

A Randomized Quasi-Monte Carlo Simulation Method for Markov Chains

Pierre L'Ecuyer

GERAD and Département d'Informatique et de Recherche Opérationnelle, Université de Montréal, C.P. 6128,
Succ. Centre-Ville, Montréal, Québec, Canada H3C 3J7, lecuyer@iro.umontreal.ca

Christian Lécot

Laboratoire de Mathématiques, Université de Savoie, 73376 Le Bourget-du-Lac Cedex, France,
christian.lecot@univ-savoie.fr

Bruno Tuffin

IRISA-INRIA, Campus Universitaire de Beaulieu, 35042 Rennes Cedex, France, bruno.tuffin@irisa.fr

We introduce and study a randomized quasi-Monte Carlo method for the simulation of Markov chains up to a random (and possibly unbounded) stopping time. The method simulates n copies of the chain in parallel, using a $(d + 1)$ -dimensional, highly uniform point set of cardinality n , randomized independently at each step, where d is the number of uniform random numbers required at each transition of the Markov chain. The general idea is to obtain a better approximation of the state distribution, at each step of the chain, than with standard Monte Carlo. The technique can be used in particular to obtain a low-variance unbiased estimator of the expected total cost when state-dependent costs are paid at each step. It is generally more effective when the state space has a natural order related to the cost function.

We provide numerical illustrations where the variance reduction with respect to standard Monte Carlo is substantial. The variance can be reduced by factors of several thousands in some cases. We prove bounds on the convergence rate of the worst-case error and of the variance for special situations where the state space of the chain is a subset of the real numbers. In line with what is typically observed in randomized quasi-Monte Carlo contexts, our empirical results indicate much better convergence than what these bounds guarantee.

Subject classifications: simulation: efficiency; probability: Markov processes.

Area of review: Simulation

History: Received April 2005; revisions received February 2006, August 2006, January 2007, March 2007; accepted March 2007.

1. Introduction

A wide variety of real-life systems can be modeled as Markov chains with a large state space. In most interesting situations, analytic formulas are not available for these Markov chains, and matrix-based numerical methods require too much time, so Monte Carlo simulation becomes the standard way of estimating performance measures for these systems.

In this paper, we propose a randomized quasi-Monte Carlo (RQMC) algorithm, based on deterministic methods introduced by Lécot and Ogawa (2002) and Lécot and Tuffin (2004), to improve simulation efficiency for discrete-time Markov chains. The algorithm simulates n copies of the chain in parallel and induces negative dependence between the corresponding sample paths by using some form of generalized antithetic variates (Wilson 1983, Ben-Ameur et al. 2004). We would like the empirical distribution of the states of these n chains, at any given step j , to be a better approximation of the corresponding theoretical distribution than if the n chains were simulated independently. As a result, performance measure estimators obtained by

taking an average across the n copies of the chain will typically have much smaller variance.

Markov Chain Setting. We consider a Markov chain $\{X_j, j \geq 0\}$ with state space $\mathcal{X} \subseteq \mathbb{R}^l \cup \{\infty\}$ for some integer $l > 0$. The chain evolves according to the stochastic recurrence

$$X_0 = x_0, \quad X_j = \varphi_j(X_{j-1}, \mathbf{U}_j), \quad j \geq 1, \quad (1)$$

where the \mathbf{U}_j are independent random vectors uniformly distributed over the d -dimensional unit cube $[0, 1]^d$ (henceforth denoted $\mathbf{U}_j \sim U[0, 1]^d$), for some functions $\varphi_j: \mathcal{X} \times [0, 1]^d \rightarrow \mathbb{R}$. Each \mathbf{U}_j represents the d uniform random numbers required to simulate step j of the chain. Every discrete-time Markov chain that can be simulated on a computer fits this framework. It is also always possible to take $d = 1$ by defining the chain so that each newly generated uniform random number in the simulation corresponds to one step of the chain, although this representation is not always natural. Here, for more flexibility, we simply assume that d is a finite constant.

We want to estimate the expected total cost $\mu = E[Y]$, where

$$Y = \sum_{j=0}^{\tau} c_j(X_j), \quad (2)$$

$c_j: \mathcal{X} \rightarrow \mathbb{R}$ is a cost function at step j , τ is a stopping time with respect to the filtration $\{\mathcal{F}_j, j \geq 0\}$ generated by $\{(j, X_j), j \geq 0\}$, $E[\tau] < \infty$, and we assume that the functions φ_j and c_j are easy to evaluate at any given point for each j . We also assume (implicitly) that \mathcal{X} , the φ_j s, and the c_j s satisfy the appropriate measure-theoretic requirements so that all objects of interest in this paper are well defined. The state ∞ is an absorbing state used to indicate that we have already reached the stopping time τ at the previous step or earlier; i.e., $X_j = \infty$ for $j > \tau$. We also have $c_j(\infty) = 0$, although the cost at the stopping time τ can be nonzero.

Given a sequence of vectors $\{\mathbf{U}_j, j \geq 1\}$ in $[0, 1]^d$, we can simulate the chain until hitting the stopping time τ by computing $X_j = \varphi_j(X_{j-1}, \mathbf{U}_j)$ iteratively, and we can add the $c_j(X_j)$ s to compute the realization of Y . Let $s = \sup_{\omega} \tau d$, where the supremum is taken over all possible sample paths ω , so $s = \infty$ if τ is unbounded. Then, Y can be written as $Y = f(U_1, \dots, U_s)$ for some function f , where $\mathbf{U}_j = (U_{(j-1)d+1}, \dots, U_{jd})$ for each j . In the first three methods to be discussed, we simulate n copies of Y ; for the i th copy, we compute $Y_i = f(\mathbf{V}_i)$, where $\mathbf{V}_i = (U_{i,1}, \dots, U_{i,s}) \in [0, 1]^s$, $i = 0, \dots, n-1$, and we take the average \bar{Y}_n of Y_0, \dots, Y_{n-1} as an estimator of μ .

Estimation by Monte Carlo (MC). In the standard MC method, all the $U_{i,j}$ s are taken as realizations of independent uniform random variables over $[0, 1]$, so the \mathbf{V}_i are independent and uniformly distributed in $[0, 1]^s$. Then, the sample mean \bar{Y}_n and the sample variance S_n^2 of Y_0, \dots, Y_{n-1} are unbiased estimators of μ and of the variance of Y , respectively. From this, we can compute a confidence interval on μ , e.g., via the central limit theorem.

Quasi-Monte Carlo (QMC). The classical QMC method replaces the n independent random points \mathbf{V}_i used in MC by a deterministic set of distinct points $P_n = \{\mathbf{v}_0, \dots, \mathbf{v}_{n-1}\}$ that cover the unit cube $[0, 1]^s$ more uniformly than typical random points. There are many ways of measuring the uniformity of P_n . One of them is the star discrepancy, defined as

$$D_n^*(P_n) = \sup_{\mathbf{x} \in [0, 1]^s} |\text{card}(P_n \cap [\mathbf{0}, \mathbf{x}]) / n - \text{vol}([\mathbf{0}, \mathbf{x}])|,$$

where $[\mathbf{0}, \mathbf{x}]$ is the s -dimensional rectangular box with opposite corners $\mathbf{0}$ and \mathbf{x} , vol denotes its volume, and $\text{card}(P)$ means the cardinality of the set P (so $\text{card}(P_n \cap [\mathbf{0}, \mathbf{x}])$ is the number of points of P_n that fall in the box). A classical worst-case error bound for numerical integration by QMC is given by the Koksma-Hlawka inequality

$$|\bar{Y}_n - \mu| \leq V(f) D_n^*(P_n), \quad (3)$$

where $V(f)$ is the Hardy-Krause total variation of the integrand f (see Niederreiter 1992 for the details). For a one-dimensional function $g: [a, b] \rightarrow \mathbb{R}$, where $[a, b] \subset \mathbb{R}$, the total variation is

$$V(g) = \sup_{L \geq 1, a=x_0 < \dots < x_L=b} \sum_{j=1}^L |g(x_j) - g(x_{j-1})| = \int_a^b |dg(x)|. \quad (4)$$

If $[a, b]$ is replaced by \mathbb{R} , then $V(g) = \int_{-\infty}^{\infty} |dg(x)|$. For $s > 1$, the total variation is much more complicated; we just point out that $V(f)$ is infinite whenever f has a discontinuity not aligned with one of the axes (Owen 2005). This happens frequently. Moreover, in most practical situations, the bound (3) is too loose to be useful and is also much too difficult to compute. Hlawka (1971) derives another error bound in which $V(f)$ is replaced by a different notion of variation, called the mean oscillation of f . In fact, a large family of bounds can be derived by adopting different definitions of discrepancy, together with corresponding definitions of function variation (Hickernell 1998, L'Ecuyer and Lemieux 2002). These bounds can be used to get asymptotic convergence rates, but are rarely convenient for practical error assessment.

Randomized Quasi-Monte Carlo (RQMC). The usual RQMC method (we call it classical RQMC, to distinguish it from our new algorithm) turns QMC into a variance reduction method by carefully randomizing P_n . This can be seen as an application of the generalized antithetic variates (GAV) principle (Wilson 1983), as explained in Ben-Ameur et al. (2004). The idea is to induce “negative dependence” between the points \mathbf{V}_i (for example, assuming that each $\text{Cov}(\mathbf{V}_i, \mathbf{V}_j)$ is a diagonal matrix, we may ask for all diagonal elements of the diagonal matrix $\sum_{(i,j): i \neq j} \text{Cov}(\mathbf{V}_i, \mathbf{V}_j)$ to be negative) by generating them in such a way that

- (a) each point \mathbf{V}_i has the uniform distribution over $[0, 1]^s$, and
- (b) the point set $\{\mathbf{V}_0, \dots, \mathbf{V}_{n-1}\}$ covers $[0, 1]^s$ more uniformly (in some sense) than a set of independent random points.

We call a point set that satisfies these two conditions an RQMC point set. This definition is incomplete: we still need to select a specific measure of uniformity so that “more uniformly” has a well-defined meaning. We leave this choice open for now because we want to keep the generality. Specific measures and conditions are always adopted when proving bounds on the convergence speed of the error and variance for particular RQMC methods, and we will do the same when we prove such results for our method, but the RQMC method in general and the algorithm proposed in this paper are defined independently of any measure of uniformity. Examples of RQMC point sets include randomly shifted lattice rules, scrambled digital nets, digital nets with a random digital shift, a Latin hypercube sample

or a stratified sample followed by a random permutation of the points, and so forth (Owen 1998, L'Ecuyer and Lemieux 2002, Glasserman 2004). For the stratified sample, we partition the unit cube into n boxes of equal volume and generate one point randomly and uniformly in each box. Then, we permute the n points randomly, so that for each fixed i , V_i has probability $1/n$ of being in any given box. The same permutation trick applies to Latin hypercube sampling.

Under conditions (a) and (b) on the randomization, the estimator \bar{Y}_n is still unbiased for μ and, although Y_0, \dots, Y_n are not independent, an unbiased estimator of $\text{Var}[\bar{Y}_n]$ can be obtained by computing m independent replicates of \bar{Y}_n , e.g., with independent randomizations of the same point set, and taking the sample variance of these m averages. Further details on this classical RQMC approach (and high-dimensional contexts) can be found in Owen (1998), L'Ecuyer and Lemieux (2000, 2002), and other references given there. This approach is typically more efficient than MC when s is small (e.g., fewer than 20 or so) or if the function f has *low effective dimension* in some sense, as explained in Owen (1998) and L'Ecuyer and Lemieux (2002).

Introductory overviews of the QMC and RQMC methods can be found in Kollig and Keller (2002), Owen (2003a), Glasserman (2004), L'Ecuyer (2004a), and the user's guide of the package hups in SSJ (L'Ecuyer 2004b), for example. For more advanced material, see Niederreiter (1992), Owen (1998), L'Ecuyer and Lemieux (2002), and the proceedings of the biennial MCQMC Conference.

Array-RQMC. The RQMC method proposed in this paper, called *array-RQMC*, operates differently. We simulate n copies of the chain in parallel. To simulate step j for all copies, we use a randomized $(d + 1)$ -dimensional highly uniform (or “low-discrepancy”) point set $P'_{n,j}$ of cardinality n , as explained in the next section, where $d \ll s$ typically. These point sets are randomized independently at the different steps in such a way that the sample path of any given copy of the chain obeys the correct probability law (the same as with the MC method). As a result, we have an unbiased estimator \bar{Y}_n for the average cost μ . The aim of the proposed method is to induce dependence across the n copies so that the empirical distribution of the n realizations of X_j (at step j) gives a better approximation of the distribution F_j of the random variable X_j than if the chains were simulated independently.

The original deterministic method of Lécot and Tuffin (2004) was designed to approximate transient measures over a fixed number of steps for discrete-time and discrete-state Markov chains with a totally ordered state space, and with $d = 1$. That method uses a $(0, 2)$ -sequence in base 2 (Niederreiter 1992). At step j of the chain, it reorders the n copies according to their current states and “simulates” the transitions (next states) for the n copies by employing the elements n_j to $n_j + n - 1$ of the $(0, 2)$ -sequence

in place of uniform random numbers to drive the simulation. Convergence to the correct value (in the deterministic sense) was proved by Lécot and Tuffin (2004) under a condition on the structure of the transition probability matrix of the Markov chain. In contrast, our method is a randomized algorithm that provides an unbiased estimator. It also applies to Markov chains with a more general state space, with a random and unbounded number of steps, and allows $d > 1$. It thus covers a much broader range of applications.

We prove bounds on the convergence rate of the variance of the mean estimator (as $n \rightarrow \infty$) only for simplified special cases of the algorithm, with $\mathcal{X} \subseteq (-\infty, \infty]$ and $d = 1$. For a setting where each $P'_{n,j}$ is constructed by stratified sampling in the unit square, and under mild additional assumptions, we prove that the variance converges as $O(n^{-3/2})$. Under a less restrictive assumption than stratified sampling (see Assumption 2), we show that the worst-case error converges as $O(n^{-1/2})$ (this is a stronger result than the probabilistic rate of $O(n^{-1/2})$ for MC). Under even milder assumptions on the sampling (we only assume that the chain transitions are generated by inversion), but only for chains with finite state spaces, we obtain an upper bound of $O(n^{-1} \log n)$ on the convergence rate, also in the worst case. Related results have been obtained by Lécot (1996) for general real-valued functions defined over the unit cube and having bounded variation, in a deterministic setting, under the stronger assumption that P_n is a (t, m, s) -net, and with different methods of proof. We conduct empirical experiments with other types of RQMC point sets as well, with a variety of examples. The results indicate that the variance goes down much faster (as a function of n) with the proposed method than for standard MC and faster than classical RQMC in many situations. In our experiments, popular RQMC point sets such as randomized lattice rules and Sobol' nets also perform much better than the stratified sampling mentioned earlier. This gap between theoretical and empirical results (we often get much better variance improvements in practice than what we can prove) is common with the QMC and RQMC methods (Tuffin 1996, Tezuka 2002), where the theoretical bounds have limited applicability in general and the surprisingly good empirical results are often only heuristically justified. The theoretical results are nevertheless a first step toward a better understanding of the proposed algorithm. We also believe that they could be extended in several directions (different sampling schemes, higher dimensions, etc.), and some of the ideas in our proofs could be useful for that purpose.

2. The Array-RQMC Algorithm

We now assume that there is a *sorting function* $h: \mathcal{X} \rightarrow \mathbb{R} \cup \{\infty\}$ that assigns a real number to each state of the chain, with $h(\infty) = \infty$. This h is used to order the states: we say that state x_1 is *smaller* than state x_2 if $h(x_1) < h(x_2)$ and that the two states are *h-equivalent* if $h(x_1) = h(x_2)$.

The h -equivalent states can be placed in arbitrary order by the sort. A similar type of function h is used in the splitting methodology for rare-event simulation, where it is called the *importance function* (Garvels et al. 2002, Glasserman et al. 1999). Just like for splitting, a good choice of h is crucial for the performance of the algorithm, especially when \mathcal{X} has more than one dimension.

The array-RQMC algorithm works by simulating n copies of the chain in parallel as follows. In step 1, we take an RQMC point set $P_{n,1} = \{\mathbf{u}_{0,1}, \dots, \mathbf{u}_{n-1,1}\}$ in $[0, 1)^d$, define $X_{i,1} = \varphi_1(x_0, \mathbf{u}_{i,1})$ for $i = 0, \dots, n-1$, and estimate the distribution function F_1 of X_1 by the empirical distribution \hat{F}_1 of $X_{0,1}, \dots, X_{n-1,1}$. This gives

$$F_1(x) = \int_{[0,1]^d} I(\varphi_1(x_0, \mathbf{u}) \leq x) d\mathbf{u} \quad (5)$$

$$\approx \frac{1}{n} \sum_{i=0}^{n-1} I(\varphi_1(x_0, \mathbf{u}_{i,1}) \leq x) \quad (6)$$

$$= \frac{1}{n} \sum_{i=0}^{n-1} I(X_{i,1} \leq x) \stackrel{\text{def}}{=} \hat{F}_1(x),$$

where I denotes the indicator function. The approximation in (6) amounts to estimating the integral (5) by RQMC. If φ_1 is well behaved, this should give better accuracy than using standard MC.

In step j , we start with the empirical distribution \hat{F}_{j-1} of $X_{0,j-1}, \dots, X_{n-1,j-1}$ as an approximation of the distribution function F_{j-1} of X_{j-1} , and want to compute a good approximation \hat{F}_j of F_j . We can write

$$F_j(x) = \int_{\mathcal{X}} \int_{[0,1]^d} I(\varphi_j(y, \mathbf{u}) \leq x) d\mathbf{u} dF_{j-1}(y) \quad (7)$$

$$\approx \int_{\mathcal{X}} \int_{[0,1]^d} I(\varphi_j(y, \mathbf{u}) \leq x) d\mathbf{u} d\hat{F}_{j-1}(y) \quad (8)$$

$$= \frac{1}{n} \sum_{i=0}^{n-1} E[I(\varphi_j(X_{i,j-1}, \mathbf{U}_{i,j}) \leq x)], \quad (9)$$

where the approximation in (8) is obtained by replacing F_{j-1} in (7) with \hat{F}_{j-1} . When the n copies of the Markov chain are simulated independently via standard MC, (9) is estimated by its realization (the expectation is removed), where the $\mathbf{U}_{i,j}$ s are independent and uniformly distributed over $[0, 1)^d$. Our aim here is to estimate the $(d+1)$ -dimensional integral (8) by RQMC instead.

To do that, we introduce a $(d+1)$ -dimensional *modified RQMC point set* $P'_{n,j}$ defined as $P'_{n,j} = \{\mathbf{u}'_{i,j} = ((i+0.5)/n, \mathbf{u}_{i,j}), 0 \leq i < n\}$, where $P_{n,j} = \{\mathbf{u}_{0,j}, \dots, \mathbf{u}_{n-1,j}\}$ is an RQMC point set in $[0, 1)^d$, and with the following two properties:

(c) $\mathbf{u}_{i,j}$ is a random vector uniformly distributed over $[0, 1)^d$ for each i , and

(d) $P'_{n,j}$ is “highly uniform” in $[0, 1)^{d+1}$, in a sense that we leave open for now (as in our definition of RQMC point set).

This $P'_{n,j}$ is not quite an RQMC point set because the first coordinate is not a random variable with the uniform distribution over $[0, 1)$. A typical way of defining $P'_{n,j}$ is to take a $(d+1)$ -dimensional RQMC point set, sort the points by order of their first coordinate, and then replace the first coordinate of the i th point by $(i+0.5)/n$ for each i . Item (b) is not needed for the proofs of Propositions 1 and 2 later in this section, but is crucial to obtain a variance reduction in the array-RQMC algorithm.

To explain how the integral (8) is approximated by RQMC with the point set $P'_{n,j}$, we first consider the special case where $l = 1$ (so $\mathcal{X} \subseteq (-\infty, \infty]$) and $h(x) = x$. Let $X_{(0),j-1}, \dots, X_{(n-1),j-1}$ be the states at step $j-1$ sorted by increasing order. Define the inverse of \hat{F}_{j-1} by $\hat{F}_{j-1}^{-1}(v) = \inf\{x \in \mathcal{X} : \hat{F}_{j-1}(x) \geq v\}$ for all $v \in [0, 1]$. We approximate (8) as follows:

$$\int_{\mathcal{X}} \int_{[0,1]^d} I(\varphi_j(y, \mathbf{u}) \leq x) d\mathbf{u} d\hat{F}_{j-1}(y) = \int_{[0,1]^{d+1}} I(\varphi_j(\hat{F}_{j-1}^{-1}(v), \mathbf{u}) \leq x) d\mathbf{u} dv \quad (10)$$

$$\approx \frac{1}{n} \sum_{i=0}^{n-1} I(\varphi_j(\hat{F}_{j-1}^{-1}((i+0.5)/n), \mathbf{u}_{i,j}) \leq x) \quad (11)$$

$$= \frac{1}{n} \sum_{i=0}^{n-1} I(\varphi_j(X_{(i),j-1}, \mathbf{u}_{i,j}) \leq x) \quad (12)$$

$$= \frac{1}{n} \sum_{i=0}^{n-1} I(X_{i,j} \leq x) \stackrel{\text{def}}{=} \hat{F}_j(x).$$

In (11), we approximate the integral in (10) by RQMC over $[0, 1)^{d+1}$ with the point set $P'_{n,j}$. The sorting at step $j-1$ is needed for (12) to be true; i.e., the i th point of $P'_{n,j}$ must be assigned to the chain whose state is $X_{(i),j-1}$. Observe that this point set gives a perfect stratification of the distribution \hat{F}_{j-1} , with exactly one observation per stratum (the strata are the jumps of \hat{F}_{j-1}). On the other hand, these observations are generally not independent across the strata. The distribution F_j is estimated by the empirical distribution \hat{F}_j of the realizations $X_{i,j} = \varphi_j(X_{(i),j-1}, \mathbf{u}_{i,j})$, $i = 0, \dots, n-1$. This RQMC approximation in (11) is expected to be more accurate than using standard MC. The approximation error in (8) depends on how well F_{j-1} is approximated by \hat{F}_{j-1} , i.e., on an accumulation of integration errors in (11) over the previous stages. The rationale of the array-RQMC method is to reduce this integration error by using RQMC at each stage.

In the case where $l > 1$, the inverse function \hat{F}_{j-1}^{-1} is not well defined, so (10) no longer stands. However, suppose that the function h is selected so that whenever two states $X_{1,j}$ and $X_{2,j}$ are h -equivalent, the distribution of the future costs $c_{j'}(X_{j'})$ for $j' > j$ conditional on $X_{i,j}$ is the same for $i = 1$ and $i = 2$. In this idealistic setting, knowing the value of $Z_j = h(X_j)$ at Step j provides as much relevant information as knowing the state X_j , so it

suffices to approximate the (univariate) distribution function G_j of Z_j instead of approximating F_j . Assuming that $Z_j = \tilde{\varphi}_j(Z_{j-1}, \mathbf{U}_j)$ for some functions $\tilde{\varphi}_j$, denoting $Z_{i,j-1} = h(X_{i,j-1})$, and denoting \hat{G}_j the empirical distribution of the sorted values $Z_{(0),j} \leq \dots \leq Z_{(n-1),j}$, we can use the same argument as above with G instead of F :

$$\begin{aligned} G_j(z) &= \int_{-\infty}^{\infty} \int_{[0,1]^d} I(\tilde{\varphi}_j(y, \mathbf{u}) \leq z) d\mathbf{u} dG_{j-1}(y) \\ &\approx \int_{-\infty}^{\infty} \int_{[0,1]^d} I(\tilde{\varphi}_j(y, \mathbf{u}) \leq z) d\mathbf{u} d\hat{G}_{j-1}(y) \\ &= \int_{[0,1]^{d+1}} I(\tilde{\varphi}_j(\hat{G}_{j-1}^{-1}(v), \mathbf{u}) \leq z) d\mathbf{u} dv \\ &\approx \frac{1}{n} \sum_{i=0}^{n-1} I(\tilde{\varphi}_j(\hat{G}_{j-1}^{-1}((i+0.5)/n), \mathbf{u}_{i,j}) \leq z) \\ &= \frac{1}{n} \sum_{i=0}^{n-1} I(\tilde{\varphi}_j(Z_{(i),j-1}, \mathbf{u}_{i,j}) \leq z) = \hat{G}_j(z). \end{aligned}$$

In practice, it is usually not possible (or too difficult) to select h so that h -equivalent states X are exactly equivalent in the sense that they give the same distributions for the future costs, but a good choice of h should try to approximate this. The performance of the method depends on this choice. With a bad choice of h , the variance may not be reduced at all or may even increase, but reasonable heuristic choices may suffice to reduce the variance, as our examples will illustrate. The following algorithm and Propositions 1 and 2 that follow are valid regardless of h .

The array-RQMC algorithm can be summarized as follows. We first select a d -dimensional QMC point set $\tilde{P}_n = (\tilde{\mathbf{u}}_0, \dots, \tilde{\mathbf{u}}_{n-1})$ and a randomization method for \tilde{P}_n such that the randomized version P_n defines a modified RQMC point set P'_n in the sense of (c) and (d). Then, we simulate in parallel n copies of the chain, numbered $0, \dots, n-1$, as follows (the braces delimit the scope of the “while” loops):

Array-RQMC Algorithm

```

Let  $X_{0,0} \leftarrow x_0, \dots, X_{n-1,0} \leftarrow x_0$ , and  $j \leftarrow 1$ ;
While  $X_{0,j-1} < \infty$  do {
  Randomize  $\tilde{P}_n$  afresh into  $P_{n,j} = \{\mathbf{u}_{0,j}, \dots, \mathbf{u}_{n-1,j}\}$ 
  (i.e.,  $P_{n,j}$  is a new randomization of  $\tilde{P}_n$ ,
  independent of all previous randomizations);
  Let  $i \leftarrow 0$ ;
  While ( $i < n$  and  $X_{i,j-1} < \infty$ ) do {
     $X_{i,j} \leftarrow \varphi_j(X_{i,j-1}, \mathbf{u}_{i,j})$ ;  $i \leftarrow i + 1$ ;
  }
  Sort the states  $X_{0,j}, \dots, X_{n-1,j}$  by increasing order
  of their values of  $h(X_{i,j})$ ,
  and renumber them in this order, i.e.,
  so that  $h(X_{0,j}) \leq \dots \leq h(X_{n-1,j})$ ;
   $j \leftarrow j + 1$ ;
}
    
```

Return the average \bar{Y}_n of the n realizations of Y as an estimator of μ .

Let $\bar{Y}_{n,j}$ denote the average cost at step j , across the n copies. We have $\bar{Y}_n = \sum_{j=1}^{\infty} \bar{Y}_{n,j}$. The randomness of τ may, of course, affect the performance of the algorithm because the number of copies of the chain that remain alive decreases with the number of steps, so it could happen that there remain just a few copies for several steps near the end, in which case only a few points from P_n are used in these steps. Two of our numerical examples, in §§4.3 and 4.6, have a random τ .

PROPOSITION 1. *The averages $\bar{Y}_{n,j}$ and \bar{Y}_n are unbiased estimators of $E[c_j(X_j)]$ and μ , respectively.*

The successive steps of the chain use independent randomizations. Therefore, for each chain, from the assumption made in item (c) of the definition of a modified RQMC point set, the vectors that take the place of the \mathbf{U}_j s in the recurrence (1) to generate the successive steps j of the chain are independent random variables uniformly distributed over $[0, 1]^d$. Thus, any given copy of the chain obeys the correct probabilistic model defined by (1) and (2), so $c_j(X_j)$ and Y have the correct expectations, $E[c_j(X_j)]$ and μ , and their averages over the n copies as well.

To estimate the variance of \bar{Y}_n and compute a confidence interval on μ , we can replicate this entire procedure independently m times. That is, across the m replications, all randomizations are independent. Because this gives m independent unbiased estimators of μ , we have:

PROPOSITION 2. *The empirical variance of the m copies of \bar{Y}_n is an unbiased estimator of $\text{Var}[\bar{Y}_n]$.*

This proposition implies that the variance of the overall average converges as $O(1/m)$ when $m \rightarrow \infty$. In the next section, we examine the convergence rate as a function of n when $n \rightarrow \infty$ for simplified cases.

3. Convergence

How can we prove theoretical results on the convergence speed (or rate) of $\bar{Y}_{n,j}$ and \bar{Y}_n to their expectations? A first idea would be to bound the integration error represented by the approximation signs in (6), (8), and (11), via the Koksma-Hlawka inequality (3). This is generally ineffective because the integrand $I(\varphi_j(\hat{F}_{j-1}^{-1}(v), \mathbf{u}) \leq x)$ in (10) may have infinite Hardy-Krause variation: this integrand is equal to one in part of the unit cube, zero elsewhere, and the shape and complexity of the boundary between these two regions depends on $\varphi_1, \dots, \varphi_j$. This boundary (on which f is discontinuous) is often not aligned with one of the axes. In the appendix, we prove bounds on the convergence rate directly from first principles for special cases for which $l = d = 1$ and $P'_{n,j}$ has special properties. Here we summarize the main results. Let

$$\Delta_j = \sup_{x \in \mathcal{X}} |\hat{F}_j(x) - F_j(x)| = \sup_{x \in \mathbb{R}} |\hat{F}_j(x) - F_j(x)|, \tag{13}$$

INFORMS holds copyright to this article and distributed this copy as a courtesy to the author(s). Additional information, including rights and permission policies, is available at http://journals.informs.org/.

the Kolmogorov distance between \widehat{F}_j and F_j (with $\Delta_0 = 0$), and

$$V(c_j) = \int_{-\infty}^{\infty} |dc_j(x)|, \quad (14)$$

the total variation of the cost function c_j . The error on the expected cost at step j is bounded by the product of these two quantities:

PROPOSITION 3. We have $|\bar{Y}_{n,j} - E[c_j(X_j)]| \leq \Delta_j V(c_j)$.

When c_j has bounded variation, this proposition tells us that the square error of $\bar{Y}_{n,j}$ converges in the worst case at the same rate as Δ_j^2 when $n \rightarrow \infty$. If τ is bounded, this implies that $|\bar{Y}_n - \mu|^2$ converges as $O(\sum_{j=1}^{\tau} \Delta_j^2)$ in the worst case. The next results provide bounds on Δ_j , under two different sets of assumptions. Proposition 4 gives a better asymptotic rate than Proposition 5, under a milder assumption on the sampling, but the bound also increases linearly with the cardinality of the state space, whereas in Proposition 5 it is independent of the cardinality.

ASSUMPTION 1. We have $l = 1$, and at each step j of the chain, we use inversion from a single uniform random variable to generate the next state X_j from its conditional distribution given X_{j-1} . This means that each φ_j is nondecreasing with respect to its second argument.

PROPOSITION 4. Let Assumption 1 hold. Suppose that the Markov chain has a finite state space $\mathcal{X} = \{1, \dots, L\}$ and that the star discrepancy of each $P'_{n,j}$ satisfies $D_n^*(P'_{n,j}) = O(n^{-1} \log n)$ w.p.1 (this can easily be achieved, for example, by constructing $P'_{n,j}$ as follows: take the first n points of a $(0, 2)$ -sequence in some base b , either randomized or deterministic, sort them by their first coordinate, and replace the first coordinate of point i by $(i + 0.5)/n$ for each i). Then,

$$\Delta_j \leq LK_j^{(1)} n^{-1} \log n \quad (15)$$

for all j , for some constants $K_j^{(1)}$ that depend only on j .

ASSUMPTION 2. Suppose that n is a square number, so \sqrt{n} is an integer, and that if we partition the unit square into n subsquares of size $n^{-1/2} \times n^{-1/2}$, each of those subsquares contains exactly one point from $P'_{n,j}$.

PROPOSITION 5. Under Assumptions 1 and 2, we have

$$\Delta_j \leq K_j^{(2)} n^{-1/2} \quad (16)$$

for all j , for some constants $K_j^{(2)}$ that depend only on j .

Several well-known RQMC point sets satisfy Assumption 2; for example, a digitally shifted version of the two-dimensional Hammersley point set (or Sobol' net, or Faure net) in prime base b , or an affine matrix scramble of it, when n is a power of b (Matoušek 1999, Owen 2003b), or even a stratified sample of n points in the n subsquares.

ASSUMPTION 3. Let Assumption 2 hold, and let the random variables $v_i = \mathbf{u}_i \bmod n^{-1/2}$ be pairwise independent and uniformly distributed over $[0, n^{-1/2})$ (the \mathbf{u}_i s are one-dimensional in this case). This means that the second coordinate of each point of $P'_{n,j}$ is uniformly distributed over the (vertical) interval determined by the box that contains this point, independently of the position of any other point in its box.

PROPOSITION 6. Under Assumptions 1 and 3,

$$E[(\bar{Y}_{n,j} - E[c_j(X_j)])^2] \leq \Gamma_j V(c_j)^2 n^{-3/2}. \quad (17)$$

for all j , for some constants Γ_j that depend only on j . If the chain is stochastically increasing, we can take $\Gamma_j = j/2$.

This gives a convergence rate of $O(n^{-3/2})$ for the variance of the average cost at step j . It beats the $O(n^{-1})$ rate of MC. A point set that satisfies Assumption 3 is easily obtained by stratified sampling over the unit square: generate one point in each subsquare of the partition, uniformly over the subsquare and independently across the different subsquares, sort these points by the first coordinate, and then replace the first coordinate of the i th point by $(i + 0.5)/n$. Any point set that satisfies Assumption 3 is also equivalent to this stratified sampling construction in the array-RQMC algorithm.

Proving the result without the independence assumption is more difficult because it is more complicated to bound the variance of the array-RQMC estimator. Intuitively, with RQMC point sets that have more uniformity and negative dependence (e.g., good lattice rules and nets), we would expect an even better variance reduction than with stratified sampling. This is indeed what we have observed in all our numerical experiments (see §4). It may also be possible to prove error bounds that are at least as good, or even tighter. Potential generalizations of these results to higher-dimensional settings are discussed in the appendix.

In our numerical experiments of §4, instead of generating the points independently in the different squares as in Assumption 3, we will generate them according to RQMC schemes that provide more uniformity, with the aim of inducing more negative dependence (and more variance reduction). Some of these RQMC schemes (e.g., the two-dimensional Sobol' nets) have one point per subsquare as in Assumption 3, but none of them fully satisfies this assumption because the locations of the points in two different squares are not independent. These RQMC schemes turn out to work better in practice.

4. Numerical Illustrations

We compare MC, classical RQMC, and array-RQMC on a set of examples. We first describe the different RQMC point sets used in our examples; then we explain our experimental setting.

For the RQMC methods, we use Korobov lattice rules and Sobol' point sets. A Korobov rule is defined by two

parameters $0 < a < n$, and the corresponding s -dimensional point set is

$$P_n = \{\mathbf{v}_i = (i/n, (ia \bmod n)/n, \dots, (ia^{s-1} \bmod n)/n), \\ i = 0, \dots, n-1\}$$

(Niederreiter 1992, Sloan and Joe 1994, L'Ecuyer and Lemieux 2000). Given n and a , there is no limit on s , so P_n can be viewed as an infinite-dimensional point set (L'Ecuyer and Lemieux 2000). This P_n is randomized by applying a random shift modulo 1 to all the points simultaneously. We also consider applying the baker's transformation, which transforms each coordinate u to $2u$ if $u < 1/2$ and to $2(1-u)$ if $u \geq 1/2$, after the random shift. This transformation has a "locally antithetic" effect and provably improves the convergence rate when integrating a smooth function with a randomly shifted lattice rule (Hickernell 2002). For the parameters, we take n equal to the largest prime number smaller than 2^k for $k = 10, 12, \dots, 20$, a equal to the odd integer nearest to $n/1.61803399$ when $s = 2$ (so a/n is near the golden ratio; the rule always performs well in the two-dimensional spectral test with this choice), and a from Table 1 of L'Ecuyer and Lemieux (2000) for point sets in dimensions $s > 2$ (which is always the case for classical RQMC).

Our second type of point set is a Sobol' net with $n = 2^k$ points for $k = 10, 12, \dots, 20$ (Bratley and Fox 1988), randomized by a left (upper-triangular) matrix scrambling followed by a random digital shift (Matoušek 1999, L'Ecuyer and Lemieux 2002, Owen 2003b, Glasserman 2004), as implemented in the SSJ software (L'Ecuyer 2004b). Available implementations of Sobol' nets have an upper bound on the dimension s of the point set, so we cannot use these point sets when the dimension is large or unbounded.

With classical RQMC, we consider the following point sets (in each case, the name in parentheses will refer to the corresponding combination): a randomly shifted Korobov rule (*Classical-Korobov*); a randomly shifted Korobov rule with the baker's transformation (*Classical-Korobov-Baker*); and a randomized Sobol' point set (*Classical-Sobol*) when the dimension s is small.

For array-RQMC, the randomization is not applied to the first coordinate, but only to the subsequent coordinates, as explained when we gave our conditions on $P'_{n,j}$. The first coordinate is skipped and is only used to enumerate the points. The Korobov points are always enumerated by order of their (skipped) first coordinate, which takes the values $0, 1/n, \dots, (n-1)/n$ in succession (in the implicit definition of P'_n , we add $0.5/n$ to these values). For example, if $d = 1$, the i th point of P_n before the shift is $(ia \bmod n)/n$. The points of a Sobol' net are usually enumerated by order of their Gray code because this is a bit faster than enumerating them in their natural order (see, e.g., Antonov and Saleev 1979, Bratley and Fox 1988, and L'Ecuyer 2004b). Enumerating the points of P_n by their Gray code is equivalent to applying a permutation to the second and further

coordinates of the points of P'_n , i.e., it performs an additional scrambling for these coordinates. It is also possible to enumerate the points by order of their first coordinate. We tried both. (Note that for classical RQMC, the order of enumeration is irrelevant.)

In summary, for array-RQMC, we have the following types of d -dimensional RQMC point sets for P_n : a $(d+1)$ -dimensional Korobov lattice rule with its first coordinate skipped, randomly shifted (*Array-Korobov*); the same Korobov rule with the random shift followed by a baker's transformation (*Array-Korobov-Baker*); the first n points of a randomized Sobol' sequence where the points are enumerated by order of their Gray code (*Array-Sobol*); the same randomized Sobol' point set but with the points enumerated in their natural order (*Array-Sobol-NoGray*); and for $d = 1$, in the first example, we also tried a stratified sample that satisfies Assumption 3 (*Array-Stratification*).

For each combination and each value of n considered, we estimate the *variance reduction factor* (VRF) compared with standard MC, defined as $\text{Var}[Y]/(n \text{Var}[\bar{Y}_n])$, where \bar{Y}_n is the estimator with the RQMC method considered. We estimate $\text{Var}[\bar{Y}_n]$ by the sample variance of $m = 100$ independent copies of \bar{Y}_n and $\text{Var}[Y]$ by the empirical variance of a very large number of independent replicates of Y (up to several millions). Thus, $\text{Var}[Y]$ is very accurately estimated, but $\text{Var}[\bar{Y}_n]$ is not. As a result, the VRF estimates are noisy; as a crude assessment, they have about 10% to 20% relative accuracy. They are all rounded to the nearest integers. The simulations were performed in Java, using the SSJ simulation library (L'Ecuyer 2004b).

If we measure the work by the number of simulated copies of the chain, then the VRF as we have defined it also measures the efficiency improvement of the RQMC method compared with MC. Another possibility would be to measure the work in terms of CPU time. In fact, the CPU times to simulate the chains (excluding the sorting) are about the same for all the methods, with one exception: there is an overhead of 10% to 40% for the baker's transformation, presumably because we have implemented it as an additional layer that applies to any point set and involves more method calls. On the other hand, sorting the chains at each step of the array-RQMC algorithm brings significant additional overhead. It requires an $O(n \log n)$ effort as $n \rightarrow \infty$, whereas the simulation itself requires $O(n)$ time. In our experiments, we just used the general-purpose sorting procedures available in the standard Java libraries of JDK 1.5 from SUN. When the state space was one-dimensional (in the examples of §§4.1, 4.2, 4.3, and 4.6), we stored the states of the chains in an array of real numbers and just sorted that array at each step of array-RQMC. Then, the additional CPU time for sorting the chains ranged from approximately 50% (for $n = 2^{10}$) to 90% (for $n = 2^{20}$), so the overall CPU time was multiplied by a factor between 1.5 and 1.9 with array-RQMC. For higher-dimensional state spaces (the examples of §§4.4 and 4.5), the chains were stored as an array of Java objects,

and in that case the sorting algorithm invokes a comparison method each time it compares two objects. This is very general and flexible, but certainly not very efficient. In this case, using array-RQMC multiplied the CPU time by a factor from 2 to 5 with $n = 2^{10}$ and from 8 to 14 with $n = 2^{20}$ (depending on the example). These factors could be reduced by implementing specialized sorting procedures with many fewer method invocations, but we did not do that and made no serious attempt at optimizing our programs for speed. Just as a quick test after completing our experiments, we tried replacing the sorting procedure of JDK 1.5 by another general-purpose procedure taken from the Colt library (Hoschek 2004) (which also invokes a comparison method but uses a different sorting algorithm). For array-RQMC with $n = 2^{20}$, on the same computer, the CPU times were approximately halved. To obtain efficiency improvement factors in terms of CPU times, the VRFs reported in the tables would have to be divided by the factors just mentioned. We did not divide by these factors in the tables because they are highly dependent on the Java libraries and Java virtual machine that are used, the choice of programming language, the size of cache memory on the computer, etc., and they can probably be reduced. The VRFs given in the tables are independent of programming implementations.

For array-RQMC with the stratification, the overhead is even more important because the randomization is more expensive, and because we have to sort the points explicitly at each step of the chain. In our implementation (for the first example only), the stratified version takes about the same time as the other array-RQMC versions for $n = 2^{10}$ and about twice the time for $n = 2^{20}$.

4.1. An $M/M/1$ Queue with $d = 1$

Consider a single-server queue with i.i.d. exponential inter-arrival times A_j with mean 1 and i.i.d. exponential service times S_j with mean $\rho < 1$. This ρ is also the utilization factor of the server. We want to estimate the expected average waiting time of the first t customers, denoted by μ . We could compute μ numerically without simulation; we just use this simple academic example to illustrate our method.

Let W_j denote the waiting time of customer j in this system, where the first customer (who arrives to the empty system) has number 0. These W_j s satisfy the Lindley recurrence: $W_0 = 0$ and $W_j = \max(0, W_{j-1} + S_{j-1} - A_j)$ for $j \geq 1$ (Kleinrock 1975). We estimate μ by the sample average $Y = (W_0 + \dots + W_{t-1})/t$. To compute Y , we need to generate the $2(t-1)$ random variates $S_0, A_1, \dots, S_{t-1}, A_t$. This estimator Y is unbounded (so the Koksma-Hlawka inequality gives an infinite bound for it), but it has bounded variance.

We define a Markov chain that moves by one step each time one of these random variates is generated. That is, $X_0 = W_0$, $X_1 = W_0 + S_0$, $X_2 = W_1$, $X_3 = W_1 + S_1$, and so on. In this case, $d = 1$ and $s = 2(t-1)$. Later, we will consider the case where the chain moves by one step every

$d/2$ customers (where d is even), so $X_j = W_{jd/2}$ and $s = 2(t-1)/d$. In all cases, this Markov chain is stochastically increasing.

Our results are for $t = 100$, with $\rho = 0.2, 0.5$, and 0.8 . The MC variance per run σ^2 was estimated by making 100×2^{18} independent simulation runs and our best estimates of μ were obtained via array-RQMC with $n = 2^{20}$. These estimates are (with an accuracy up to the given digits): $\mu = 0.04922$ and $\sigma^2 = 0.0005393$ for $\rho = 0.2$, $\mu = 0.48000$ and $\sigma^2 = 0.06307$ for $\rho = 0.5$, and $\mu = 2.48004$ and $\sigma^2 = 3.1544$ for $\rho = 0.8$.

Table 1 reports the empirical VRFs for this example. The array-RQMC methods clearly outperform both MC and classical RQMC, even though classical RQMC is already significantly more efficient than MC (up to 100 times more efficient in one case). The improvement factor is larger when the queue has more traffic (i.e., for larger ρ , which is also when the variance is larger) and larger for the Sobol' nets than for the Korobov rules. For the Korobov rules, adding the baker's transform really makes a difference for both the classical and array methods. This shows that the transformation can be very effective even when the integrand is not smooth (as in the array-RQMC case). We have seen that in all the examples we tried. For this reason, we do not report results for Korobov rules without the baker's transform in what follows (with one exception).

For the Sobol' nets, the results are better when the points are enumerated in Gray code order. The corresponding scrambling appears helpful, yet we do not have a clear explanation of why.

The stratification provides a significant improvement compared with MC and is even better than classical RQMC, but not quite competitive with the other array-RQMC implementations. We implemented it only for the present example. An interesting fact is that the VRF with Array-Stratification is multiplied roughly by 2 when n is multiplied by 4 (i.e., from one column to the next), so it seems to be proportional to \sqrt{n} . This means that the variance appears to be proportional to $n^{-3/2}$, which corresponds exactly to the upper bound proved in Proposition 6.

For classical RQMC, the improvement is much better for the Korobov rules with the baker's transform than for the Sobol' nets. Without the baker's transform (not shown in the table), the Korobov rules are just slightly better than the Sobol' nets.

4.2. Increasing d

Suppose that at each step of the Markov chain, we generate d random variates to compute the waiting times of $d/2$ customers ($d = 1$ represents the case examined in the previous subsection). Then, the integral (10) approximated by array-RQMC at each step is a $(d+1)$ -dimensional integral and we anticipate that the integration error will increase with d . This is confirmed by the following results.

Table 2 shows the empirical VRFs for various values of d , with $n \approx 2^{18}$. For classical RQMC, the exact VRF

Table 1. Empirical VRFs of RQMC with respect to MC for the average waiting time of 100 customers in an $M/M/1$ queue with utilization factor ρ , with $n \approx 2^k$ points.

ρ		$k = 10$	$k = 12$	$k = 14$	$k = 16$	$k = 18$	$k = 20$
	a for classical Korobov-Baker	306	1,397	5,693	944	118,068	802,275
	a for array-Korobov	633	2,531	10,125	40,503	162,013	648,055
0.2	Classical-Korobov-Baker	5	8	15	16	59	117
	Classical-Sobol	1	1	3	1	13	28
	Array-Korobov	18	55	49	292	850	2,169
	Array-Korobov-Baker	43	159	306	991	3,168	10,590
	Array-Sobol	87	282	836	3,705	10,640	47,850
	Array-Sobol-NoGray	46	112	276	874	2,914	7,429
	Array-Stratification	12	17	31	57	104	183
0.5	Classical-Korobov-Baker	10	7	13	6	14	14
	Classical-Sobol	2	1	4	5	9	10
	Array-Korobov	14	46	33	231	686	2,034
	Array-Korobov-Baker	44	200	241	1,155	3,540	15,650
	Array-Sobol	123	504	1,083	5,651	13,830	55,160
	Array-Sobol-NoGray	55	130	302	1,188	3,507	11,260
	Array-Stratification	14	23	40	81	142	281
0.8	Classical-Korobov-Baker	11	2	15	17	21	26
	Classical-Sobol	3	2	4	6	10	11
	Array-Korobov	15	85	33	337	727	5,119
	Array-Korobov-Baker	70	463	287	2,225	10,080	75,920
	Array-Sobol	370	1,281	3,240	19,730	57,290	233,100
	Array-Sobol-NoGray	117	288	996	4,580	13,210	48,660
	Array-Stratification	21	40	77	153	246	535

does not depend on d and the variation observed in the table is only statistical noise; it gives an idea of the accuracy of our VRF estimators. For Array-Sobol, the VRF decreases with d , but not so fast. Moreover, the “NoGray” version becomes comparable to the regular one for $d > 2$. The VRFs are still substantial even for $d = 8$, where the array-RQMC method approximates 9-dimensional integrals at each step of the Markov chain.

Table 2. Estimated VRFs of classical RQMC and d -dimensional array-RQMC with respect to MC for the mean waiting time of 100 customers in an $M/M/1$ queue with utilization factor ρ , for selected values of d and $n \approx 2^{18}$.

ρ		$d = 1$	$d = 2$	$d = 4$	$d = 8$
0.2	Classical-Korobov-Baker	59	70	73	78
	Classical-Sobol	13	13	12	12
	Array-Korobov-Baker	3,168	571	283	137
	Array-Sobol	10,640	4,329	2,247	352
	Array-Sobol-NoGray	2,914	5,294	2,476	403
0.5	Classical-Korobov-Baker	14	22	16	18
	Classical-Sobol	9	6	9	7
	Array-Korobov-Baker	3,540	918	152	150
	Array-Sobol	13,830	8,067	5,566	667
	Array-Sobol-NoGray	3,507	6,206	5,205	702
0.8	Classical-Korobov-Baker	21	22	20	28
	Classical-Sobol	10	12	14	10
	Array-Korobov-Baker	10,080	2,296	1,074	597
	Array-Sobol	57,290	33,360	22,550	2,515
	Array-Sobol-NoGray	13,210	23,850	15,570	2,117

4.3. Random Dimension: A Regenerative System

So far in this example, s was fixed at $2(t - 1)$. We now modify the example so that $s = \infty$ (variable stopping time). Recall that the $M/M/1$ queue is a *regenerative system* that regenerates whenever a customer arrives to an empty system. Each regenerative cycle contains a random and unbounded number of customers. Suppose that we want to estimate $\mu = E[Y]$, where we take the following two possibilities for Y : (i) the total waiting time of all customers in a regenerative cycle, and (ii) the number of customers in a cycle whose waiting time exceeds c for some constant $c > 0$. In this case, slightly changing the values of U_j in (1) may split or merge regenerative cycles, making the function f highly discontinuous in both cases. Moreover, in case (ii), Y is integer valued, so it is not as smooth as in case (i). For our numerical illustration of case (ii), we take $c = 1$. The exact values of μ for case (i) are 0.0625, 1, and 16 for $\rho = 0.2, 0.5,$ and 0.8 (computed via standard queueing formulas). For case (ii), they are approximately 0.00458, 0.368, and 3.115 for $\rho = 0.2, 0.5,$ and 0.8 (estimated by simulation).

Tables 3 and 4 give the estimated VRFs of classical RQMC and array-RQMC compared with standard MC, again with $m = 100$. The improvement factors are not as large as in the two previous tables, but they are still significant, increase with n , and are much larger for the array versions than for the classical ones. The smaller improvements observed here could be due in part to the randomness of τ , as explained earlier: some chains reach τ much later than others, and the advantage of a good uniformity

INFORMS holds copyright to this article and distributed this copy as a courtesy to the author(s). Additional information, including rights and permission policies, is available at http://journals.informs.org/.

Table 3. Estimated VRFs for the regenerative $M/M/1$ queue with utilization factor ρ , case (i).

ρ		$k = 10$	$k = 12$	$k = 14$	$k = 16$	$k = 18$	$k = 20$
0.2	Classical-Korobov-Baker	3	5	6	5	14	24
	Array-Korobov-Baker	13	28	49	116	289	1,093
	Array-Sobol	7	21	46	99	239	756
0.5	Classical-Korobov-Baker	2	3	3	1	6	5
	Array-Korobov-Baker	11	16	37	79	159	438
	Array-Sobol	6	11	24	72	228	469
0.8	Classical-Korobov-Baker	1	1	2	1	2	2
	Array-Korobov-Baker	6	12	22	36	151	237
	Array-Sobol	3	5	19	32	92	225

of the point set $P'_{n,j}$ decreases with the number of chains left because we use a smaller and smaller subset of the points. The gain decreases with ρ in case (i) and increases with ρ in case (ii). Note that in case (ii), $\{Y > 0\}$ is a rare event when ρ is very small, so in that case something else (such as importance sampling) would have to be done to reduce the variance. The fact that the gain decreases with ρ in case (i) here whereas the opposite was true in Table 1 for the same performance measure might be an indication of a loss of efficiency due to a larger variance of τ when ρ increases. For classical RQMC, we need infinite-dimensional RQMC point sets because the number of steps of the chain is unbounded, so we cannot use Sobol' nets.

4.4. Markovian Queues in Series

For an example with a higher-dimensional state space, we consider a system of l Markovian queues in series. Customers arrive to the first queue according to a Poisson process with rate λ , go through queue 1, then queue 2, etc., in FIFO order. Service rate at queue q is $\mu_q > \lambda$, for $q = 1, \dots, l$. We uniformize this continuous-time Markov chain (CTMC) model so that the global transition rate is always $\nu = \lambda + \mu_1 + \dots + \mu_l$. At each step, the next transition is an arrival with probability λ/ν and a service completion at queue q with probability μ_q/ν . A service completion at an empty queue is just a dummy event that does not change the state of the system. The embedded discrete-time Markov chain is $\{X_j = (N_{j,1}, \dots, N_{j,l}), j \geq 0\}$, where $N_{j,q}$ is the number of customers in queue q just after the j th transition. Therefore, if we are in state X_j at step j , with

probability λ/ν the next state is $X_{j+1} = (N_{j,1} + 1, \dots, N_{j,l})$, and with probability μ_q/ν the next state is

$$X_{j+1} = \begin{cases} (N_{j,1}, \dots, N_{j,l}) & \text{(unchanged) if } N_{j,q} = 0, \\ (N_{j,1}, \dots, N_{j,q} - 1, N_{j,q+1} + 1, \dots, N_{j,l}) & \text{if } q < l \text{ and } N_{j,q} > 0, \\ (N_{j,1}, \dots, N_{j,l} - 1) & \text{if } q = l \text{ and } N_{j,l} > 0. \end{cases}$$

We assume that the system starts empty and evolves for t transitions, so $X_0 = (0, \dots, 0)$ and $\tau = t$ in (2). The performance measure μ that we want to estimate is the expected average number of customers in the system over the first t transitions. The corresponding cost function at each step is $c_j(X_j) = N_{j,1} + \dots + N_{j,l}$. Each transition is simulated from a single uniform $U_j \in [0, 1)$ (so $d = 1$) as follows: if $U_j > 1 - \lambda/\nu$, we have an arrival; otherwise, we have a service completion at queue $q = \min\{i: U_j \leq (\mu_1 + \dots + \mu_i)/\nu\}$.

The sorting function h is defined by

$$h(N_{j,1}, \dots, N_{j,l}) = N_{j,l}(1 + \epsilon) + N_{j,l-1}(1 + \epsilon^2) + \dots + N_{j,1}(1 + \epsilon^l),$$

where ϵ is a very small constant (e.g., 10^{-6}). This has the effect of sorting the states primarily according to the total number of customers in the system, then (in case of equalities) according to the number of customers in the last queue, then by the number in the next-to-last queue, and so on. This choice is only a heuristic to regroup "similar" states.

Table 4. Estimated VRFs for the regenerative $M/M/1$ queue with utilization factor ρ , case (ii).

ρ		$k = 10$	$k = 12$	$k = 14$	$k = 16$	$k = 18$	$k = 20$
0.2	Classical-Korobov-Baker	1	1	2	2	3	2
	Array-Korobov-Baker	3	5	15	22	72	113
	Array-Sobol	2	5	9	23	46	108
0.5	Classical-Korobov-Baker	3	3	4	2	7	6
	Array-Korobov-Baker	22	35	146	253	540	1,655
	Array-Sobol	13	33	85	245	645	1,847
0.8	Classical-Korobov-Baker	2	1	3	2	3	3
	Array-Korobov-Baker	16	40	100	76	442	997
	Array-Sobol	10	27	81	198	629	1,844

Table 5. Estimated VRFs for the average number of customers in a network of l queues in series, with $n = 2^k$ points.

l		$k = 10$	$k = 12$	$k = 14$	$k = 16$	$k = 18$	$k = 20$
2	Classical-Korobov-Baker	6	1	16	11	16	12
	Classical-Sobol	3	3	5	8	11	9
	Array-Korobov-Baker	34	82	316	476	1,733	3,233
	Array-Sobol	40	143	480	1,104	3,785	14,340
3	Classical-Korobov-Baker	8	2	15	13	20	15
	Classical-Sobol	4	5	8	7	12	14
	Array-Korobov-Baker	13	78	196	190	884	783
	Array-Sobol	21	55	169	479	1,884	4,663

We tried two numerical examples. The first one has $l = 2$ (a tandem queue), $\lambda = 1$, $\mu_1 = 1.75$, $\mu_2 = 1.25$, and $t = 200$. The second one has $l = 3$, $\lambda = 6$, $\mu_1 = 10$, $\mu_2 = 9$, $\mu_3 = 7$, and $t = 200$. Table 5 reports the estimated VRFs of RQMC compared with standard MC. Array-RQMC clearly outperforms classical RQMC for both examples, and Array-Sobol does better than Array-Korobov-Baker. We also observe a smaller variance reduction for the three-dimensional example than for the two-dimensional one.

We then repeated the same experiment, but with μ defined as the fraction of the time where the number of customers at the last queue exceeds some constant K ; that is, with $c_j(X_j) = I[N_{j,l} > K]$. We took $K = 4$ for the tandem queue example and $K = 6$ for the three-queue example. The results are in Table 6. They are quite similar to those of Table 5.

4.5. Pricing an Asian Option

We consider the pricing of an Asian option on a single asset. Under the risk-neutral measure, the asset price at time t , denoted by $S(t)$, obeys a geometric Brownian motion: $dS(t) = rS(t)dt + \sigma S(t)dB(t)$, where r is the risk-free interest rate, σ is the (risk-neutral) volatility parameter, and $B(\cdot)$ is a standard Brownian motion. The option's value can be written as $\mu = E[e^{-rT} C_a(T) | S(0)]$, where

$$C_a(T) = \max \left[0, \left(\frac{1}{s} \sum_{j=1}^s S(t_j) \right) - K \right]$$

and $0 < t_1 < \dots < t_s = T$ are the discrete observation times. See, e.g., Hull (2000) and Glasserman (2004) for further details. We have

$$S(t_j) = S(t_{j-1}) \exp[(r - \sigma^2/2)(t_j - t_{j-1}) + \sigma(t_j - t_{j-1})^{1/2} \Phi^{-1}(U_j)] \tag{18}$$

for $j = 1, \dots, s$, where the U_j s are independent $U(0, 1)$ random variables and Φ is the standard normal distribution function. To get an unbiased estimator of μ , it suffices to generate $S(t_1), \dots, S(t_s)$ via (18), with s i.i.d. $U(0, 1)$ random variates, and compute the estimator $X = e^{-rT} C_a(T)$.

To apply the array-RQMC method, we define the state of the chain at step j as the two-dimensional vector $X_j = (S(t_j), \bar{S}_j)$, where $\bar{S}_j = (S(t_1) + \dots + S(t_j))/j$. We order these states simply by increasing order of their value of $S(t_j)$, i.e., we define the sorting function by $h(x_1, x_2) = x_1$ (there are other possibilities, probably better ones).

For a numerical illustration, let $S(0) = 100$, $K = 90$, $T = 240/365$ (in years), $t_j - t_{j-1} = 1/365$ for all j , $t_1 = T - (s - 1)/365$, $r = \ln 1.09$, $\sigma = 0.2$, and $s = 10, 60, 120$, and 240 . Table 7 gives the estimated VRFs of RQMC compared with standard MC. The classical RQMC methods uses the straightforward simulation approach described above. Efficiency can be further improved by combining RQMC with bridge sampling and other variance-reduction techniques such as control variates and importance sampling (Caffisch and Moskowitz 1995, Glasserman 2004, L'Ecuyer and Lemieux 2000, L'Ecuyer 2004a), but we do not go in that direction.

Table 6. Estimated VRFs for the fraction of the time where the number of customers at the last queue exceeds K for a network of l queues in series, with $n = 2^k$ points.

l		$k = 10$	$k = 12$	$k = 14$	$k = 16$	$k = 18$	$k = 20$
2	Classical-Korobov-Baker	4	4	9	10	11	7
	Classical-Sobol	4	6	4	9	9	6
	Array-Korobov-Baker	39	161	595	929	1,477	6,408
	Array-Sobol	55	226	582	2,438	8,659	38,030
3	Classical-Korobov-Baker	4	2	4	5	5	5
	Classical-Sobol	4	3	3	4	7	5
	Array-Korobov-Baker	11	46	134	358	610	607
	Array-Sobol	16	50	132	450	2,091	6,082

INFORMS holds copyright to this article and distributed this copy as a courtesy to the author(s). Additional information, including rights and permission policies, is available at http://journals.informs.org/.

Table 7. Estimated VRFs for an Asian option with s observation times, with $n = 2^k$ points.

s		$k = 10$	$k = 12$	$k = 14$	$k = 16$	$k = 18$	$k = 20$
10	Classical-Korobov	188	594	2,601	5,505	18,050	11,040
	Classical-Korobov-Baker	2,629	10,600	5,104	83,450	27,560	93,620
	Classical-Sobol	4,844	11,460	28,740	46,020	142,900	222,800
	Array-Korobov-Baker	4,783	13,280	23,960	45,990	36,670	39,950
	Array-Sobol	5,080	13,030	37,460	38,320	36,360	32,430
60	Classical-Korobov	81	17	352	406	552	497
	Classical-Korobov-Baker	481	567	919	610	1,362	1,745
	Classical-Sobol	282	488	907	787	1,654	2,413
	Array-Korobov-Baker	1,187	1,742	1,218	2,231	1,680	1,998
	Array-Sobol	1,234	1,742	2,050	2,203	2,189	1,866
120	Classical-Korobov	73	27	152	209	252	276
	Classical-Korobov-Baker	244	380	452	407	581	498
	Classical-Sobol	68	92	234	253	531	410
	Array-Korobov-Baker	816	1,263	1,355	1,736	1,456	1,635
	Array-Sobol	1,423	1,485	1,260	1,390	1,333	1,477
240	Classical-Korobov	30	9	93	116	95	148
	Classical-Korobov-Baker	76	167	233	303	375	319
	Classical-Sobol	29	32	54	69	151	217
	Array-Korobov-Baker	445	703	375	758	773	601
	Array-Sobol	744	769	670	725	702	667

Here, classical RQMC is already very effective when s is small, and array-RQMC is not really better. For larger s , however, array-RQMC eventually provides larger VRFs, especially when k is small. For $s = 240$ and $k = 10$, for example, the variance is approximately seven times smaller for Array-Sobol than for the best classical RQMC method. On the other hand, it is disappointing to see that the VRFs eventually stabilize when we increase n . This suggests a $O(1/n)$ asymptotic rate of convergence of the variance for this particular implementation and choice of sorting function h . Another important point to notice is that, once again, the baker's transformation applied on top of the Korobov rules really helps. In previous experiments with this Asian option example, the lattice rules were used only without the baker's transformation (L'Ecuyer and Lemieux 2000).

4.6. Estimating a Small Ruin Probability with Importance Sampling and Array-RQMC

A (simplified) insurance company receives premiums at constant rate $c > 0$ and claims according to a Poisson process $\{N(t), t \geq 0\}$ with rate $\lambda > 0$. The claim sizes C_j , $j \geq 1$, are i.i.d. random variables with density h . The reserve (amount of money in hand) at time t is

$$R(t) = R(0) + ct - \sum_{j=1}^{N(t)} C_j,$$

where $R(0)$ is the initial reserve. We want to estimate the ruin probability, i.e., the probability μ that $R(t)$ eventually becomes negative. Ruin can occur only at the time of a claim. The reserve just after claim j is $X_j = X_{j-1} + A_j c - C_j$, $j \geq 1$, where $X_0 = R(0)$ and A_j is the time between claims $j - 1$ and j . The process $\{X_j, j \geq 0\}$ is

a random walk. This process cannot be simulated directly to estimate the ruin probability because: (1) we cannot be 100% sure that ruin does not occur if we simulate only for a finite time; and (2) in practice, ruin occurs very rarely. We can get around these difficulties by using importance sampling (IS) with exponential twisting as follows (Asmussen 1985). Assuming that h has a finite moment generating function M_h around zero, we replace the density $h(x)$ by $h_\theta(x) = h(x)e^{\theta x} / M_h(\theta)$ and increase the rate λ of the Poisson process to $\lambda_\theta = \lambda + \theta c$, where θ is the largest solution to the Lundberg equation $M_h(\theta) = (\lambda + \theta c) / \lambda$. Under this IS scheme, ruin occurs with probability one and the (unbiased) estimator of μ is

$$e^{\theta(X_\tau - X_0)}, \tag{19}$$

where $\tau = \inf\{j: X_j < 0\}$, a random stopping time. Here, there is no need to store the intermediate values of the likelihood ratio during the simulation because its final value depends only on X_τ . Thus, the state space is one-dimensional. We are interested in seeing if a combination of IS with array-RQMC is more efficient than IS alone. A priori, because the function $f(\mathbf{u}) = e^{\theta(X_\tau - X_0)}$ is sawtooth-like (not smooth at all) with respect to each coordinate of \mathbf{u} , we do not expect RQMC to perform well.

For a numerical experiment, we take $\lambda = 1$, exponential claim sizes with mean $1/\beta = 2$, and $R(0) = 200$. We use $d = 1$, i.e., one step of the chain each time a uniform random number is generated. The number of steps before ruin occurs is random. Table 8 gives the estimated VRFs compared with MC (with IS) for $c = 3, 5$, and 10. For classical RQMC, we need an infinite-dimensional point set; this rules out classical-Sobol. The exact ruin probability μ is approximately $\mu = 2.2 \times 10^{-15}$ for $c = 3$, $\mu = 3.5 \times 10^{-27}$

Table 8. Estimated VRFs for the ruin probability example, with inflow rate c and $n = 2^k$ points.

c		$k = 10$	$k = 12$	$k = 14$	$k = 16$	$k = 18$	$k = 20$
10	Classical-Korobov-Baker	1	1	1	1	1	1
	Array-Korobov-Baker	3	3	7	3	15	27
	Array-Sobol	2	2	6	10	19	45
5	Classical-Korobov-Baker	1	1	1	1	2	1
	Array-Korobov-Baker	3	4	10	5	21	37
	Array-Sobol	2	4	8	13	33	73
3	Classical-Korobov-Baker	1	1	1	1	1	1
	Array-Korobov-Baker	2	4	8	7	24	38
	Array-Sobol	2	5	7	17	30	49

for $c = 5$, and $\mu = 3.6 \times 10^{-36}$ for $c = 10$. The gains are not as spectacular as for the previous examples, but they are nevertheless significant for large n .

5. Conclusion

We have proposed a new RQMC method for Markov chains, proved results on its convergence for special cases, and tested it numerically. The new method provides large efficiency gains compared with standard MC in our examples. It performs better than classical RQMC in the examples where the Markov chain has a one-dimensional state space and evolves over several steps. Generally speaking, the performance of the array-RQMC method tends to degrade when the integrand has higher variability, or when the dimension of the state space becomes larger than one and there is no natural (or obvious) sorting function for the states. However, even in these cases, there can be significant variance reductions compared with standard MC, and sometimes compared with classical RQMC as well. Our paper also provides novel empirical evidence of the effectiveness of applying the baker's transformation over a randomly shifted lattice rule, an idea that was studied theoretically by Hickernell (2002).

From the practical viewpoint, an interesting challenge is to find good ways of defining the sorting function h for specific classes of problems where the Markov chain has a multidimensional state space. Our ongoing and future work also includes studying the application of array-RQMC to other settings that fit a general Markov chain framework, such as Markov chain Monte Carlo methods and stochastic approximation algorithms.

Appendix

Worst-Case Error Bounds on the Expected Average Cost

PROOF OF PROPOSITION 3. Using integration by parts for the third equality, we get

$$\begin{aligned}
 |\bar{Y}_{n,j} - E[c_j(X_j)]| &= \left| \frac{1}{n} \sum_{i=0}^{n-1} c_j(X_{i,j}) - E[c_j(X_j)] \right| \\
 &= \left| \int_{-\infty}^{\infty} c_j(z) d\hat{F}_j(z) - \int_{-\infty}^{\infty} c_j(z) dF_j(z) \right|
 \end{aligned}$$

$$\begin{aligned}
 &= \left| \int_{-\infty}^{\infty} (\hat{F}_j(z) - F_j(z)) dc_j(z) \right| \\
 &\leq \int_{-\infty}^{\infty} \Delta_j |dc_j(z)| \\
 &\leq \Delta_j V(c_j). \quad \square
 \end{aligned}$$

Histogram and Integral Discrepancies

We introduce notions of histogram and integral discrepancies, and examine their properties in this subsection. This will be used to obtain bounds on Δ_j .

A histogram with L intervals over the unit square is defined as the region under a step function over the interval $[0,1]$: partition the unit interval $[0, 1)$ at the bottom of the square into L subintervals, say of lengths q_1, \dots, q_L , where $q_1 + \dots + q_L = 1$. Over the i th interval, put a rectangle of height h_i , where $0 \leq h_i \leq 1$, and with the same width as the interval. The histogram H is the union of these rectangles. We say that the histogram is *monotone* (increasing or decreasing) if $h_1 \leq \dots \leq h_L$ or $h_1 \geq \dots \geq h_L$. Let $\mathcal{H}(L)$ be the family of all histograms with L intervals over the unit square, and $\mathcal{H}^+(L)$ the subfamily of all monotone histograms. The L -histogram discrepancy of a point set P_n in the unit square is defined as

$$D_h(L, P_n) = \sup_{H \in \mathcal{H}(L)} |\text{card}(P_n \cap H)/n - \text{area}(H)|,$$

where $\text{area}(H) = \sum_{i=1}^L q_i h_i$ denotes the area of H . If $\mathcal{H}(L)$ is replaced by $\mathcal{H}^+(L)$, we get the L -staircase discrepancy of P_n , denoted $D_h^+(L, P_n)$.

LEMMA 1. Let P_n denote the first n points of a two-dimensional low-discrepancy sequence whose star discrepancy satisfies $D_n^*(P_n) = O(n^{-1} \log n)$. Then, there is a constant K such that for all L ,

$$D_h^+(L, P_n) \leq D_h(L, P_n) \leq LKn^{-1} \log n.$$

PROOF. In every histogram H , each rectangle can be written as a difference of two rectangular boxes anchored at the origin. Thus, H can be written as a sum and difference of $2L$ such boxes. However, we know that the star discrepancy of P_n is in $O(n^{-1} \log n)$, and the last inequality follows. The first inequality is obvious. \square

Several two-dimensional sequences that satisfy this requirement are readily available (Niederreiter 1992); for example, one can take the two-dimensional Sobol' sequence. However, the bound in Lemma 1 is linear in L , so it is not very useful for Markov chains with large state spaces (assuming that each state of the chain is associated with one of the L subintervals, and vice versa). The next lemma provides a bound that does not depend on L . It is based on Assumption 2 and a notion of integral discrepancy obtained when $L \rightarrow \infty$.

Let $\mathcal{F}(v)$ denote the set of functions $f: [0, 1] \rightarrow [0, 1]$ such that $V(f) \leq v$, where $V(f)$ is the variation of f , defined in (4). It is well known that a function of bounded variation over a given interval is Riemann integrable over that interval. For $f: [0, 1] \rightarrow [0, 1]$, let $H(f) = \{(x, y) \in [0, 1]^2: 0 \leq y \leq f(x)\}$ be the area under f . For a point set P_n , we define the *integral discrepancy at variation v* by

$$D_i(P_n, v) = \sup_{f \in \mathcal{F}(v)} |\text{card}(P_n \cap H(f))/n - \text{area}(H(f))|. \quad (20)$$

(An equivalent quantity is defined in Niederreiter 1992, p. 17.) If f has bounded variation, $H(f)$ can be approximated arbitrarily closely by a histogram H having L rectangles of heights $h_1 = f(x_1), \dots, h_L = f(x_L)$, where $0 < x_1 < \dots < x_L = 1$ for some large L . If $V(f) \leq v$, the step function f_H that corresponds to this H has total variation

$$V(f_H) = \sum_{i=2}^L |h_i - h_{i-1}| \leq V(f) \leq v.$$


Hence, we have that

$$D_i(P_n, v) = \sup_{L \geq 1, H \in \mathcal{H}(L), V(f_H) \leq v} |\text{card}(P_n \cap H)/n - \text{area}(H)|.$$

LEMMA 2. *Under Assumption 2, for any $v \geq 0$, we have $D_i(P_n, v) \leq (v + 1)n^{-1/2}$.*

PROOF. Consider a function $f: [0, 1] \rightarrow [0, 1]$ with $V(f) \leq v$. We define the extended graph of f , denoted $G(f)$, as the boundary between $H(f)$ and $[0, 1]^2 \setminus H(f)$. This is the graph of f , $\{(x, f(x)): 0 \leq x \leq 1\}$, to which we add vertical lines that link the graph pieces where f is discontinuous. The idea of the proof is to bound the number of subsquares that intersect $G(f)$ and then bound the error in terms of this number. The n squares are partitioned into \sqrt{n} columns that correspond to \sqrt{n} intervals on the horizontal axis. Suppose that $G(f)$ goes into l_j different subsquares in column j . Clearly, these l_j subsquares must form a connected rectangle (see Figure A.1). For any of these subsquares S , we have an overestimation of $\text{area}(H(f) \cap S)$ if the point of P_n lying in subsquare S is in $H(f)$, and underestimation otherwise. The total error in any given column is the total amount of overestimation minus the total amount of underestimation. Let $S_{t,j}$ and $S_{b,j}$ denote the top and bottom

Figure A.1. Illustration of the proof of Lemma 2.



Notes. The shaded surface is $H(f)$ for a given column j . Here, $l_j = 3$, we overestimate in each of those three subsquares, and $V_j = (b - a) + (b - c)$.

subsquares from this rectangle, respectively. They may be the same subsquare if $l_j = 1$.

Suppose that we are overestimating in $S_{t,j}$ (as in Figure A.1). If we are underestimating in $S_{b,j}$, or if the top and bottom subsquares are the same, the combined error in the top and bottom subsquares cannot exceed $1/n$, so the error in column j is at most $(l_j - 1)/n$. Otherwise—i.e., if we are also overestimating in $S_{b,j}$ and $l_j \geq 2$ (as in Figure A.1)—then the error in the bottom subsquare is the surface of this subsquare, which is above $G(f)$. This surface cannot exceed $V_{b,j}n^{-1/2}$, where $V_{b,j}$ is the variation of f in this subsquare. Then, in this second case, the error in column j is at most $(l_j - 1)/n + V_{b,j}n^{-1/2}$. However, in both cases, the total variation V_j in column j satisfies $V_j \geq (l_j - 2)n^{-1/2} + V_{b,j}$, so the error in column j cannot exceed $V_jn^{-1/2} + n^{-1}$. This same error bound can be obtained by a symmetrical argument in the case where we are underestimating in $S_{t,j}$.

By adding these inequalities over all columns, we obtain that the total error cannot exceed $V(f)n^{-1/2} + n^{1/2}n^{-1} = (V(f) + 1)n^{-1/2} \leq (v + 1)n^{-1/2}$. By taking the sup over f as in (20), the result follows. \square

Let \mathcal{F}^+ be the set of monotone functions $f: [0, 1] \rightarrow [0, 1]$ and define

$$D_i^+(P_n) = \sup_{f \in \mathcal{F}^+} |\text{card}(P_n \cap H(f))/n - \text{area}(H(f))|.$$

For $f \in \mathcal{F}^+$, $V(f)$ cannot exceed one. This gives:

COROLLARY 1. *Under Assumption 2, we have $D_i^+(P_n) \leq 2n^{-1/2}$.*

To see that this discrepancy bound is tight, consider the constant function $f(x) = n^{-1/2}(1 + \epsilon)$ for a small $\epsilon > 0$, and suppose that each column has two points in $H(f)$. Then, $V(f) = 0$ and the error is $(1 - \epsilon)n^{-1/2}$, which can be arbitrarily close to the bound $n^{-1/2}$. Theorem 1 of Hlawka (1971) yields an alternative bound for this situation: if we apply it with $s = 2$ and a two-dimensional function that equals one in the histogram and zero elsewhere, we get the looser bound $D_1^+(P_n) \leq 18n^{-1/2}$.

Instead of asking only for the points of P_n to be evenly distributed among the n subsquares, we could have the stronger requirement that they form a $(0, k, 2)$ -net in base 2, assuming that $n = 2^k$ for some integer k . This means that for every partition of the unit square into n rectangles of width 2^{-q} and height 2^{-k+q} for some $q = 0, 1, \dots, k$, every rectangle contains exactly one point. Consider the function $f(x) = x$. Proposition 5 of Lécot (1996) shows that for this example, there is a $(0, k, 2)$ -net (namely the Hammersley point set in base 2) for which the error is $n^{-1/2}/2$. In other words, the rate of the bound of Lemma 2 is tight even under this stronger assumption.

Error Bounds on the State Distribution

At step j of the Markov chain, define $\Delta_j(z) = \widehat{F}_j(z) - F_j(z)$, so $\Delta_j = \sup_{z \in \mathbb{R}} |\Delta_j(z)|$. For $x \in \mathcal{X}$ and $z \in \mathbb{R}$, let

$$F_j(z | x) = P[X_j \leq z | X_{j-1} = x],$$

$$\Lambda_j(z) = V(F_j(z | \cdot)) \stackrel{\text{def}}{=} \int_{-\infty}^{\infty} |dF_j(z | x)|, \quad \text{and}$$

$$\Lambda_j = \sup_{z \in \mathbb{R}} \Lambda_j(z),$$

where the differential $dF_j(z | x)$ in the definition of $\Lambda_j(z)$ is with respect to x . Thus, $\Lambda_j(z)$ is the total variation of the function $F_j(z | \cdot)$. If the Markov chain is stochastically monotone, which means that $F_j(z | y) = P[X_j \leq x | X_{j-1} = y]$ is monotone in y for each j , then Λ_j cannot exceed one.

Let \widehat{F}_j be the empirical distribution of the states of the n copies of the chain at step j and

$$\widetilde{F}_j(z) = \int_{-\infty}^{\infty} F_j(z | x) d\widehat{F}_{j-1}(x) = \frac{1}{n} \sum_{i=0}^{n-1} F_j(z | X_{(i), j-1}),$$


so \widetilde{F}_j is the conditional distribution function of X_j when X_{j-1} is generated from \widehat{F}_{j-1} . The value of $\widetilde{F}_j(z)$ is equal to the area of the histogram $H_{j,z}$ whose height over the interval $[i/n, (i+1)/n]$ is $F_j(z | X_{(i), j-1})$ for $i = 0, \dots, n-1$. This histogram $H_{j,z}$ is inside the unit square $[0, 1]^2$ (see Figure A.2 for an illustration). We also have that

$$V[f_{H_{j,z}}] = \sum_{i=1}^{n-1} |F_j(z | X_{(i-1), j-1}) - F_j(z | X_{(i), j-1})| \leq \Lambda_j(z) \leq \Lambda_j, \tag{21}$$

where we have exploited the fact that the states are sorted by increasing order.

The proofs of Propositions 4 and 5 are based on the following lemma.

Figure A.2. An example of a histogram $H_{j,z}$ for $n = 4$.



Notes. The histogram is the shaded area, with height $F_j(z | X_{(i), j-1})$, and $\widehat{F}_j(z)$ is the fraction of $P'_{n,j}$ that falls in $H_{j,z}$.

LEMMA 3. Under Assumption 1, we have

$$\Delta_j \leq \sum_{k=1}^j \sup_{z \in \mathbb{R}} |\widehat{F}_k(z) - \widetilde{F}_k(z)| \prod_{i=k+1}^j \Lambda_i, \tag{22}$$

where an empty product is assumed to be one. We also have

$$\widehat{F}_j(z) - \widetilde{F}_j(z) = \text{card}(P'_{n,j} \cap H_{j,z})/n - \text{area}(H_{j,z}) \tag{23}$$

and

$$|\widehat{F}_j(z) - \widetilde{F}_j(z)| \leq D_i(P'_{n,j}, V(H_{j,z})) \leq D_i(P'_{n,j}, \Lambda_j(z)). \tag{24}$$

PROOF. At step j , we have

$$|\Delta_j(z)| = |\widehat{F}_j(z) - F_j(z)| \leq |\widehat{F}_j(z) - \widetilde{F}_j(z)| + |\widetilde{F}_j(z) - F_j(z)|. \tag{25}$$

To bound Δ_j , we will bound the two quantities on the right of (25). We have

$$\begin{aligned} \widetilde{F}_j(z) - F_j(z) &= \int_{-\infty}^{\infty} F_j(z | x) d\widehat{F}_{j-1}(x) - \int_{-\infty}^{\infty} F_j(z | x) dF_{j-1}(x) \\ &= \int_{-\infty}^{\infty} (F_{j-1}(x) - \widehat{F}_{j-1}(x)) dF_j(z | x), \end{aligned} \tag{26}$$

where the second equality is obtained via integration by parts. Therefore,

$$\begin{aligned} |\widetilde{F}_j(z) - F_j(z)| &\leq \int_{-\infty}^{\infty} |\widehat{F}_{j-1}(x) - F_{j-1}(x)| |dF_j(z | x)| \\ &\leq \Lambda_j(z) \Delta_{j-1} \leq \Lambda_j \Delta_{j-1}. \end{aligned} \tag{27}$$

From (27) and (25), we have

$$\Delta_j = \sup_{z \in \mathbb{R}} |\widehat{F}_j(z) - F_j(z)| \leq \Lambda_j \Delta_{j-1} + \sup_{z \in \mathbb{R}} |\widehat{F}_j(z) - \widetilde{F}_j(z)|.$$

Together with the fact that $\Delta_0 = 0$, this gives (22).

We recall that $\tilde{F}_j(z)$ is the area of the histogram $H_{j,z}$ and observe that $\hat{F}_j(z)$ is the fraction of the points of $P'_{n,j}$ that fall in $H_{j,z}$ (Figure A.2). This gives (23), which implies (24) using (20). \square

PROOF OF PROPOSITION 4. Using Lemmas 1 and 3, we find that $|\hat{F}_j(z) - \tilde{F}_j(z)| \leq LKn^{-1} \log n$ for some constant K , and then

$$\Delta_j \leq LKn^{-1} \log n \sum_{k=1}^j \prod_{i=k+1}^j \Lambda_i \quad (28)$$

for all j . \square

PROOF OF PROPOSITION 5. Assuming that $P'_{n,j}$ satisfies Assumption 2, Lemmas 2 and 3 assure us that for all $j \geq 1$ and $z \in \mathbb{R}$, $|\hat{F}_j(z) - \tilde{F}_j(z)| \leq (\Lambda_j(z) + 1)n^{-1/2}$ and

$$\Delta_j \leq n^{-1/2} \sum_{k=1}^j (\Lambda_k + 1) \prod_{i=k+1}^j \Lambda_i. \quad (29)$$

If the Markov chain is also stochastically increasing, i.e., $P[X_j \geq x | X_{j-1} = y]$ is nondecreasing in y for each j , then $F_j(z | y)$ is nondecreasing in y , so $\Lambda_j \leq 1$ for each j , and the bound (29) becomes $\Delta_j \leq 2jn^{-1/2}$. \square

Variance Bounds for Array-RQMC

Proposition 5 gives a worst-case deterministic bound of $O(1/n)$ for the square error of $\bar{Y}_{n,j}$, in contrast with an expected square error of $O(1/n)$ for ordinary Monte Carlo. In what follows, we obtain a better bound on the convergence rate of the variance by exploiting randomization in a setting where Assumption 3 holds.

PROOF OF PROPOSITION 6. Define

$$\Gamma_j = \frac{1}{4} \sum_{k=1}^j (\Lambda_k + 1) \prod_{i=k+1}^j \Lambda_i^2. \quad (30)$$

We will prove that for each j and z ,

$$\text{Var}[\hat{F}_j(z) - \tilde{F}_j(z)] \leq (\Lambda_j + 1)n^{-3/2}/4, \quad (31)$$

$$\text{Var}[\hat{F}_j(z) - F_j(z)] \leq \Gamma_j n^{-3/2}, \quad (32)$$

and use this to show the final result.

To prove (31), we first recall that $\tilde{F}_j(z)$ is the area of the histogram $H_{j,z}$, whereas $\hat{F}_j(z)$ is the fraction of $P'_{n,j}$ that falls in this histogram. We enumerate the n subsquares with $i = 0, \dots, n-1$. Let S_i be the i th subsquare and $\delta_j(z, i) = \text{card}(P'_{n,j} \cap H_{j,z} \cap S_i) - n \text{area}(H_{j,z} \cap S_i)$. We have

$$\begin{aligned} \hat{F}_j(z) - \tilde{F}_j(z) &= \text{card}(P'_{n,j} \cap H_{j,z})/n - \text{area}(H_{j,z}) \\ &= \sum_{i=0}^{n-1} \delta_j(z, i)/n. \end{aligned}$$

For any given j and each i , $\delta_j(z, i)$ is a Bernoulli random variable minus its mean, so $E[\delta_j(z, i)] = 0$ and

$\text{Var}[\delta_j(z, i)] \leq 1/4$. These Bernoulli random variables have nonzero variance only for the subsquares that intersect the histogram boundary (that separates $H_{j,z}$ from $[0, 1]^2 \setminus H_{j,z}$) because for the other subsquares they are constant.

Here we consider a fixed j and drop the subscript j for a moment. Let $S_{b,c}$ and $S_{t,c}$ denote the lowest and highest subsquares that intersect the histogram boundary in column c , and let $V_{b,c}$ and $V_{t,c}$ be $n^{1/2}$ times the variation of the histogram in these two subsquares. Let p_c be n times the area of $S_{t,c}$ contained in the histogram and q_c be n times the area of $S_{b,c}$ not contained in the histogram. We suppose, until indicated otherwise, that $S_{b,c}$ and $S_{t,c}$ are not the same subsquare. Then, the Bernoulli variables $\delta_j(z, i)$ that correspond to these two subsquares are independent and have variances bounded by $p_c(1-p_c)$ and $q_c(1-q_c)$ (which cannot exceed $1/4$), respectively. We will now prove the following bound on the sum of their variances:

$$p_c(1-p_c) + q_c(1-q_c) \leq \frac{V_{b,c} + V_{t,c} + 1}{4}. \quad (33)$$

Denote $v = V_{b,c} + V_{t,c}$. If $v \geq 1$, (33) holds trivially, so let us assume that $v \leq 1$. At any given point on the horizontal axis, the histogram boundary cannot be in $S_{b,c}$ and $S_{t,c}$ at the same time. Let ρ be the fraction of the horizontal interval in column c where the histogram is in $S_{t,c}$. Then, we have $p_c \leq \rho V_{t,c}$ and $q_c \leq (1-\rho)V_{b,c}$. We now observe that either $\rho V_{t,c} \leq v/4$ or $(1-\rho)V_{b,c} \leq v/4$. To see this, consider a rectangle of width one and height v , with the bottom-left corner at point $(0, 0)$, divided into four subrectangles by a vertical line at ρ and a horizontal line at $V_{t,c}$. The quantities $\rho V_{t,c}$ and $(1-\rho)V_{b,c}$ are the surfaces of two opposite subrectangles of this rectangle, so these surfaces cannot both be larger than a quarter of the rectangle's surface v . Indeed, suppose that $\rho V_{t,c} > v/4$. Then, $V_{t,c} > v/(4\rho)$, and therefore $(1-\rho)(v - V_{t,c}) > v/4$ would imply that $(1-\rho)(v - v/(4\rho)) > v/4$, i.e., $0 < (1-\rho)(4\rho - 1) - \rho = -(2\rho - 1)^2$, which is a contradiction. Therefore, we have either $p_c(1-p_c) \leq v/4$ or $q_c(1-q_c) \leq v/4$, and because these two quantities never exceed $1/4$, the bound (33) follows. If $S_{b,c}$ and $S_{t,c}$ are the same subsquare, the variance of $\delta_j(z, i)$ in this subsquare cannot exceed $1/4$.

If other subsquares intersect the histogram in column c , between $S_{b,c}$ and $S_{t,c}$, then in each of these subsquares the histogram variation is at least $n^{-1/2}$ and the variance of the corresponding $\delta_j(z, i)$ is at most $1/4$. By adding the above inequalities over all the columns, we obtain that the sum (over i) of variances of all Bernoulli variables $\delta_j(z, i)$ is bounded by $(\Lambda_j + 1)n^{1/2}/4$.

Because these $\delta_j(z, i)$ s are pairwise independent across the different values of i , we obtain

$$\begin{aligned} \text{Var}[\hat{F}_j(z) - \tilde{F}_j(z)] &= \sum_{i=0}^{n-1} \text{Var}[\delta_j(z, i)/n] \\ &\leq n^{1/2}(\Lambda_j + 1)/(4n^2) \\ &= (\Lambda_j + 1)n^{-3/2}/4, \end{aligned}$$

which proves (31).

We now prove (32) by induction on j . It obviously holds for $j = 0$ because $\hat{F}_0 = F_0$. Suppose that it holds for $j - 1$ for all z . Observe that from the proof of Proposition 1, $E[\hat{F}_j(x)] = P[X_j \leq x] = F_j(x)$, so $E[\Delta_j(x)] = 0$. Then,

$$E[\Delta_{j-1}(x)\Delta_{j-1}(y)] = \text{Cov}[\Delta_{j-1}(x), \Delta_{j-1}(y)] \\ \leq \sup_{z \in \mathbb{R}} \text{Var}[\Delta_{j-1}(z)] \leq \Gamma_{j-1}n^{-3/2}$$

for all states x, y . Therefore, using (26) for the first equality and assuming that we can interchange the expectation and integral in the third equality,

$$E[(\tilde{F}_j(z) - F_j(z))^2] \\ = E\left[\left(\int_{-\infty}^{\infty} (\hat{F}_{j-1}(x) - F_{j-1}(x)) dF_j(z|x)\right)^2\right] \\ = E\left[\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} (\Delta_{j-1}(x)\Delta_{j-1}(y)) dF_j(z|x) dF_j(z|y)\right] \\ = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} E[\Delta_{j-1}(x)\Delta_{j-1}(y)] dF_j(z|x) dF_j(z|y) \\ \leq \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \Gamma_{j-1}n^{-3/2} dF_j(z|x) dF_j(z|y) \\ \leq \Lambda_j^2(z)\Gamma_{j-1}n^{-3/2}.$$

Combining this with (31) and (25), and observing that $\hat{F}_j(z) - \tilde{F}_j(z)$ has mean zero and is uncorrelated with $\tilde{F}_j(z) - F_j(z)$, we obtain that

$$E[(\hat{F}_j(z) - F_j(z))^2] = E[(\hat{F}_j(z) - \tilde{F}_j(z))^2] \\ + E[(\tilde{F}_j(z) - F_j(z))^2] \\ \leq \Lambda_j^2(z)\Gamma_{j-1}n^{-3/2} + (\Lambda_j + 1)n^{-3/2}/4 \\ \leq \Gamma_j n^{-3/2},$$

and this completes the induction. To prove the proposition's statement, we have

$$E\left[\left|\frac{1}{n} \sum_{i=0}^{n-1} c_j(X_{i,j}) - E[c_j(X_j)]\right|^2\right] \\ = E\left[\left|\int_{-\infty}^{\infty} c_j(z) d\hat{F}_j(z) - \int_{-\infty}^{\infty} c_j(z) dF_j(z)\right|^2\right] \\ = E\left[\left|\int_{-\infty}^{\infty} (\hat{F}_j(z) - F_j(z)) dc_j(z)\right|^2\right] \\ \leq \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} E(\Delta_j(x)\Delta_j(y)) |dc_j(x)| |dc_j(y)| \\ \leq \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \Gamma_j n^{-3/2} |dc_j(x)| |dc_j(y)| \\ \leq (V(c_j))^2 \Gamma_j n^{-3/2}. \quad \square$$

This result could be generalized in principle to higher-dimensional settings, although we quickly hit the curse of dimensionality when l or d increases. A potential extension

would need a counterpart of Lemma 2. To illustrate the idea, we sketch how this can be done for $d = 2$ and $l = 1$ (the state space is still one-dimensional, but we need two uniforms per step) for a stochastically increasing Markov chain, so that we have a two-dimensional increasing histogram in the unit cube. Partition the cube into n subcubes by partitioning each axis into $n^{1/3}$ equal parts (assuming that $n^{1/3}$ is an integer). The histogram boundary is now a surface. If we fix one of the two horizontal coordinates to a multiple of $n^{-1/3}$, this determines a vertical plane, and the intersection of this plane with the histogram boundary can be covered by at most $2n^{1/3}$ subcubes in a similar manner as in the proof of Lemma 2. We can do this for each multiple of $n^{-1/3}$, and repeat in each direction. We find that the histogram boundary can be covered by at most $Kn^{2/3}$ subcubes for some constant K . Press et al. (1992, p. 314) sketch a similar argument in a setting where the integral of an indicator function over the three-dimensional unit cube is approximated by averaging over the (deterministic) points of a Sobol' point set. However, their argument is highly heuristic; it would only stand if the locations of the points in the n subcubes were independent uniform random variables instead of being deterministic.

In general, for a $d + 1$ -dimensional cube and a stochastically increasing Markov chain, the histogram boundary can be covered by $Kn^{-d/(d+1)}$ subcubes for some constant K that may depend on d but not on n . This can be turned into a variance bound of $O(n^{2-d/(d+1)}) = O(n^{1+1/(d+1)})$. The result could also be generalized to the case where c_j has infinite variation (e.g., unbounded state spaces and cost functions) if we assume that large-variation areas have low probability.

Acknowledgments

The work of the first author was supported by NSERC-Canada grant ODGP0110050, NATEQ-Québec grant 02ER3218, and a Canada Research Chair. Part of this work was done while the first author was a visitor at IRISA in Rennes, France. The work of the third author was supported by EuroNGI Network of Excellence and the SurePath ACI sécurité project. A (short) preliminary version of this paper appeared as L'Ecuyer et al. (2006). Richard Simard helped in preparing the software for the numerical experiments, and Charles Sanvido implemented the stratification. The authors thank Area Editor Michael Fu, the associate editor, and two referees for their helpful comments.

References

Antonov, I. A., V. M. Saleev. 1979. An economic method of computing LP $_r$ -sequences. *Zh. Vychisl. Mat. i. Mat. Fiz.* **19** 243–245. In Russian.
 Asmussen, S. 1985. Conjugate processes and the simulation of ruin problems. *Stochastic Processes and Their Appl.* **20** 213–229.
 Ben-Ameur, H., P. L'Ecuyer, C. Lemieux. 2004. Combination of general antithetic transformations and control variables. *Math. Oper. Res.* **29**(4) 946–960.

- Bratley, P., B. L. Fox. 1988. Algorithm 659: Implementing Sobol's quasirandom sequence generator. *ACM Trans. Math. Software* **14**(1) 88–100.
- Cafisch, R. E., B. Moskowitz. 1995. Modified Monte Carlo methods using quasi-random sequences. H. Niederreiter, P. J.-S. Shiu, eds. *Monte Carlo and Quasi-Monte Carlo Methods in Scientific Computing. Lecture Notes in Statistics*, Vol. 106. Springer-Verlag, New York, 1–16.
- Garvels, M. J. J., D. P. Kroese, J.-K. C. W. Van Ommeren. 2002. On the importance function in splitting simulation. *Eur. Trans. Telecomm.* **13**(4) 363–371.
- Glasserman, P. 2004. *Monte Carlo Methods in Financial Engineering*. Springer-Verlag, New York.
- Glasserman, P., P. Heidelberger, P. Shahabuddin, T. Zajic. 1999. Multi-level splitting for estimating rare event probabilities. *Oper. Res.* **47**(4) 585–600.
- Hickernell, F. J. 1998. A generalized discrepancy and quadrature error bound. *Math. Comput.* **67** 299–322.
- Hickernell, F. J. 2002. Obtaining $o(n^{-2+\epsilon})$ convergence for lattice quadrature rules. K.-T. Fang, F. J. Hickernell, H. Niederreiter, eds. *Monte Carlo and Quasi-Monte Carlo Methods 2000*. Springer-Verlag, Berlin, 274–289.
- Hlawka, E. 1971. Discrepancy and Riemann integration. L. Mirsky, ed. *Studies in Pure Mathematics: Papers in Combinatorial Theory, Analysis Geometry, Algebra and the Theory of Numbers*. Academic Press, New York, 121–129.
- Hoschek, W. 2004. *The Colt Distribution: Open Source Libraries for High Performance Scientific and Technical Computing in Java*. CERN, Geneva, Switzerland. <http://dtd.lbl.gov/~hoschek/colt/>.
- Hull, J. 2000. *Options, Futures, and Other Derivative Securities*, 4th ed. Prentice-Hall, Englewood Cliffs, NJ.
- Kleinrock, L. 1975. *Queueing Systems*, Vol. 1. Wiley, New York.
- Kollig, T., A. Keller. 2002. Efficient bidirectional path-tracing by randomized quasi-Monte Carlo integration. K.-T. Fang, F. J. Hickernell, H. Niederreiter, eds. *Monte Carlo and Quasi-Monte Carlo Methods 2000*. Springer-Verlag, Berlin, 290–305.
- Lécot, C. 1996. Error bound for quasi-Monte Carlo integration with nets. *Math. Comput.* **65**(213) 179–187.
- Lécot, C., S. Ogawa. 2002. Quasirandom walk methods. K.-T. Fang, F. J. Hickernell, H. Niederreiter, eds. *Monte Carlo and Quasi-Monte Carlo Methods 2000*. Springer-Verlag, Berlin, 63–85.
- Lécot, C., B. Tuffin. 2004. Quasi-Monte Carlo methods for estimating transient measures of discrete time Markov chains. H. Niederreiter, ed. *Monte Carlo and Quasi-Monte Carlo Methods 2002*. Springer-Verlag, Berlin, 329–343.
- L'Ecuyer, P. 2004a. Quasi-Monte Carlo methods in finance. R. G. Ingalls, M. D. Rossetti, J. S. Smith, B. A. Peters, eds. *Proc. 2004 Winter Simulation Conf.* IEEE Press, Piscataway, NJ, 1645–1655.
- L'Ecuyer, P. 2004b. *SSJ: A Java Library for Stochastic Simulation*. Software user's guide. <http://www.iro.umontreal.ca/~lecuyer>.
- L'Ecuyer, P., C. Lemieux. 2000. Variance reduction via lattice rules. *Management Sci.* **46**(9) 1214–1235.
- L'Ecuyer, P., C. Lemieux. 2002. Recent advances in randomized quasi-Monte Carlo methods. M. Dror, P. L'Ecuyer, F. Szidarovszky, eds. *Modeling Uncertainty: An Examination of Stochastic Theory, Methods, and Applications*. Kluwer Academic Publishers, Boston, 419–474.
- L'Ecuyer, P., C. Lécot, B. Tuffin. 2006. Randomized quasi-Monte Carlo simulation of Markov chains with an ordered state space. H. Niederreiter, D. Talay, eds. *Monte Carlo and Quasi-Monte Carlo Methods 2004*. 331–342.
- Matoušek, J. 1999. *Geometric Discrepancy: An Illustrated Guide*. Springer-Verlag, Berlin.
- Niederreiter, H. 1992. *Random Number Generation and Quasi-Monte Carlo Methods, SIAM CBMS-NSF Regional Conf. Ser. Appl. Math.*, Vol. 63. SIAM, Philadelphia.
- Owen, A. B. 1998. Latin supercube sampling for very high-dimensional simulations. *ACM Trans. Model. Comput. Simulation* **8**(1) 71–102.
- Owen, A. B. 2003a. Quasi-Monte Carlo sampling. H. W. Jensen, ed. *Monte Carlo Ray Tracing: Siggraph 2003, Course 44*. <http://www-stat.stanford.edu/~owen/reports/siggraph03.pdf>. 69–88.
- Owen, A. B. 2003b. Variance with alternative scramblings of digital nets. *ACM Trans. Model. Comput. Simulation* **13**(4) 363–378.
- Owen, A. B. 2005. Multidimensional variation for quasi-Monte Carlo. J. Fan, G. Li, eds. *Internat. Conf. Statist. Honour of Professor Kai-Tai Fang's 65th Birthday*, Hong Kong. 49–74. <http://www-stat.stanford.edu/~owen/reports/>.
- Press, W. H., S. A. Teukolsky, W. T. Vetterling, B. P. Flannery. 1992. *Numerical Recipes in C: The Art of Scientific Computing*, 2nd ed. Cambridge University Press, New York.
- Sloan, I. H., S. Joe. 1994. *Lattice Methods for Multiple Integration*. Clarendon Press, Oxford, UK.
- Tezuka, S. 2002. Quasi-Monte Carlo—Discrepancy between theory and practice. K.-T. Fang, F. J. Hickernell, H. Niederreiter, eds. *Monte Carlo and Quasi-Monte Carlo Methods 2000*. Springer-Verlag, Berlin, 124–140.
- Tuffin, B. 1996. On the use of low discrepancy sequences in Monte Carlo methods. *Monte Carlo Methods Appl.* **2**(4) 295–320.
- Wilson, J. R. 1983. Antithetic sampling with multivariate inputs. *Amer. J. Math. Management Sci.* **3** 121–144.