ECE 555 Control of Stochastic Systems Fall 2005

Handout: ***Control Variates in Simulation***

In the past few lectures we have considered the general stochastic approximation recursion,

$$\theta(k+1) = \theta(k) + a_k[g(\theta(k)) + \Delta(k+1)], \qquad k \geq 0.$$

Under general conditions, verified by considering various ODEs, it is known that $\{\theta(k)\}$ converges to the set of zeros of g.

The remaining problem is that *convergence can be very slow.* These notes summarize the control variate method for speeding convergence in simulation. It is highly likely that this technique can be generalized to other recursive algorithms.

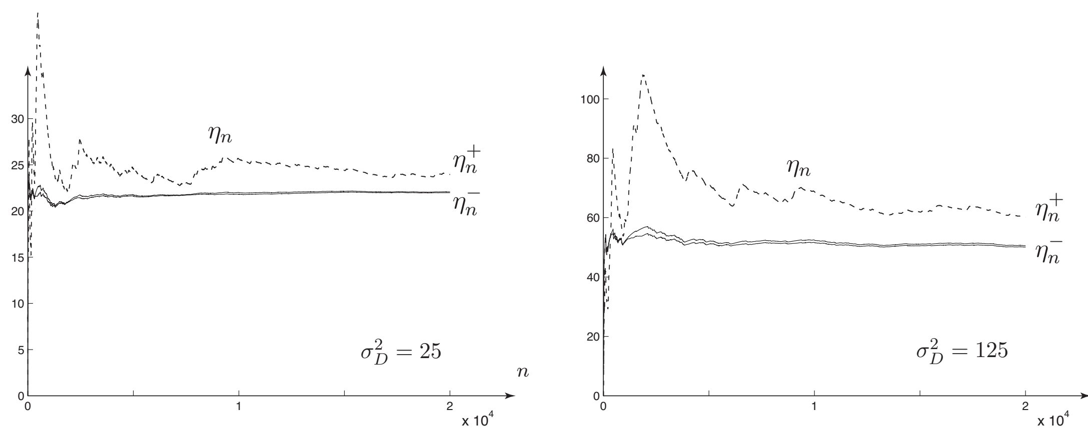

Figure 1: Simulation using the standard estimator, and the two controlled estimators. The plot at left shows results with $\sigma_D^2 = 25$, and at right the variance is increased to $\sigma_D^2 = 125$. In each case the estimates obtained from the standard Monte-Carlo estimator are significantly larger than those obtained using the controlled estimator, and the bound $\eta_n^- < \eta_n^+$ holds for all large n.

Simulating a Markov Chain Suppose that $\boldsymbol{X}$ is a Markov chain on a state space X with invariant distribution π. For background see [8] (as well as [10, 3, 8, 4].)

For a given function $F\colon \mathsf{X} \to \mathbb{R}$ we denote,

$$L_n(F) := \frac{1}{n}\sum_{k=0}^{n-1} F(X(k)) \qquad n \geq 1\,.$$

One can hope to establish the following limit theorems,

The Strong Law of Large Numbers, or SLLN: For each initial condition,

$$L_n(F) \to \pi(F), \qquad a.s.,\ n \to \infty. \tag{1}$$

The Central Limit Theorem, or CLT: For some $\sigma \geq 0$ and each initial condition,

$$\sqrt{n}[L_n(F) - \eta] \xrightarrow{\mathrm{w}} \sigma W, \qquad n \to \infty, \tag{2}$$

where W is a standard normal random variable, and the convergence is in distribution.

It is assumed here that the chain is *ergodic*, which means that the SLLN holds for any bounded function $F\colon \mathsf{X} \to \mathbb{R}$.

Suppose that $F\colon \mathsf{X} \to \mathbb{R}$ is a π-integrable function. Under ergodicity the SLLN can be generalized to any such function. Our interest is to efficiently estimate the finite mean $\eta = \pi(F)$. The standard estimator is the sample path average,

$$\eta_n = L_n(F) \qquad n \geq 1\,. \tag{3}$$

Its performance is typically gauged by the associated asymptotic variance σ^2 used in (2). Below are two well known representations in terms of the centered function $\widetilde{F} := F - \eta$.

Limiting variance:

$$\sigma^2 = \lim_{n\to\infty} n\mathsf{Var}_x(L_n(F)) := \lim_{n\to\infty} \mathsf{E}_x\bigl[L_n(\widetilde{F})^2\bigr] \tag{4}$$

Sum of the correlation function:

$$\sigma^2 = \sum_{k=-\infty}^{\infty} \mathsf{E}_\pi\bigl[\widetilde{F}(X(k))\widetilde{F}(X(0))\bigr] \tag{5}$$

The following operator-theoretic representation holds more generally. Let Z denote a version of the *fundamental kernel*, defined so that $\widehat{F} = ZF$ solves Poisson's equation for some class of functions F,

$$P\widehat{F} = \widehat{F} - F + \eta. \tag{6}$$

It will be convenient to apply the following bilinear and quadratic forms, defined for measurable functions $F, G\colon \mathsf{X} \to \mathbb{R}$,

$$\langle\!\langle F, G\rangle\!\rangle := P(FG) - (PF)(PG), \qquad \mathcal{Q}(F) := \langle\!\langle F, F\rangle\!\rangle.$$

Using this notation we have the following representation for the asymptotic variance,

$$\sigma^2(F) = \pi(\mathcal{Q}(\widehat{F})). \tag{7}$$

Recall that the resolvent is expressed $R := \sum_0^\infty 2^{-n-1}P^n$. The function $s : \mathsf{X} \to (0,1]$ and the probability measure ν are called *small* if the *minorization condition* holds,

$$R(x, A) \geq s(x)\nu(A), \qquad x \in \mathsf{X},\ A \in \mathcal{B}(\mathsf{X}).$$

The following is the general state space version of Condition (V3):

For functions $V : \mathsf{X} \to (0,\infty]$, $f : \mathsf{X} \to [1,\infty)$, a small function s, a small measure ν, and a constant $b < \infty$,
$$\mathcal{D}V \leq -f + bs \tag{V3}$$

The following result is taken from [8, 6]:

Proposition. Suppose that $\boldsymbol{X}$ satisfies (V3) with $\pi(V^2) < \infty$. Then, the SLLN and CLT hold for any $F \in L_\infty^f$, and the asymptotic variance $\sigma^2(F)$ exists, and can be expressed as (4), (5), or (7) above. □

Control-variates The purpose of the control-variate method is to reduce the variance of the standard estimator (3). See [7, 9, 2, 1] for background on the general control-variate method.

Suppose that $H\colon \mathsf{X} \to \mathbb{R}$ is a π-integrable function with known mean, and finite asymptotic variance. By normalization we can assume that $\pi(H) = 0$. Then, for a given $\vartheta \in \mathbb{R}$ and with $F_\vartheta := F - \vartheta H$, the sequence $\{L_n(F_\vartheta)\}$ provides an asymptotically unbiased estimator of $\pi(F)$. The asymptotic variance of the controlled estimator is given by

$$\sigma^2(F_\vartheta) = \mathcal{Q}(\widehat{F}_\vartheta) = \pi\big(\langle\!\langle ZF, ZF\rangle\!\rangle - 2\vartheta\langle\!\langle ZF, ZH\rangle\!\rangle + \vartheta^2\langle\!\langle ZH, ZH\rangle\!\rangle\big).$$

Minimizing over $\vartheta \in \mathbb{R}$ gives the estimator with minimal asymptotic variance,

$$\vartheta^* = \frac{\pi(\langle\!\langle ZF, ZH\rangle\!\rangle)}{\pi(\langle\!\langle ZH, ZH\rangle\!\rangle)}.$$

For a Markov chain it is easy to construct a function with zero mean: consider $H = J - PJ$ where J is known to have finite mean. Our goal then is to choose J so that it approximates the solution to Poisson's equation (6): The idea is that if $J = \widehat{F}$, then the resulting controlled estimator with $\vartheta = 1$ has *zero* asymptotic variance. This approach has been successfully applied in queueing models by taking J equal to the associated fluid value function described in lecture.

Consider the simple reflected random walk on $\mathbb{R}_+$, defined by the recursion

$$X(k+1) = [X(k) + D(k+1)]_+, \qquad k \geq 0, \tag{8}$$

with $[x]_+ = \max(x, 0)$ for $x \in \mathbb{R}$, and $\boldsymbol{D}$ i.i.d.. The fluid model is given by,

$$q(t) = [q(0) - \delta]_+, \qquad t \geq 0,$$

where $-\delta = \mathsf{E}[D(k)]$ is assumed to be negative. The fluid value function is the quadratic,

$$J(x) = \int_0^\infty q(t)\, dt = \tfrac{1}{2}\delta^{-1}x^2, \qquad x = q(0) \in \mathbb{R}_+.$$

Consider the special case in which $\boldsymbol{D}$ has common marginal distribution,

$$D(k) = \begin{cases} 1 & \text{with probability } \alpha; \\ -1 & \text{with probability } 1 - \alpha. \end{cases}$$

The Markov chain $\boldsymbol{X}$ is then a discrete-time model of the M/M/1 queue with state space $\mathsf{X} = \mathbb{Z}_+$. When $F(x) \equiv x$ we have seen that $\widehat{F}(x) = \frac{1}{2}\delta^{-1}(x^2 + x)$, so that the error $\widehat{F} - J$ is linear in x. Moreover, the representation (7) can be written,

$$\sigma^2(F) = \pi(\mathcal{Q}(\widehat{F})) = 2\pi(\widetilde{F}\widehat{F}) - \pi(\widetilde{F}^2) = \mathsf{E}[\tfrac{1}{2}\delta^{-1}\widetilde{X}^3 - \widetilde{X}^2]$$

which grows like δ^{-4} as $\delta \downarrow 0$ (equivalently, $\rho \uparrow 1$.)

Returning to the random walk (8), consider the following special case in which the sequence $\boldsymbol{D}$ is of the form $D(k) = A(k) - S(k)$, where $\boldsymbol{A}$ and $\boldsymbol{S}$ are mutually independent, i.i.d. sequences, with mean α, μ respectively. We let $\kappa > 0$ denote a variability parameter, and define

$$\begin{aligned} &\mathsf{P}\{S(k) = (1+\kappa)\mu\} = 1 - \mathsf{P}\{S(k) = 0\} = (1+\kappa)^{-1} \\ &\mathsf{P}\{A(k) = (1+\kappa)\alpha\} = 1 - \mathsf{P}\{A(k) = 0\} = (1+\kappa)^{-1} \end{aligned}$$

Consequently, we have $-\delta = \mathsf{E}[A(k)] - \mathsf{E}[S(k)] = -(\mu - \alpha)$, and $\sigma_D^2 = \sigma_A^2 + \sigma_S^2 = (\mu^2 + \alpha^2)\kappa$.

The simulation results shown use $\mu = 4$ and $\alpha = 3$, so that $\delta = 1$. Two estimators $\{\eta_n^-, \eta_n^+\}$ were constructed based on the parameter values $\vartheta_- = 1.05$ and $\vartheta_+ = 1$. The plot at left in Figure 1 illustrates the resulting performance with $\kappa = 2$ ($\sigma_D^2 = 25$), and the plot at right shows the controlled and uncontrolled estimators with $\kappa = 5$, and hence $\sigma_D^2 = 125$.

Note that the bounds $\eta_n^- < \eta_n^+ < \eta_n$ hold for all large n, even though all three estimators are asymptotically unbiased.

A network model The *Kumar-Seidman-Rybko-Stolyar* (KSRS) network shown in Figure 2 is described in Chapter 1 of the course notes.

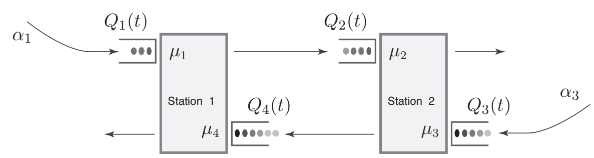

Figure 2: The Kumar-Seidman-Rybko-Stolyar (KSRS) network.

Consider the following policy based on a vector $\bar{w} \in \mathbb{R}_+^2$ of *safety-stock* values: Serve $Q_1 \geq 1$ at Station I if and only if $Q_4 = 0$, or

$$\mu_2^{-1} Q_2 + \mu_3^{-1} Q_3 \leq \bar{w}_2. \tag{9}$$

An analogous condition holds at Station II.

A simulation experiment was conducted to estimate the steady-state mean customer population. So, with $\mathsf{X} = \mathbb{Z}_+^4$, we let $F \colon \mathsf{X} \to \mathbb{R}_+$ denote the ℓ_1 norm on $\mathbb{R}^4$. A CRW network model was constructed in which the elements of $(\boldsymbol{A}, \boldsymbol{S})$ were taken Bernoulli (see course lecture notes.) Details can be found in [5].

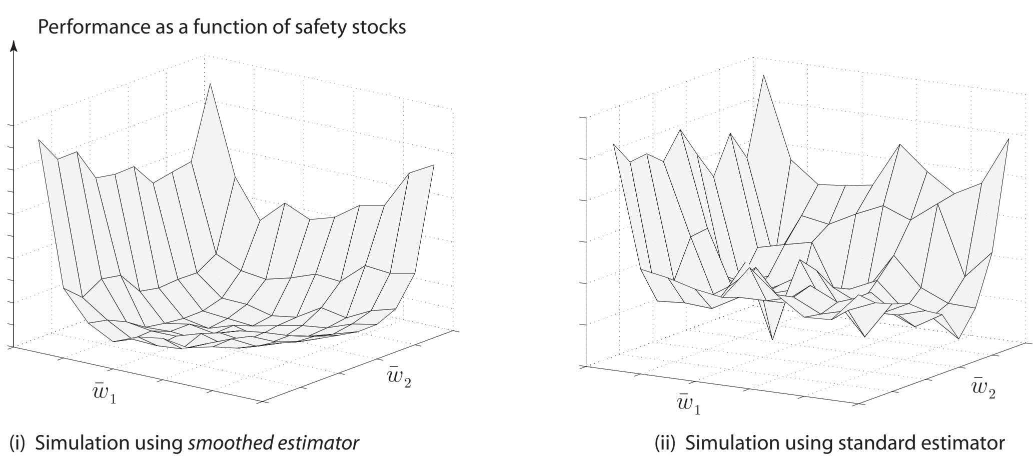

Figure 3: Estimates of the steady-state customer population in the KSRS model as a function of 100 different safety-stock levels using the policy (9). Two simulation experiments are shown, where in each case the simulation runlength consisted of $N = 200,000$ steps. The left hand side shows the results obtained using the smoothed estimator; the right hand side shows results with the standard estimator.

Shown in Figure 3 are estimates of the steady-state customer population in Case I for the family of policies (9), indexed by the safety-stock level $\bar{w} \in \mathbb{R}_+^2$. Shown at left are estimates obtained using the "smoothed estimator" based on a fluid value function. The plot at right shows estimates obtained using the standard estimator.

References

[1] P. Glasserman. *Monte Carlo methods in financial engineering.* Applications of Mathematics, 53. Springer, York, NY, 2004.

[2] P. Glynn and R. Szechtman. Some new perspectives on the method of control variates. In K.T. Fang, F.J. Hickernell, and H. Niederreiter, editors, *Monte Carlo and Quasi-Monte Carlo Methods 2000: Proceedings of a Conference held at Hong Kong Baptist University, Hong Kong SAR, China*, pages 27–49, Berlin, 2002. Springer-Verlag.

[3] P. W. Glynn and S. P. Meyn. A Liapounov bound for solutions of the Poisson equation. *Ann. Probab.*, 24(2):916–931, 1996.

[4] P. W. Glynn and W. Whitt. Necessary conditions for limit theorems in cumulative processes. *Stoch. Proc. Applns.*, 98:199–209, 2002.

[5] S. G. Henderson, S. P. Meyn, and V. B. Tadić. Performance evaluation and policy selection in multiclass networks. *Discrete Event Dynamic Systems: Theory and Applications*, 13(1-2):149–189, 2003. Special issue on learning, optimization and decision making (invited).

[6] I. Kontoyiannis and S. P. Meyn. Spectral theory and limit theorems for geometrically ergodic Markov processes. *Ann. Appl. Probab.*, 13:304–362, 2003. Presented at the INFORMS Applied Probability Conference, NYC, July, 2001.

[7] A. M. Law and W. D. Kelton. *Simulation Modeling and Analysis.* McGraw-Hill, New York, 3rd edition, 2000.

[8] S. P. Meyn and R. L. Tweedie. *Markov Chains and Stochastic Stability.* Springer-Verlag, London, 1993. online: `http://black.csl.uiuc.edu/~meyn/pages/book.html`.

[9] B. L. Nelson. Control-variate remedies. *Operations Res.*, 38(4):974–992, 1990.

[10] E. Nummelin. *General Irreducible Markov Chains and Nonnegative Operators.* Cambridge University Press, Cambridge, 1984.