^* = x$.
Plugging u^* into (1) then gives

$$c(x) = \eta^* - \begin{cases} x^2 + \big(x^2 - \mathsf{E}[A(1)^2]\big), & x < \alpha; \\ \frac{1}{4}(x+\alpha)^2 + \big(x^2 - \mathsf{E}[\{x - \frac{1}{2}(x+\alpha) + A(1)\}^2]\big), & x \ge \alpha. \end{cases}$$

To ensure positivity we can take $\eta^* = \mathsf{E}[A(1)^2]$.

It is not worth expanding out this expression for c, but we can argue that c is approximated by the quadratic $\frac{1}{2}x^2$ for large x:

$$\lim_{x\to\infty} \frac{c(x)}{x^2} = \lim_{x\to\infty} \frac{\frac{1}{4}x^2 + (x^2 - (\frac{1}{2}x)^2)}{x^2} = \tfrac{1}{2}$$

If our goal was to minimize the average cost with $c(x,u) = \frac{1}{2}x^2 + u^2$, then this would likely provide a good approximate solution.