Recent Developments in Randomized MLMC

Peter W. Glynn
Stanford University

Joint work with Chang-Han Rhee, Zeyu Zheng
and Jose Blanchet

International Conference on Monte Carlo Techniques
Paris, July 5-8, 2016

General Setting:

- Compute $\alpha = EY_\infty$
- Given Y_n such that
  - $EY_n \to EY_\infty$
  - Y_n's are coupled so that $||Y_n - Y_\infty||_2^2 \to 0$ as $n \to \infty$

  e.g. $Y_n \to Y_\infty$ in L^2
- Success depends on the quality of coupling

Some Applications

- Stochastic differential equations (SDE's)
- Partial differential equations with random coefficients
- Many others

Today's talk: Focus on Markov chains

Markov Chains

$X = (X_n : n \geq 0)$ $S-$valued Markov chain

- $P(x, dy) \triangleq P(X_{n+1} \in dy \mid X_n = x)$
- $P_x(\cdot) \triangleq P(\cdot \mid X_0 = x)$
- $E_x(\cdot) \triangleq E(\cdot \mid X_0 = x)$

We will focus on computations related to equilibrium properties, so we will assume that X has a unique stationary distribution π

$$\pi(\cdot) \triangleq P(X_\infty \in \cdot)$$

Goal 1: Compute $\alpha = E(Y_\infty)$ where $Y_\infty = r(X_\infty)$ $(r : S \to \mathbb{R}_+)$

The Relative Value Function

In approximate dynamic programming, the goal is to compute the solution of the HJB equation

$$\alpha + v(x) = \max_{a \in \mathcal{A}(x)} \left[r(x,a) + \int_S P_{a(x)}(x,dy)v(y) \right]$$

The value function satisfies *Poisson's equation* under the optimal a^*,

$$v(x) = r(x, a^*(x)) - \alpha + \int_S P_{a^*(x)}(x,dy)v(y)$$

Under suitable regularity conditions,

$$v(x) = \sum_{n=0}^{\infty} (E_x^* r(X_n) - E^* r(X_\infty))$$

<u>Goal 2:</u> Compute $v(x)$ efficiently

The Time-average Variance Constant (TAVC)

In approximating the *distribution* of total reward over n periods, we take advantage of the central limit theorem:

$$\sum_{j=0}^{n-1} r(X_j) \overset{D}{\approx} nEr(X_\infty) + \sqrt{n}\sigma N(0,1)$$

<u>Goal 3:</u> Compute the TAVC σ^2 efficiently

- The quantity σ^2 also plays a role in computing confidence intervals for $Er(X_\infty)$ when using the sample mean $n^{-1}\sum_{j=0}^{n-1} r(X_j)$.
- The quantity σ^2 is a special case of the spectral density

$$f(\lambda) = \frac{1}{2\pi}\sum_{j=-\infty}^{\infty} e^{i\lambda j}\mathrm{cov}_\pi(r(X_0), r(X_j))$$

In particular,

$$\sigma^2 = 2\pi f(0)$$

The Gradient of an Equilibrium Expectation

Suppose that the transition kernel P depends on a parameter θ

$$P = P(\theta)$$

Goal 4: Compute $\frac{d}{d\theta} E_\theta r(X_\infty)$ efficiently

- Arises in statistical settings with uncertain input parameters

$$E_{\hat{\theta}} r(X_\infty) \overset{D}{\approx} E_{\theta^*} r(X_\infty) + \frac{d}{d\theta} E_{\theta^*} r(X_\infty)(\hat{\theta} - \theta^*)$$

- Needed in numerical algorithms for optimizing steady-state reward over decision parameter space

A Simple Randomization Idea

Goal: Compute EY, where Y is difficult / impossible to generate exactly

Suppose that we have a sequence $(Y_m : m \geq 0)$ for which

- Y_m can be generated exactly
- $Y_m \stackrel{L^2}{\rightarrow} Y$ as $m \rightarrow \infty$

Then:

$$Y = \lim_{m\to\infty} \big(\underbrace{Y_0}_{\Delta_0} + \sum_{i=1}^{m} (\underbrace{Y_i - Y_{i-1}}_{\Delta_i})\big) = \sum_{i=0}^{\infty} \Delta_i.$$

If

$$\sum_{i=1}^{\infty} \mathbf{E}|\Delta_i| < \infty,$$

then

$$\mathbf{E}Y = \sum_{i=1}^{\infty} \mathbf{E}\Delta_i = \sum_{i=1}^{\infty} \mathbf{E}\Delta_i \frac{\mathbf{P}(N \geq i)}{\mathbf{P}(N \geq i)} = \mathbf{E}\sum_{i=1}^{\infty} \Delta_i \frac{\mathbb{I}(N \geq i)}{\mathbf{P}(N \geq i)} = \mathbf{E}\sum_{i=1}^{N} \frac{\Delta_i}{\mathbf{P}(N \geq i)} \triangleq \mathbf{E}Z.$$

Note that Z is an unbiased estimator of $\mathbf{E}Y$. We call Z the *summed* estimator.

Then:

$$Y = \lim_{m\to\infty}\Big(\underbrace{Y_0}_{\Delta_0} + \sum_{i=1}^{m}(\underbrace{Y_i - Y_{i-1}}_{\Delta_i})\Big) = \sum_{i=0}^{\infty}\Delta_i.$$

If

$$\sum_{i=1}^{\infty}\mathbf{E}|\Delta_i| < \infty,$$

then

$$\mathbf{E}Y = \sum_{i=1}^{\infty}\mathbf{E}\Delta_i = \sum_{i=1}^{\infty}\mathbf{E}\Delta_i\frac{\mathbf{P}(N\geq i)}{\mathbf{P}(N\geq i)} = \mathbf{E}\sum_{i=1}^{\infty}\Delta_i\frac{\mathbb{I}(N\geq i)}{\mathbf{P}(N\geq i)} = \mathbf{E}\sum_{i=1}^{N}\frac{\Delta_i}{\mathbf{P}(N\geq i)} \triangleq \mathbf{E}Z.$$

Note that Z is an unbiased estimator of $\mathbf{E}Y$. We call Z the *summed* estimator.

Then:

$$Y = \lim_{m\to\infty} \big(\underbrace{Y_0}_{\Delta_0} + \sum_{i=1}^{m} (\underbrace{Y_i - Y_{i-1}}_{\Delta_i})\big) = \sum_{i=0}^{\infty} \Delta_i.$$

If

$$\sum_{i=1}^{\infty} \mathbf{E}|\Delta_i| < \infty,$$

then

$$\mathbf{E}Y = \sum_{i=1}^{\infty} \mathbf{E}\Delta_i = \sum_{i=1}^{\infty} \mathbf{E}\Delta_i \frac{\mathbf{P}(N \geq i)}{\mathbf{P}(N \geq i)} = \mathbf{E}\sum_{i=1}^{\infty} \Delta_i \frac{\mathbb{I}(N \geq i)}{\mathbf{P}(N \geq i)} = \mathbf{E}\sum_{i=1}^{N} \frac{\Delta_i}{\mathbf{P}(N \geq i)} \triangleq \mathbf{E}Z.$$

Note that Z is an unbiased estimator of $\mathbf{E}Y$. We call Z the *summed* estimator.

Then:

$$Y = \lim_{m \to \infty} \Big(\underbrace{Y_0}_{\Delta_0} + \sum_{i=1}^{m} \underbrace{(Y_i - Y_{i-1})}_{\Delta_i} \Big) = \sum_{i=0}^{\infty} \Delta_i.$$

If

$$\sum_{i=1}^{\infty} \mathbf{E}|\Delta_i| < \infty,$$

then

$$\mathbf{E}Y = \sum_{i=1}^{\infty} \mathbf{E}\Delta_i = \sum_{i=1}^{\infty} \mathbf{E}\Delta_i \frac{\mathbf{P}(N \geq i)}{\mathbf{P}(N \geq i)} = \mathbf{E} \sum_{i=1}^{\infty} \Delta_i \frac{\mathbb{I}(N \geq i)}{\mathbf{P}(N \geq i)} = \mathbf{E} \sum_{i=1}^{N} \frac{\Delta_i}{\mathbf{P}(N \geq i)} \triangleq \mathbf{E}Z.$$

Note that Z is an unbiased estimator of $\mathbf{E}Y$. We call Z the *summed* estimator.

Then:

$$Y = \lim_{m\to\infty} \big(\underbrace{Y_0}_{\Delta_0} + \sum_{i=1}^{m} (\underbrace{Y_i - Y_{i-1}}_{\Delta_i})\big) = \sum_{i=0}^{\infty} \Delta_i.$$

If

$$\sum_{i=1}^{\infty} \mathbf{E}|\Delta_i| < \infty,$$

then

$$\mathbf{E}Y = \sum_{i=1}^{\infty} \mathbf{E}\Delta_i = \sum_{i=1}^{\infty} \mathbf{E}\Delta_i \frac{\mathbf{P}(N \geq i)}{\mathbf{P}(N \geq i)} = \mathbf{E} \sum_{i=1}^{\infty} \Delta_i \frac{\mathbb{I}(N \geq i)}{\mathbf{P}(N \geq i)} = \mathbf{E} \sum_{i=1}^{N} \frac{\Delta_i}{\mathbf{P}(N \geq i)} \triangleq \mathbf{E}Z.$$

Note that Z is an unbiased estimator of $\mathbf{E}Y$. We call Z the *summed* estimator.

Then:

$$Y = \lim_{m\to\infty}\Big(\underbrace{Y_0}_{\Delta_0} + \sum_{i=1}^{m}\underbrace{(Y_i - Y_{i-1})}_{\Delta_i}\Big) = \sum_{i=0}^{\infty}\Delta_i.$$

If

$$\sum_{i=1}^{\infty}\mathbf{E}|\Delta_i| < \infty,$$

then

$$\mathbf{E}Y = \sum_{i=1}^{\infty}\mathbf{E}\Delta_i = \sum_{i=1}^{\infty}\mathbf{E}\Delta_i\frac{\mathbf{P}(N\geq i)}{\mathbf{P}(N\geq i)} = \mathbf{E}\sum_{i=1}^{\infty}\Delta_i\frac{\mathbb{I}(N\geq i)}{\mathbf{P}(N\geq i)} = \mathbf{E}\sum_{i=1}^{N}\frac{\Delta_i}{\mathbf{P}(N\geq i)} \triangleq \mathbf{E}Z.$$

Note that Z is an unbiased estimator of $\mathbf{E}Y$. We call Z the *summed* estimator.

Then:

$$Y = \lim_{m\to\infty}\big(\underbrace{Y_0}_{\Delta_0} + \sum_{i=1}^{m}(\underbrace{Y_i - Y_{i-1}}_{\Delta_i})\big) = \sum_{i=0}^{\infty}\Delta_i.$$

If

$$\sum_{i=1}^{\infty}\mathbf{E}|\Delta_i| < \infty,$$

then

$$\mathbf{E}Y = \sum_{i=1}^{\infty}\mathbf{E}\Delta_i = \sum_{i=1}^{\infty}\mathbf{E}\Delta_i\frac{\mathbf{P}(N\geq i)}{\mathbf{P}(N\geq i)} = \mathbf{E}\sum_{i=1}^{\infty}\Delta_i\frac{\mathbb{I}(N\geq i)}{\mathbf{P}(N\geq i)} = \mathbf{E}\sum_{i=1}^{N}\frac{\Delta_i}{\mathbf{P}(N\geq i)} \triangleq \mathbf{E}Z.$$

Note that Z is an unbiased estimator of $\mathbf{E}Y$. We call Z the *summed* estimator.

We want:

$$c^{1/2}(\widehat{Y}(c) - \mathbf{E}Y) \Rightarrow \eta N(0,1)$$

as $c \to \infty$.

For this, we need:

- $\mathrm{var}(Z) < \infty$
- $\mathbf{E}(\text{time to generate } Z) < \infty$

Then,

$$c^{1/2}(\widehat{Y}(c) - \mathbf{E}Y) \Rightarrow \sqrt{\mathrm{var}(Z) \cdot \mathbf{E}(\text{time to generate } Z)}\, N(0,1)$$

as $c \to \infty$

Variance of Z:

$$\mathrm{var}(Z) = \sum_{i=0}^{\infty} \gamma_i / \mathbf{P}(N \geq i)$$

where $\gamma_i = O(\|Y - Y_i\|_2^2)$

Mean time to generate Z:

$$\mathbf{E}(\text{time to generate } Z) = \mathbf{E}\left(\sum_{i=0}^{N} t_i\right)$$
$$= \sum_{i=0}^{\infty} t_i \, \mathbf{P}(N \geq i)$$

where $t_i =$ (incremental) time to generate Y_i

Rhee and G (2012, 2013)
McLeish (2011)

Single-Term Estimator

Note that

$$\begin{aligned}\mathbf{E}Y &= \sum_{i=1}^{\infty} \mathbf{E}\Delta_i = \sum_{i=1}^{\infty} \mathbf{E}\Delta_i \frac{\mathbf{P}(N=i)}{\mathbf{P}(N=i)} = \mathbf{E}\sum_{i=1}^{\infty} \Delta_i \frac{\mathbb{I}(N=i)}{\mathbf{P}(N=i)} \\ &= \mathbf{E}\Delta_N / p_N \\ &\triangleq \mathbf{E}\check{Z} \qquad \text{"single-term estimator"}\end{aligned}$$

where $p_n = \mathbf{P}(N=n)$. Note that $\check{Z}$ is also an unbiased estimator of $\mathbf{E}Y$

Optimal Choice of Randomization Distribution N for Summed Estimator

Find N that minimizes the work-variance product:

$$\begin{aligned} \text{minimize} \quad & g(\bar{F}) = \left(\sum_{i=0}^{\infty} t_i \overline{F}_i\right)\left(\sum_{i=0}^{\infty} \gamma_i / \overline{F}_i\right) \\ \text{subject to} \quad & \overline{F}_n \geq \overline{F}_{n+1}, \quad \forall n \geq 0 \\ & \overline{F}_n > 0, \quad \forall n \geq 0 \\ & \overline{F}_0 = 1 \end{aligned}$$

where $\overline{F}_i = \mathbf{P}(N \geq i)$.

If $\sqrt{\gamma_i / t_i}$ is well-defined and non-increasing:

$$\overline{F}_i^* = \frac{\sqrt{\gamma_i / t_i}}{\sqrt{\gamma_0 / t_0}}$$

McLeish (2011)
Rhee and G (2012)

Optimal Choice of Randomization Distribution N for Single-term Estimator

Theorem

Optimal N for the single-term estimator $\check{Z}$ has the probability mass function $(p_n^* : n \geq 0)$, where

$$p_n^* = \sqrt{\frac{\mathbf{E}\Delta_n^2}{\alpha^2 + c^* t_n}}$$

for $n \geq 0$. Here c^* is the unique root of the equation

$$\sum_{n=0}^{\infty} \sqrt{\frac{\mathbf{E}\Delta_n^2}{\alpha^2 + c^* t_n}} = 1.$$

The Connection to MLMC

- Our summed estimator takes the form

$$\sum_{i=1}^{\infty} \hat{\Delta}_i \frac{\hat{P}(N \geq i)}{P(N \geq i)},$$

where $\hat{\Delta}_i = \frac{1}{\hat{m}_i} \sum_{j=1}^{\hat{m}_i} \Delta_{ji} \quad (\hat{m}_1 \geq \hat{m}_2 \geq \cdots)$

- We can *post-stratify* this estimator:

$$\sum_{i=1}^{\infty} \hat{\Delta}_i (\triangleq \hat{W})$$

- This post-stratified estimator reduces the variance

$$\text{var}(\hat{W}) \approx \frac{1}{n} \frac{\sum_{i=1}^{\infty} [\text{var}\Delta_i + 2\text{cov}(\Delta_i, Y_\infty - Y_i)]}{P(N \geq i)}$$

If we further stratify the $\hat{\Delta}_i$'s:

Replace $\hat{\Delta}_i$ by $\frac{1}{\lfloor nP(N\geq i)\rfloor}\sum_{j=1}^{\lfloor nP(N\geq i)\rfloor}\Delta_{ji}$

we effectively end up with MLMC (see Vihola (2015))

- This last step does not reduce asymptotic variance

Goal 1: Compute $Er(X_\infty)$ efficiently

- The natural approximation to $Y_\infty = r(X_\infty)$ is $Y_n = r(X_n)$
- But the $r(X_n)$'s do not couple in the sense that $||r(X_n) - r(X_\infty)||_2^2 \to 0$ as $n \to \infty$

Random Iterated Function Perspective

Set

$$X_{n+1} = f(X_n, Z_{n+1})$$

where the Z_i's are iid. View as

$$X_{n+1} = \phi_{n+1}(X_n)$$

where the ϕ_i's are iid random functions. So,

$$X_n = (\phi_n \circ \phi_{n-1} \cdots \circ \phi_1)(X_0).$$

But since the ϕ_i's are iid,

$$\phi_n \circ \cdots \circ \phi_1 \stackrel{D}{=} \phi_1 \circ \cdots \circ \phi_n$$

so $\tilde{X}_n \stackrel{D}{=} X_n$, where

$$\tilde{X}_n = (\phi_1 \circ \phi_2 \cdots \circ \phi_n)(X_0).$$

Set $\tilde{Y}_n = r(\tilde{X}_n)$. Fortunately, the $\tilde{Y}_n$'s often couple nicely, in the sense that

$$||\tilde{Y}_n - \tilde{Y}_{n-1}||_2^2 \to 0 \quad \text{as } n \to \infty$$

Setting 1: Contraction Chains

If:

- φ_i is contracting on average

$$\sup_{x \neq y} \mathbf{E} \frac{|\varphi_i(x) - \varphi_i(y)|^2}{|x-y|^2} \leq \rho < 1$$

e.g. $\quad \varphi_i(x) = A_i x + Z_i \qquad (0 < \mathbf{E}|A_i|^2 < 1)$

- $\mathbf{E}|x_0 - \varphi_i(x_0)|^2 < \infty$, for some x_0

Then,

$$\mathbf{E}_x|\widetilde{X}_m - \widetilde{X}_\infty|^2 \leq \gamma_x \rho^n, \qquad n \geq 0$$

Setting 2: Monotone Chains

Suppose that $\phi_i(\cdot)$ is non-decreasing in $\preceq$ and S has a least element (0, say) under $\preceq$. Then,

$$(\phi_1 \circ \cdots \circ \phi_n)(0) \nearrow$$

in $\preceq$. So, if r is increasing in $\preceq$,

$$\tilde{Y}_n \nearrow \tilde{Y}_\infty \quad a.s.$$

Numerical Examples

Table: $\alpha = \mathbf{E}f(W_\infty)$, $W_{n+1} = [W_n + X_n]^+$, $X_i \sim N(-0.5, 1)$, $f(x) = x \wedge 1$, 100 Samples

IRE	90% Confidence Interval	RMSE/α	Work	Work$\times$MSE
0.200	0.3152012 ± 0.0100708	1.90×10^{-1}	1.9×10^4	72.0
0.100	0.3283592 ± 0.0045089	8.43×10^{-2}	8.9×10^4	67.5
0.050	0.3244757 ± 0.0024279	4.54×10^{-2}	3.3×10^5	73.3
0.020	0.3248810 ± 0.0011379	2.14×10^{-2}	2.0×10^6	99.6
0.010	0.3265324 ± 0.0004189	8.00×10^{-3}	9.5×10^6	64.5
0.005	0.3257757 ± 0.0002199	4.13×10^{-3}	3.4×10^7	60.8

Table: $\alpha = \mathbf{E}f(X_\infty)$, $X_{n+1} = \frac{1}{2}X_n + \xi_n$, $\xi_i \sim N(1, 1)$, $f(x) = x$, 100 Samples

IRE	90% Confidence Interval	RMSE/α	Work	Work$\times$MSE
0.050	2.0022907 ± 0.0116508	3.54×10^{-2}	2.6×10^3	13.2
0.020	1.9925651 ± 0.0062383	1.93×10^{-2}	1.3×10^4	19.9
0.010	1.9999375 ± 0.0028700	8.71×10^{-3}	6.7×10^4	20.3
0.005	2.0005492 ± 0.0015075	4.58×10^{-3}	2.3×10^5	19.2
0.002	2.0000365 ± 0.0005775	1.75×10^{-3}	1.7×10^6	21.3
0.001	1.9999333 ± 0.0002688	8.16×10^{-4}	5.9×10^6	15.6

Unbiased Estimation for Equilibrium Expectations

- Use the above coupling in our randomized MLMC methodology
- We call this class of algorithms *exact estimation* algorithms, to be contrasted with *exact simulation* algorithms (like Propp-Wilson)
- Exact estimation is possible in many settings where exact simulation is hard or impossible (G and Rhee (2014))

Another Approach:

If X exhibits regenerative structure,

$$Er(X_\infty) = \frac{E\sum_{j=0}^{\tau-1} r(X_j)}{E\tau}$$

where $\tau =$ regeneration time. So,

$$Er(X_\infty) = g(E\sum_{j=0}^{\tau-1} r(X_j), E\tau)$$

where $g(x, y) = x/y$. Our randomized MLMC scheme can be extended to obtain unbiased estimators of such ratio quantities

Blanchet and G (2015)

Computing the Relative Value Function

Consider the relative value function:

$$v(x) = \sum_{j=0}^{\infty} [E_x(r(X_n) - Er(X_\infty)]$$

Conventional estimator must trade-off bias versus variance, leading to a sub-canonical (slower than $c^{-1/2}$) convergence rate

The Multi-level Alternative:

- Set $b_n(x) = \mathbf{E}_x r(X_n)$ and note that if

$$y_n \triangleq \sum_{j=0}^{n-1} (j+1)(b_j(x) - b_{j+1}(x)) = \sum_{j=0}^{n-1} b_j(x) - nb_n(x) \to h(x)$$

exponentially rapidly, when X is contractive. Also, $y_k = E\tilde{Y}_k$, where

$$\tilde{Y}_k = \sum_{j=0}^{k-1} (j+1) \left[r((\phi_0 \circ \cdots \circ \phi_{-j+1})(x)) - r((\phi_0 \circ \cdots \circ \phi_{-j})(x)) \right],$$

$\tilde{Y}_0 = r(x)$ and $||\tilde{Y}_n - \tilde{Y}_{n-r}||_2 = O(\beta^{n-r})$ for $0 < \beta < 1$.

- Leads to a square root convergent estimator (Zheng and G (2016))

A Numerical Example

Autoregressive model of order 1:

$$X_{n+1} = 0.5 * X_n + Z_{n+1}$$

where Z_i's are iid $N(0,1)$; $r_1(x) = x$ and $r_2(x) = x^2$.

Table: $P(N \geq n) = 0.95^n$, 10^5 # samples, 100 replications

$r(\cdot)$	x	$P(X_\infty > x)$	True $h(x)$	Est $h(x)$	Emp Bias	90% C.I.
r_1	0	0.5	0	0.0001	1.3×10^{-4}	[-0.0251, 0.0311]
	0.7788	0.75	1.5577	1.5568	-8.2×10^{-4}	[1.5250, 1.5859]
	1.8993	0.95	3.7986	3.7983	-3.7×10^{-4}	[3.7690, 3.8356]
	2.6862	0.99	5.3725	5.3698	-2.7×10^{-3}	[5.3331, 5.4085]
r_2	0	0.5	-1.778	-1.7829	-5.1×10^{-3}	[-1.8331, -1.7320]
	0.7788	0.75	-0.9690	-0.9696	-5.8×10^{-4}	[-1.0360, -0.9028]
	1.8993	0.95	3.0321	3.0460	-1.4×10^{-3}	[2.9418, 3.1266]
	2.6862	0.99	7.8434	7.8413	-2.1×10^{-3}	[7.7170, 7.9729]

Est $h(x)$ denotes the averaged estimations; Emp Bias denotes Empirical Bias

A Numerical Example: Square-root Convergence Rate

Table: Rate of Convergence

# of Samples	True $h(x)$	90% Confidence Interval
1.0×10^3	3.798	3.798 $\pm 1.1 \times 10^{-1}$
2.0×10^3	3.798	3.792 $\pm 7.4 \times 10^{-2}$
5.0×10^3	3.798	3.801 $\pm 5.4 \times 10^{-2}$
1.0×10^4	3.798	3.798 $\pm 4.0 \times 10^{-2}$
2.0×10^4	3.798	3.798 $\pm 2.9 \times 10^{-2}$
5.0×10^4	3.798	3.798 $\pm 1.8 \times 10^{-2}$
1.0×10^5	3.798	3.798 $\pm 1.0 \times 10^{-2}$
2.0×10^5	3.798	3.798 $\pm 7.4 \times 10^{-3}$

We present the results of computations with the 95% steady-state quantile $x = 1.8993$ and $r(\cdot) = r_1$.

A Numerical Example: Square-root Convergence Rate

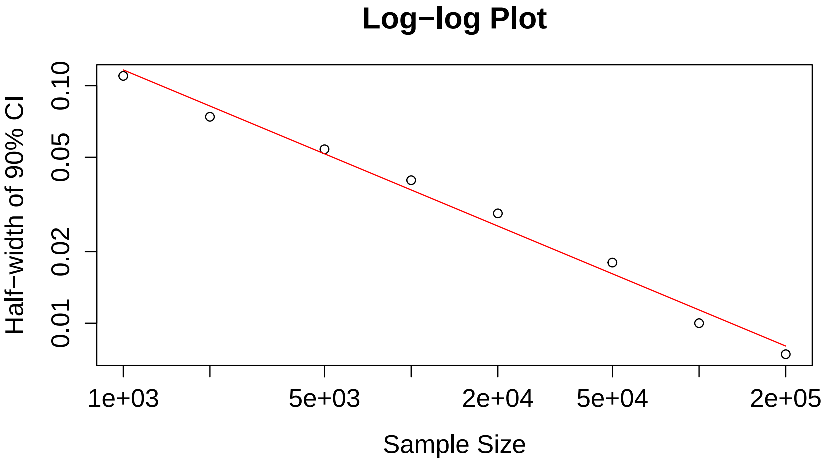

Figure: The redline is a fitted linear regression, with slope = -0.506

Computing the Spectral Density

- The spectral density satisfies

$$2\pi f(\lambda) = \mathrm{var}_\pi r(X_0) + \sum_{j=1}^{\infty} (e^{i\lambda j} + e^{-i\lambda_j}) \mathrm{cov}_\pi (r(X_0), r(X_j))$$

- A coupling that works here is:

$$Y_n = \left[r^2((\phi_1 \circ \cdots \circ \phi_n)(x)) - r((\phi_1 \circ \cdots \circ \phi_n)(x)) r((\tilde{\phi}_1 \circ \cdots \circ \tilde{\phi}_n)(x)) \right]$$

$$+ r((\phi_1 \circ \cdots \circ \phi_n)(x)) \sum_{j=1}^{n} (e^{i\lambda j} + e^{-i\lambda j}) \cdot$$

$$[r((\tilde{\phi}_1 \circ \cdots \circ \tilde{\phi}_j \circ \phi_1 \circ \cdots \circ \phi_n)(x)) - r((\tilde{\phi}_1 \circ \cdots \circ \tilde{\phi}_n)(x))]$$

- Yields a canonical convergence rate (conventional estimator is sub-canonical)

Computing Steady-state Gradients Efficiently

Setting: $X = (X_n : n \geq 0)$ positive recurrent Markov chain with transition kernel $P(\theta)$ having stationary distribution $\pi(\theta)$

Problem: Given a reward function $r : S \to \mathbb{R}_+$, compute the gradient

$$\nabla_\theta \int_S \pi(\theta, dx) r(x)$$

Idea:

$$\begin{aligned}
\pi(\theta) &= \pi(\theta) P(\theta) \\
\pi'(\theta)(I - P(\theta)) &= \pi(\theta) P'(\theta) \\
\pi'(\theta) &= \pi(\theta) P'(\theta) \left[\sum_{n=0}^{\infty} P^n(\theta) - \pi(\theta) \right] \\
\pi'(\theta) r &= \sum_{k=1}^{\infty} \mathbf{E}^{\theta}_{\pi(\theta)} \frac{p'(\theta, X_0, X_1)}{p(\theta, X_0, X_1)} \left(r(X_k) - \mathbf{E} r(X_\infty) \right)
\end{aligned}$$

Can be justified under suitable Lyapunov conditions on r (Rhee and G (2016))

The Multi-level Alternative

- A coupling that works here is

$$Y_n = \sum_{j=1}^{n} \frac{p'(\theta, (\phi_1 \circ \cdots \circ \phi_n)(x), (\phi_0 \circ \cdots \circ \phi_n)(x))}{p(\theta, (\phi_1 \circ \cdots \circ \phi_n)(x), (\phi_0 \circ \cdots \circ \phi_n)(x))}$$
$$\cdot j \left[r((\tilde{\phi}_1 \circ \cdots \circ \tilde{\phi}_{j-1} \circ \phi_0 \circ \cdots \circ \phi_n)(x)) - r((\tilde{\phi}_1 \circ \cdots \circ \tilde{\phi}_j \circ \phi_0 \circ \cdots \circ \phi_n)(x)) \right]$$

- Yields a canonical convergence rate, even for non-regenerative systems

Conclusions:

- "De-biasing" a sequence of estimators and MLMC are closely related
- Exact estimation vs exact simulation
- Several applications to Markov chains that improve the convergence rate to "square root" rate:
  - solutions to Poisson's equation
  - spectral density computation
  - gradient computation

Questions?