Available online at www.sciencedirect.com

ScienceDirect

Mathematics and Computers in Simulation 81 (2010) 476–489

MATHEMATICS AND COMPUTERS IN SIMULATION

www.elsevier.com/locate/matcom

Coupling from the past with randomized quasi-Monte Carlo

P. L'Ecuyer*, C. Sanvido

DIRO, Université de Montréal, CP 6128, Succ. Centre-Ville, Montréal H3C 3J7, Canada

Received 21 August 2007; received in revised form 17 April 2009; accepted 7 September 2009
Available online 26 September 2009

Abstract

The coupling-from-the-past (CFTP) algorithm of Propp and Wilson permits one to sample exactly from the stationary distribution of an ergodic Markov chain. By using it n times independently, we obtain an independent sample from that distribution. A more representative sample can be obtained by creating negative dependence between these n replicates; other authors have already proposed to do this via antithetic variates, Latin hypercube sampling, and randomized quasi-Monte Carlo (RQMC). We study a new, often more effective, way of combining CFTP with RQMC, based on the array-RQMC algorithm. We provide numerical illustrations for Markov chains with both finite and continuous state spaces, and compare with the RQMC combinations proposed earlier.

Keywords: Variance reduction; Randomized quasi-Monte Carlo; Markov chain; Perfect sampling; Coupling from the past

1. Introduction

The Monte Carlo (MC) method is a key tool for estimating mathematical expectations of the form

$$\mu = \int_{\mathcal{S}} c(x)\,\mathrm{d}\pi(x), \tag{1}$$

where π is a probability measure defined over some measurable space $(\mathcal{S}, \mathcal{F})$, and $c : \mathcal{S} \to \mathbb{R}$ is a measurable cost (or reward) function. If we know how to sample exactly and efficiently from the measure π, then we can sample n independent realizations of a random variable X from this measure, compute $c(X)$ for each of them, and take the average. This is the Monte Carlo method.

However, there are many practical situations where we have no direct way of sampling exactly from π. An important type of setting where this occurs naturally is when we want to estimate the steady-state average cost per step for a system whose evolution is modeled by an ergodic Markov chain $\{X_j,\ j \geq 0\}$, often with large state space $\mathcal{S}$ and complicated dynamics; see [10,12] and [1, Chapter IV], for example. Typically, in this case, we have little a priori clue of how π might look like but we can easily simulate the Markov chain. For example, the Markov chain might represent the day-to-day evolution of an inventory system whose state on day j is X_j, $c(X_j)$ is the cost for that day given X_j (we assume that X_j contains enough information for the model to be Markovian and for the cost to be

* Corresponding author. Tel.: +1 514 343 2143.
E-mail address: lecuyer@iro.umontreal.ca (P. L'Ecuyer).

0378-4754/$36.00 © 2009 IMACS. Published by Elsevier B.V. All rights reserved.
doi:10.1016/j.matcom.2009.09.003