Proceedings of the 2005 Winter Simulation Conference
M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines, eds.

IMPORTANCE SAMPLING IN MARKOVIAN SETTINGS

Werner Sandmann

Department of Information Systems and Applied Computer Science
Communication Services, Telecommunication Systems, and Computer Networks
University of Bamberg
D–96045 Bamberg, GERMANY

ABSTRACT

Rare event simulation for stochastic models of complex systems is still a great challenge even for Markovian models. We review results in importance sampling for Markov chains, provide new viewpoints and insights, and we pose some future research directions.

1 INTRODUCTION

Model-based performance evaluation is an important tool in a variety of application areas. In particular, when some uncertainty in system behavior is involved, stochastic models are adequate and have shown to be useful. Dynamical systems evolve over time, and very often stochastic processes are used to model and analyze them. Although a lot of analytical and numerical techniques are known for the analysis of stochastic processes, due to the increasing complexity of today's real–world systems, these techniques are often not applicable. They may be computationally too costly or space complexity may be intractable. Even in the relatively simple case of Markov chains, state space explosion causes serious problems. In such cases simulation can provide an alternative.

A further difficulty is introduced when rare events are involved. In many systems rare events must not be neglected. Typical examples are data packet losses in communication networks, bit errors in digital communications, or ruins in finance and insurance risk. Such rare but important events substantially determine system performance and reliability. As a matter of fact, direct rare event simulation is impossible, since rare events occur with such an extremely small probability, that they occur too infrequently in simulations to provide reliable estimates in reasonable time. Hence, to estimate rare events, efficient simulation modeling and analysis is required. Simulation speed-up in the sense of reducing the computer time necessary to obtain estimates with a pre-defined accuracy can be achieved by reducing the variance of the estimator. One of the most promising variance reduction methods for efficient rare event simulation is importance sampling.

Importance sampling is already known for a long time, originally developed in the early 1940s in the context of Monte Carlo integration arising from problems in nuclear physics, see (Hammersley and Handscomb 1964) for a classical treatment, and (Halton 1970) for a comprehensive survey. Since around two decades it has been recognized that importance sampling is also a potentially powerful technique for simulating systems governed by stochastic processes, including Markov processes, generalized semi–Markov processes (GSMP), queueing or reliability models. Importance sampling is particularly useful for such models in the presence of rare events, see for example Glynn and Iglehart (1989) and Heidelberger (1995).

The basic idea of importance sampling is to transform the original model into a modified model, which should behave better with regard to the simulation goals. The probability measure underlying the model under consideration, for instance the density of random variables or the law driving a stochastic process, is changed in such a way that the variance of the resulting estimators is significantly reduced. The systematically biased estimator is then weighted by a correcting factor, the so-called likelihood ratio, which is simply the ratio of the original probability density or measure and the changed importance sampling probability density or measure. This yields an unbiased estimator for the property of interest, and it is known that there always exists an optimal zero variance estimator, which unfortunately depends on the unknown quantity to be estimated. In the particular context of rare event simulation importance sampling is often called a rare event provoking technique, since it then aims in generating more of the rare events of interest to obtain accurate confidence intervals in reasonable time.

The crucial point in a successful application of importance sampling is to find a change of measure providing accurate estimates, which means estimators with much smaller variance than the direct simulation estimators. When

importance sampling is applied properly, the variance of the resulting estimator can be many orders of magnitude less than that of the direct simulation estimator, but on the other hand, importance sampling can also lead to variance increase, even to infinite variance, if the chosen estimator is too far from optimal. Therefore, success and efficiency heavily depend on the change of measure determining the estimator, and consequently, the main difficulty and the art of importance sampling is to find such changes of measure that considerably decrease the estimator's variance. Thereby, it is by no means obvious whether or not a specific change of measure is a good one. The main part of the importance sampling literature is concerned with the change of measure and the development of estimators with certain accuracy and efficiency criteria.

In applications of importance sampling, such as in the already mentioned literature and also in more recent publications, as for example in Juneja (2004) or Sandmann (2004), the focus is most often on estimators providing bounded relative error or asymptotic optimality. Moreover, most often importance sampling methods are developed for higher-level model descriptions, usually for Markovian queueing networks. There are significantly less direct applications to Markov chains than to queueing models. Sometimes, in direct applications to Markov chains a specific structure is assumed, such as small transition probabilities (Juneja and Shahabuddin 2000, 2001). Many available techniques appear to be very efficient even if they are not close to the optimum. But, of course, exploiting knowledge of the optimal estimator to derive a change of measure should be a good approach. Investigating the optimal zero variance importance sampling estimator for Markov chains has been neglected.

The remainder of the paper is organized as follows. In section 2, we review importance sampling for discrete-time Markov chains. Sections 3 and 4 deal with the form of the optimal importance sampling distribution for finite horizon and steady state probabilities, respectively. Section 5 recalls a result for absorbing Markov chains, which is then used in section 6 to demonstrate that zero variance estimation of probabilities related to sequences of Bernoulli trials can be performed by simulating a discrete-time Markov chain. Finally, section 7 concludes the paper and discusses some directions for further research.

2 IMPORTANCE SAMPLING FOR DISCRETE-TIME MARKOV CHAINS

We review the mathematical basis of importance sampling for discrete-time Markov chains. Thereby we start and comprehensively describe the basics for discrete-time Markov chains, and then we present analogous formulae for continuous-time ones. Additional details and generalizations can be found in (Glynn and Iglehart 1989). For a comprehensive survey on importance sampling for queueing and reliability models see (Heidelberger 1995), for Markovian reliability models also see (Goyal et al 1992).

Let (X_n) be a homogeneous discrete–time Markov chain with state space $\mathcal{S}$, transition probability matrix $\mathbf{P} = (p_{ij})$ and initial distribution μ. Denote by $p(x_0, \ldots, x_K)$ the path probability determined by $\mathbf{P}$ and μ, that is $p(x_0, \ldots, x_K) = \mu(x_0) p_{x_0,x_1} \cdots p_{x_{K-1},x_K}$, where K is in general a random stopping time, which particularly includes the deterministic case of a constant path length. Let $g : \mathcal{S}^{K+1} \to \mathbf{R}$ be a real-valued function, and define $\gamma := \mathbf{E_P}[g(X_0, \ldots, X_K)]$, where $\mathbf{E_P}$ denotes the expectation taken with respect to $\mathbf{P}$ (and μ) on the space of sample paths of (X_n). Then for states $x_0, \ldots, x_K \in \mathcal{S}$

$$\gamma = \sum_{x_0,\ldots,x_K} g(x_0, \ldots, x_K) p(x_0, \ldots, x_K).$$

Note, that we distinguish random variables from their realizations by using upper case and lower case letters, respectively. The goal is to estimate γ via simulation. This would be usually done by the sample mean of $g(X_0, \ldots, X_K)$, but if γ is related to a rare event such a direct simulation is impossible. Think for example of g as the indicator function of an infrequently visited state or set of states. Now, let $p^* : \mathcal{S}^{K+1} \to [0, 1]$ denote a probability distribution on $\mathcal{S}^{K+1}$, which means $p^*(x_0, \ldots, x_K) = P\{X_0 = x_0, \ldots, X_K = x_K\}$. Assume, the implication

$$\begin{aligned} &g(x_0, \ldots, x_K) \cdot p(x_0, \ldots, x_K) > 0 \\ &\quad \Rightarrow p^*(x_0, \ldots, x_K) > 0 \end{aligned} \tag{1}$$

holds for all $(x_0, \ldots, x_K) \in \mathcal{S}^{K+1}$, and consider the likelihood ratio given by $L(X_0, \ldots, X_K) = p(X_0, \ldots, X_K)/p^*(X_0, \ldots, X_K)$. Then

$$\begin{aligned} \gamma &= \mathbf{E_P}[g(X_0, \ldots, X_K)] \\ &= \sum_{x_0,\ldots,x_K} \left(g(x_0, \ldots, x_K) \frac{p(x_0, \ldots, x_K)}{p^*(x_0, \ldots, x_K)} \right. \\ &\qquad \left. \times\ p^*(x_0, \ldots, x_K) \right) \\ &= \mathbf{E}_{p^*}\left[g(X_0, \ldots, X_K) L(X_0, \ldots, X_K)\right], \end{aligned}$$

where $\mathbf{E}_{p^*}$ denotes the expectation on paths with probability distribution p^*. This expectation can be estimated by repeated generations of $(X_0, \ldots, X_K)$, denoted by $(X_0^{(1)}, \ldots, X_K^{(1)}), \ldots, (X_0^{(N)}, \ldots, X_K^{(N)})$, according to p^*, building the sample mean and weighting by the likelihood

ratio. Thus, the importance sampling estimator

$$\hat{\gamma}_{IS} = \frac{1}{N}\sum_{i=1}^{N} g(X_0^{(i)}, \dots X_K^{(i)}) L(X_0^{(i)}, \dots X_K^{(i)}) \qquad (2)$$

is an unbiased estimator of γ. Note, that if K is a random stopping time, realizations of sample paths may have different lengths. We particularly emphasize that arbitrary probability distributions on $\mathcal{S}^{K+1}$ with (1) are allowed for p^*. There are no explicit conditions on single transition probabilities, but only on path probabilities. The importance sampling distributions may even be inhomogeneous or non–Markovian´, but the generation of paths then can become much more complicated than for homogeneous discrete-time Markov chains. Using a discrete-time Markov chain as importance sampling distribution follows as a special case of the general formulae. For inhomogeneous discrete-time Markov chains as importance sampling distributions we have with initial distribution ν and transition probability matrices $\mathbf{P}^*(0), \dots \mathbf{P}^*(K-1)$

$$p^*(x_0, \dots, x_K) = \nu(x_0) \prod_{i=0}^{K-1} p^*_{x_i, x_{i+1}}(i).$$

For the special case of homogeneous discrete–time Markov chains with initial distribution ν and transition probability matrix $\mathbf{P}^* = (p^*_{ij})$ as importance sampling distribution we get the same expression except for the dependence on i in the transition probabilities $p^*_{x_i, x_{i+1}}$. This corresponds to the so–called exponential change of measure for Markov chains (Bucklew 1990). The advantage of performing importance sampling using a fixed alternative Markov chain is that the effort for generating paths remains approximately the same as for the original Markov chain. Then the likelihood ratio is given by

$$L(x_0, \dots, x_K) = \frac{\mu(x_0)}{\nu(x_0)} \prod_{i=0}^{K-1} \frac{p_{x_i, x_{i+1}}}{p^*_{x_i, x_{i+1}}}$$

for the homogeneous case, and for the inhomogeneous case we have an additional dependence on i similar as in the corresponding formula for path probabilities above. This specific form renders successive computation by actualizing the likelihood ratio after each simulated state transition possible. A serious problem is caused when simulation runs become very long, when the time horizon becomes very large or even potentially infinite, as it is typically necessary for estimating steady–state performance measures. Indeed, it can be shown that the likelihood ratio almost surely vanishes when simulation run length tends to infinity (Glynn and Iglehart 1989). Thus, in the case of very long simulation runs, which are usual in simulating realistic network models, simply taking one fixed alternative Markovian transition probability matrix may not be satisfactory, and advanced techniques must be applied. In particular, regenerative simulation has proven useful, but although each Markov chain state is regenerative, the choice of suitable – in the sense that they are reached sufficiently often – regeneration points is by no means trivial.

2.1 Optimal Importance Sampling

The variance of the importance sampling estimator (2) is given by

$$\sigma^2(\gamma_{IS}) = \frac{\mathbf{E}_{p^*}[g(X_0, \dots, X_K)^2 L(X_0, \dots, X_K)^2] - \gamma^2}{N}.$$

From that, one can easily determine the uniquely defined optimal importance sampling distribution

$$p^*_{opt}(x_0, \dots, x_K) = \frac{g(x_0, \dots, x_K) p(x_0, \dots, x_K)}{\gamma},$$

which unfortunately explicitly depends on the unknown quantity to be estimated. Nevertheless, from its form given above valuable insights can be gained. Hence, although there are also a lot of applications of importance sampling for Markov chains, where the used importance sampling distribution is not very close to the optimal one, we should be interested in investigating the optimal importance sampling distribution more detailed. For this purpose, it is not only useful but necessary to have some guidelines about how to derive it without using a priori (in practice not given) knowledge on the exact solution.

2.2 Likelihood Ratio Conditions

The optimal importance sampling distribution depends on the quantity of interest. If the quantity is known, the optimal importance sampling distribution can be computed. Of course, this is not helpful in practice, since we need not estimate a known quantity. Therefore, one should aim at being able to derive optimal importance sampling distributions without this knowledge. It has proven useful (Sandmann 2005) to exploit a special property of the likelihood ratio. Note that under optimal importance sampling it is guaranteed that only such samples are possible, where $g(x_0, \dots, x_K) \neq 0$. In the following, for ease of notation $x := (x_0, \dots, x_K)$, and $x' = (x'_0, \dots, x'_K)$, respectively. For optimal importance sampling it can be readily seen that the likelihood ratio is given by

$$\forall x \in \{x : g(x) \neq 0\} : L(x) = \frac{p(x)}{p^*_{opt}(x)} = \frac{\gamma}{g(x)}.$$

Although this cannot be directly used as it again contains γ, we get the information, that for equal values of g the values of the likelihood ratio are equal, too. That means for all $x, x' \in \{x : g(x) \neq 0\}$:

$$g(x) = g(x') \Rightarrow L(x) = L(x') = \frac{\gamma}{g(x)}.$$

For the important case of estimating the probability of an event A putting the above equations together yields

$$\forall x, x' \in A : L(x) = L(x') = \frac{p(x)}{p^*_{opt}(x)} = \gamma.$$

We call these equations the *likelihood ratio conditions*. What is useful in applying them is of course not that the likelihood ratio depends on γ, but that likelihood ratios are equal. In particular note for estimating probabilities that the likelihood ratios are constant.

3 OPTIMAL IMPORTANCE SAMPLING FOR FINITE HORIZON PROBABILITIES

Consider the probability of a rare state $r \in \mathcal{S}$ over a finite horizon $K \in \mathbf{N}^+$, starting in an arbitrary state $x_0 \in \mathcal{S}$. This probability can be estimated via the relative frequency of visits in state r and is given by

$$\mathbf{E}\left[\frac{1}{K}\sum_{i=1}^{K} I_{\{X_i=r\}}\right].$$

That means as a special case of the general term the estimate for the probability under consideration is

$$g(X_0, \ldots, X_K) = \frac{1}{K}\sum_{i=1}^{K} I_{\{X_i=r\}}.$$

For the optimal importance sampling estimator it follows (Sandmann 2005)

$$p^*_{opt}(x_0, \ldots, x_K) = \frac{\left(\sum_{i=1}^{K} I_{\{x_i=r\}}\right) p(x_0, \ldots, x_K)}{\mathbf{E}\left[\sum_{i=1}^{K} I_{\{X_i=r\}}\right]}$$

for all paths $(x_0, \ldots, x_K) \in \mathcal{S}^{K+1}$, that include state r at least once, and $p^*_{opt}(x_0, \ldots, x_K) = 0$ for all paths that do not include state r. It is readily seen that the optimal estimator for the probability of interest equals the optimal estimator for the number of visits in state r. Thus, we can restrict to the estimation of this frequency of visits by

$$g(X_0, \ldots, X_K) = \sum_{i=1}^{K} I_{\{X_i=r\}}$$

$$\gamma^{(K)} = \mathbf{E}[g(X_0, \ldots, X_K)] = \sum_{i=1}^{K} \mathbf{E}\left[I_{\{X_i=r\}}\right].$$

Now, assume the optimal importance sampling estimator for the frequency of visits over horizon $K-1$ is known. In particular, the likelihood ratio condition must hold. In the following, we consider one further transition, thereby distinguishing if this leads to a visit of the rare state or not. Let the probability of a path of length K under optimal importance sampling be $p^{(K-1)}(x_0, \ldots, x_{K-1}) \cdot p^{(K)}_{x_{K-1}x_K}$, where $p^{(K-1)}(x_0, \ldots, x_{K-1})$ denotes the optimal importance sampling measure for horizon $K-1$ and $p^{(K)}_{x_{K-1}x_K}$ denotes the probability of a transition from a state x_{K-1} to state x_K in step K under optimal importance sampling for horizon K. Then it can be shown (Sandmann 2005) that for $x_k \neq r$

$$p^{(K)}_{x_{K-1}x_K} = \frac{\gamma^{(K-1)}}{\gamma^{(K)}} \cdot p_{x_{K-1}x_K}$$

and for $x_K = r$

$$p^{(K)}_{x_{K-1}x_K} = \frac{\gamma^{(K-1)}}{\gamma^{(K)}} \cdot \left(1 + \frac{1}{g(x_0, \ldots, x_{K-1})}\right) \times p_{x_{K-1}x_K}.$$

As a major problem in optimal importance sampling we discover the unknown factor $\gamma^{(K-1)}/\gamma^{(K)}$. Therefore, the practical application and implementation is difficult. In particular, the transition probabilities under optimal importance sampling in general depend on the number of visits to the rare state r already observed in the current simulation run. Thus, we can easily see that the optimal importance sampling measure for estimating finite horizon probabilities in Markov chains is generally not a Markov chain. In other words, for estimating $\gamma^{(K)}$ optimal importance sampling can generally not be realized by a Markov chain.

4 OPTIMAL IMPORTANCE SAMPLING FOR STEADY STATE PROBABILITIES

Now, consider optimal importance sampling for steady-state probabilities of rare states. For that, the first decision has to be made on the simulation method, that means which method should be applied to determine estimators. The method of independent replications of finite horizon was the natural method to apply in case of probabilities over

a finite horizon. It can be similarly applied to steady-state probabilities, where the horizon or in other words the simulation run length must be sufficiently long. Additionally a warm up period has to be taken into account to delete the initial transient period. Alternatively, the warm up period may be canceled for very long runs due to decreasing influence of initial conditions when simulation runs evolve. Regarding the optimal importance sampling measure there is no significant difference to performance measures defined on a finite horizon except for the fact that typically after the warm up period the initial state for the main simulation and data collection is different for all runs. Since under optimal importance sampling each run yields a perfect estimate of the property of interest, the form of the optimal importance sampling measure is independent of the particular initial state and corresponds to the finite horizon case. Similarly we can argue for the batch means method, where under optimal importance sampling each batch yields a perfect estimate.

Therefore, we consider here the regenerative method, where contradictory to the just mentioned methods, regenerative cycles (simulation runs of different lengths) occur. First of all the choice of a regeneration point is required. In Markov chains each state is regenerative and from the probabilistic point of view it does not matter which state is chosen. Only the simulation effort is influenced depending on the frequency of visits to the regeneration state or in other words on the mean length of regenerative cycles. The choice of the regeneration state does not influence the optimal importance sampling measure. Note that the change of measure also changes the probability of any regeneration state. In particular, a rare state can be chosen as regeneration state. Although this may seem quite strange at a first glance it has advantages. When chosing a rare state as the regeneration state the steady-state probability we are interested in is identical to the reciprocal of the mean recurrence time. We do choose the rare state of interest as regeneration state. Hence, we are interested in the expectation of the recurrence time to a rare state r chosen as regeneration state. For a path $(x_0,\ldots,x_t) \in \mathcal{S}^{t+1}$ where $x_0 = x_t = r$ and $x_i \neq r, 0 < i < t$ we have $g(x_0,\ldots,x_t) = t$ and

$$\begin{aligned} p^*_{opt}(x_0,\ldots,x_t) &= \frac{g(x_0,\ldots,x_t)\,p(x_0,\ldots,x_t)}{\gamma} \\ &= \frac{t p_{x_0x_1} p_{x_1x_2} \cdots p_{x_{t-1}x_t}}{\gamma}. \end{aligned}$$

Under optimal importance sampling each regenerative cycle yields a perfect estimate of the mean recurrence time. In particular, for two cycles of lengths t_1 and t_2, and with likelihood ratios L_1 and L_2, respectively, $t_1 \cdot L_1 = t_2 \cdot L_2 = \gamma$. Consider a cycle of length t_1 containing at least one state in which a self-loop is possible. Then there is a cycle of length $t_2 = t_1 + 1$ such that compared to the cycle of length t_1 it has been additionally occured one such self-loop. Without loss of generality assume the self-loop has occured in state x_i. Then the likelihood ratio of the second cycle is

$$L_2 = L_1 \cdot \frac{p_{x_ix_i}}{p^*_{x_ix_i}},$$

and it follows

$$\begin{aligned} t_1 &= (t_1+1) \cdot \frac{p_{x_ix_i}}{p^*_{x_ix_i}}, \\ p^*_{x_ix_i} &= \frac{t_1+1}{t_1} \cdot p_{x_ix_i} = \left(1 + \frac{1}{t_1}\right) \cdot p_{x_ix_i}. \end{aligned}$$

Note that even in this simple case the optimal importance sampling measure depends on the cycle length, which is neither known in advance nor during the simulation run. Hence, the recurrence time cannot be perfectly estimated by a Markov chain simulation.

Consider now an arbitrary regeneration state different from r. Then under optimal importance sampling the rare state r must be reached in each cycle. All regenerative cycles containing r must be possible under optimal importance sampling, and each regenerative cycle must yield a perfect estimate of the steady-state probability of r. Denote by $\frac{v}{t_1}$ and $\frac{v}{t_2}$ the estimates for two regenerative cycles of length t_1 and t_2, respectively, where both cycles contain the rare set equally often, say v times. As before, denote by L_1 and L_2 the cycles' likelihood ratios. Then under optimal importance sampling

$$\frac{v}{t_1} \cdot L_1 = \frac{v}{t_2} \cdot L_2, \quad t_1 \cdot L_2 = t_2 \cdot L_1.$$

Now, consider again two cycles of lengths t_1 and $t_2 = t_1+1$, which only differ in one self-loop, that is

$$L_2 = L_1 \cdot \frac{p_{x_ix_i}}{p^*_{x_ix_i}}.$$

Then it follows

$$\begin{aligned} t_1 + 1 &= \frac{p_{x_ix_i}}{p^*_{x_ix_i}} \cdot t_1, \\ p^*_{x_ix_i} &= \frac{t_1}{t_1+1} \cdot p_{x_ix_i}. \end{aligned}$$

Hence, for each regeneration state different from the rare state r the optimal importance sampling measure depends on the cycle length, too. Thus, the optimal importance sampling measure for estimating steady-state probabilities in Markov chains is generally not a Markov chain. In other words, for estimating steady-state probabilities in Markov chains optimal importance sampling can generally not be realized by a Markov chain.

5 ABSORBING MARKOV CHAINS

As we have seen, in general, optimal importance sampling for Markov chains cannot be performed by a Markov chain, even not by an inhomogeneous one. Although this is a negative result the form of the optimal importance sampling distribution gives insights and provide guidelines on how to choose a good importance sampling distribution (Sandmann 2005). Besides, we can note that optimal importance sampling leads out of the class of the original distribution. A generalized class has to be considered. The question arises if it is generally hopeless to simulate with a Markov chain as an importance sampling distribution. The answer is definitely no. In many cases of practical interest we are not concerned with totally general finite horizon or steady state probabilities of general Markov chains, but with some specific probabilities or with some specifically structured Markov chains. One of the most important special cases are absorbing Markov chains, where several specialized results on importance sampling exist.

In the following we will demonstrate 1) for simplified types of Markov chains optimal importance sampling often leads to a more generalized class of distributions, too; 2) this class may belong to the class of Markov chains. In particular, we will show that we can optimally simulate probabilities in sequences of independent Bernoulli trials (which can be viewed as Markov chains) by Markov chains, but the property of being independent Bernoulli trials is lost. When deriving the optimal importance sampling distribution for our Bernoulli example we will use the following result for absorbing Markov chains from (Kuruganti and Strickland 1997). For homogeneous discrete-time Markov chains with a set F of absorbing states and a non-absorbing initial state x_0 the optimal importance sampling distribution for estimating the probability $P\{T_F < T_{x_0}\}$ of reaching set F before returning to x_0 is given by a homogeneous discrete-time Markov chain. Here T_F and T_{x_0} denote hitting times of F and x_0, respectively. This is all we need in this paper, for a detailed treatment see (Kuruganti and Strickland 1997). In particular, it it not surprising that the transition probabilities under optimal importance sampling explicitly depend on the just mentioned unknown probability.

6 SEQUENCES OF INDEPENDENT BERNOULLI TRIALS

Consider n independent Bernoulli trials with success probability $p \in (0, 1)$. We are interested in estimating the probability of at least $k < n$ successes. For small p or large k this is clearly a rare event probability. A simulation typically performs independent simulation runs, each consisting of n independent Bernoulli trials, in other words of generating n random numbers. Since most of the runs do not generate k or more successes, a huge number of runs is needed to achieve confidence intervals of reasonable accuracy.

Following the basic idea of importance sampling, we would choose $p^* > p$ to provoke more successes and therefore a larger number of simulation runs including the rare event. Giving a special version of this example with $p = 0.5$, Juneja (2004) states that the optimal importance sampling distribution is typically not implementable, since for each $p \in (0, 1)$ the probability of observing less than k successes is positive. He does not further investigate this example but turns to large deviations theory to determine efficient, but not optimal, importance sampling distributions. From the above it seems to be clear that the optimal change of measure leads out of the class of independent Bernoulli trials. Nevertheless, the existence of the optimal importance sampling distribution is guaranteed. In what follows we show how one could find this optimal distribution. We do not compute it using the known exact solution of the probability we consider here, but we utilize the likelihood ratio condition. Thereby, we additionally illustrate the way ideas come and are rejected until finally the optimal importance sampling distribution is found. This should give valuable insights not only concerning the specific optimal importance sampling distributions, but particularly in the way one can successively improve models to get the optimal importance sampling distribution. First, we ask, whether there exists a $p^* \neq p$ such that the likelihood ratio condition holds. The answer is

Lemma 1 *There exists no $p^* \in (0, 1),\ p^* \neq p$ such that the likelihood ratio condition holds.*

Proof Let $x_1, \ldots, x_n,\ x_i \in \{0, 1\},\ i = 1, \ldots, n$ be a realization of one sequence of independent Bernoulli trials, building a *Bernoulli path*. The probability of each path of length n is given by

$$p(x_1, \ldots, x_n) = p^m (1 - p)^{n-m},\ m \leq n,$$

where m denotes the number of successes. For paths including the rare event we have $m \geq k$. Importance sampling with success probability p^* yields the likelihood ratio

$$L(x_1, \ldots, x_n) = \frac{p^m (1 - p)^{n-m}}{p^{*m} (1 - p^*)^{n-m}},$$

which should be the same for all $m \geq k$. Now, choose $m = k$ and $m = k + 1$, respectively. Then the likelihood ratio condition implies

$$\frac{p^k (1 - p)^{n-k}}{p^{*k} (1 - p^*)^{n-k}} = \frac{p^{k+1} (1 - p)^{n-k-1}}{p^{*k+1} (1 - p^*)^{n-k-1}},$$

from which we get $p^* = p$. As the next step on our way to the optimal importance sampling distribution, we consider a

modification, where the success probability may vary in each trial. This can be interpreted as time dependence. Again, we ask for the likelihood condition to hold. Unfortunately,

Lemma 2 *There exists no time dependent and state independent modified Bernoulli experiment such that the likelihood ratio condition holds.*

Proof It suffices to consider only the special case $n = 2, k = 1$. Here, paths including the event of interest are $(0,1), (1,0)$ and $(1,1)$ where the path probabilities are $p(0,1) = p(1,0) = p(1-p),\ \ p(1,1) = p^2$.

Let $p^{(1)}, p^{(2)}$ denote the success probabilities in the first and in the second trial, respectively, for the modified model. Then the corresponding likelihood ratios are

$$\begin{aligned} L(0,1) &= \frac{p(1-p)}{(1-p^{(1)})p^{(2)}}, \\ L(1,0) &= \frac{p(1-p)}{p^{(1)}(1-p^{(2)})}, \\ L(1,1) &= \frac{p^2}{p^{(1)}p^{(2)}}. \end{aligned}$$

For these must be identical, it follows

$$\frac{p(1-p)}{p^{(1)}(1-p^{(2)})} = \frac{p^2}{p^{(1)}p^{(2)}}$$

and

$$\frac{p(1-p)}{(1-p^{(1)})p^{(2)}} = \frac{p(1-p)}{p^{(1)}(1-p^{(2)})},$$

which yields $p^{(2)} = p$ and $p^{(1)} = p$, respectively.

Boths variants we have tried so far do not bring us to the optimal importance sampling distribution, so we have to experiment further. After trying independent identically distributed Bernoulli trials and independent time dependent Bernoulli trials we include dependence on the state, that is on the number of already observed successes. Let $p^{(i)}(m)$ denote the success probability for the i–th trial under the condition that there were already m successes. Consider again the case $n = 2, k = 1$. In the first trial the state is always 0, since there could not be successes before starting. The change of measure under consideration allows to guarantee, that only such paths are generated, which include the event of interest. We exclude (0,0) by defining the success probability for the second trial to be equal to one, if there was no success in the first trial. Hence, $p^{(2)}(0) := 1$.

For the other paths we have

$$\begin{aligned} L(0,1) &= \frac{(1-p)p}{(1-p^{(1)}(0))p^{(2)}(0)} = \frac{(1-p)p}{1-p^{(1)}(0)}, \\ L(1,0) &= \frac{p(1-p)}{p^{(1)}(0)(1-p^{(2)}(1))}, \\ L(1,1) &= \frac{p^2}{p^{(1)}(0)p^{(2)}(1)}. \end{aligned}$$

To compute the optimal importance sampling distribution, it remains to determine $p^{(1)}(0)$ and $p^{(2)}(1)$ such that the likelihood ratio condition holds. We get

$$\frac{p(1-p)}{p^{(1)}(0)(1-p^{(2)}(1))} = \frac{p^2}{p^{(1)}(0)p^{(2)}(1)}$$

and

$$\frac{(1-p)p}{1-p^{(1)}(0)} = \frac{p^2}{p^{(1)}(0)p^{(2)}(1)},$$

from which we obtain

$$p^{(2)}(1) = p, \quad p^{(1)}(0) = \frac{1}{2-p}.$$

We have seen, that in the optimal change of measure for estimating the probability of at least a given number k of successes, the optimal success probabilities depend both on the current trial and on the number of the already observed successes, even for this very special case of at least one success in two trials. The optimal change of measure leads out of the class of independent Bernoulli trials.

Now, the question arises if in general the optimal change of measure leads to such state dependent and time dependent success probabilities, or if it even depends on the complete history of the sequence of trials. This can be checked for small values of n by performing analogous computations as we have done above and at for instance for $n = 3$ and $n = 4$ it turns out that the optimal change of measure leads to success probabilities for each trial which only depend on the number of already observed successes, not on the complete history. In the following, we examine this more systematically for the general case. For this purpose, we consider the number of success probabilities to be determined and the number of equations resulting from the likelihood ratio condition.

We start with the latter. On principle, we have to consider all Bernoulli paths of length n with at least n successes, but since the simulation can be stopped after the k–th success, it suffices to consider only such path segments of length at most n where the k–th success occurs in the last trial. The minimum length of such paths is k, and there is exactly one path of length k with k successes. For paths of length $k+1$ it is fixed that the k–th success must occur

in trial $k+1$, and the remaining $k-1$ successes can occur in the first k trials. There exist $\binom{k}{k-1}$ such paths. Similarly, it can be argued for paths of lengths $k+2, k+3, \ldots$ until finally for paths of length n where the k–th success must occur in the n–th trial, and $k-1$ successes occur in $n-1$ trials before. Thus there are $\binom{n-1}{k-1}$ paths of length n with k–th success in trial n, and there are altogether

$$|\text{Paths}| = \sum_{m=k-1}^{n-1} \binom{m}{k-1} = \binom{n}{k}$$

paths to consider. Hence, the likelihood ratio condition yields $\binom{n}{k} - 1$ equations. This is the same as the number of paths of length n with exactly k successes, because paths with more than k successes are up to the k–th success identical to paths with exactly k successes.

Let us consider now, how many unknown probabilities have to be determined. In the i–th trial these are the probabilities $p^{(i)}(0), \ldots, p^{(i)}(i-1)$. These are i probabilities, and therefore, for n trials $1+2+\cdots n = n(n+1)/2$ probabilities. Since the simulation can be stopped after k successes, probabilities for states representing at least k successes are redundant. They have not to be determined. Again, this is possible at the earliest after trial k. Then the probability $p^{(k+1)}(k)$ is redundant, and also the probabilities $p^{(k+2)}(k)$, $p^{(k+2)}(k+1)$ and so on until after trial $n-1$ the probabilities $p^{(n)}(k), \ldots, p^{(n)}(n-1)$ are redundant. Altogether there are $1+2+\cdots+(n-k) = (n-k)(n-k+1)/2$ redundant probabilities.

Moreover, in optimal importance sampling the rare event – at least k successes – must be observed in each experiment. From this, it follows, that some probabilities must be one, namely if and only if from the current state k successes are only possible, if all remaining trials are successes. This is the case for the last but k trial in state 0, in the last but $k-1$ trial in state 1, and so on. Moreover, for the last but $k-1$ trial, state 0 is infeasible under optimal importance sampling, and therefore its success probability becomes redundant, too. Similar arguments hold until the last trial, for which in state $k-1$ the success probability must be one, and states $0, \ldots, k-2$ become infeasible under optimal importance sampling. Altogether we have $1+2+\cdots+k = k(k+1)/2$ probabilities to be one or, because of their infeasibility, redundant. The remaining number of unknown probabilities to determine is $k(n-k)$. If we compare this with the number of $\binom{n}{k} - 1$ equations, we state, that it is not yet guaranteed, that these are solvable. We will show, that they are indeed solvable. Hence, the optimal importance sampling distribution is given by Bernoulli trials, that depend on time and state but not on the complete history of all previously performed trials, and this yields a perfect estimator for our probability of interest, the probability of at least k successes in n independent Bernoulli trials with identical success probabilities. We will show this by modelling sequences of independent Bernoulli trials by discrete-time Markov chains with absorbing states and applying results given by Kuruganti and Strickland (1997).

6.1 Modeling as Absorbing Markov Chain

It may be surprising that we utilize results for recurrence probabilities in absorbing Markov chains to investigate optimal importance sampling for Bernoulli trials. However, as we have seen, the optimal importance sampling distribution for estimating the probability of more than k successes in n Bernoulli trials must at least depend on both the number of the already performed trials and the already observed successes. In other words, it depends on both time and state. In particular, it leads out of the class of independent Bernoulli trials. Our notion of time and state suggests an interpretation as stochastic process. As we want to show that the dependence on states is only on the current state and not on the complete history of successive states, it seems reasonable to interpret our experiment as a Markov chain. A sequence of n independent Bernoulli trials can be modelled as a homogeneous discrete-time Markov chain, where a state is a pair consisting of the number of successes and the number of trials. Hence, the state space is a subset of $\{0, \ldots, n\}^2$, and the initial state is $(0,0)$. Since it is not possible to have more successes than trials, for each state (i,m) it follows $i \leq m$, where the first component i denotes the number of successes and the second component m denotes the number of already performed trials. Here, we are interested in at least k successes in n trials, not in the exact number of successes. Therefore, we can stop the sequence of trials after k successes. Thus the state space is

$$\mathcal{S} := \{(i,m) \in \{0, \ldots, k\} \times \{0, \ldots, n\} : i \leq m\}.$$

The set of states in which k successes have been occured can be modelled as an absorbing class defined by $F := \{(i,m) \in \mathcal{S} : i = k\}$. There are states from which k successes are not possible. That are such states, where the number of remaining trials is less than the number of necessary successes. In other words, states, where the number of unsuccessful trials is greater than the maximum possible number of unsuccessful trials. Formally, that means states in

$$\begin{aligned} A &:= \{(i,m) \in \mathcal{S} : n - m < k - i\} \\ &= \{(i,m) \in \mathcal{S} : m - i > n - k\}, \end{aligned}$$

which is exactly the set of states that must not be reached under optimal importance sampling.

In (Kuruganti and Strickland 1997) the probability of returning to the initial state before reaching an absorbing

state is considered. In our actual model a return to the initial state is obviously impossible. Therefore, we need the modification that the initial state is reached iff k successes in n trials are not any longer possible. In our actual model this is the case exactly when a transition to the set A occurs. Hence, instead of entering set A we make a transition to the initial state. We do not need the states in A and we can eliminate them from the state space. A natural interpretation of our resulting Markov chain is a restart after the failure of the overall experiment is sure. For simulation practice and the optimal importance sampling estimator this does not play any role since the simulation is stopped after reaching the absorbing set F or the initial state. In particular, under optimal importance sampling the overall experiment does not fail and no return to the initial state or the formerly used set A, respectively, is possible. The probability of at least k successes in a sequence of n independent Bernoulli trials is exactly the probability of reaching the absorbing set F before returning to the initial state $(0, 0)$ in our just described Markov chain.

Altogether we have a homogeneous discrete-time Markov chain (X_n) with state space

$$\mathcal{S} := \{(i, m) \in \{0, \ldots, k\} \times \{0, \ldots, n\} : i \leq m\} \setminus A,$$

initial distribution $\pi^{(0)}$, where

$$\pi^{(0)}_{(0,0)} = 1, \qquad \forall s \in \mathcal{S} \setminus \{(0, 0)\} : \pi^{(0)}_s = 0$$

transition probability matrix $\mathbf{P}$ with transition probabilities

$$\begin{array}{ll} p_{(i,m),(i,m)} = 1, & \text{if } (i, m) \in F, \\ p_{(i,m),(i+1,m+1)} = p, & \text{if } (i, m) \in \mathcal{S} \setminus F, \\ p_{(i,m),(i,m+1)} = 1 - p, & \text{if } (i, m) \in \mathcal{S} \setminus F \wedge (i, m + 1) \in \mathcal{S}, \\ p_{(i,m),(0,0)} = 1 - p, & \text{if } (i, m) \in \mathcal{S} \setminus F \wedge (i, m + 1) \notin \mathcal{S}. \end{array}$$

The probability of interest is $P\{T_F < T_{(0,0)}\}$. Hence, the probability of interest is exactly of the form investigated by Kuruganti and Strickland (1997). That implies that the optimal change of measure leads to a homogeneous discrete-time Markov chain. To retransfom our Markov chain into Bernoulli trials we first note that the second component of a Markov chain state represents the number of already performed Bernoulli trials, which is the time in the Bernoulli trials. As the state after some Bernoulli trials we have the number of successes. In particular, it is not necessary to consider the complete history. With the formerly used notations we get the state and time dependent success probabilities

$$p^{(m)}(i) = p^*_{(i,m-1)(i+1,m)}.$$

Since the results of Kuruganti and Strickland (1997) imply that the transition probability matrix is uniquely defined, it follows that our system of equations derived in 6 is uniquely solvable. Thus we have shown

Theorem 1 *The optimal importance sampling estimator for the probability of at least k successes in n independent Bernoulli trials is given by Bernoulli trials with state dependent and time dependent success probabilities, where a state represents the number of already observed successes and time is represented by the number of already performed trials.*

It is remarkable that the form of such a time dependent model has been derived by means of homogeneous, time independent Markov chains. For that it was essential that the investigated experiment was of finite horizon. This rendered possible to represent the time in the investigated model as state components in a Markov chain thereby keeping the state space finite. The second essential point was to abstract from the real meaning of the initial state in the Bernoulli model to apply the results from (Kuruganti and Strickland 1997).

Finally, we will show that the probability of exactly k successes in n independent Bernoulli trials can be perfectly estimated by simulating Bernoulli trials with state and time dependent success probabilities, too. Only a slight modification in our above Markov chain is necessary. Let $F = \{(k, n)\}$ be the set of states, whose probabilities are to be estimated. Here, this is only one single state. Similarly as before we get a homogeneous discrete-time Markov chain with state space

$$\mathcal{S} = \{(i, m) \in \{0, \ldots, k\} \times \{0, \ldots, n\} : i \leq m\} \setminus A,$$

initial distribution $\pi^{(0)}$, where

$$\pi^{(0)}_{(0,0)} = 1, \qquad \forall s \in \mathcal{S} \setminus \{(0, 0)\} : \pi^{(0)}_s = 0$$

and transition probability matrix $\mathbf{P}$ with transition probabilities

$$\begin{array}{ll} p_{(i,m),(i,m)} = 1, & \text{if } (i, m) \in F, \\ p_{(i,m),(i+1,m+1)} = p, & \text{if } (i, m) \in \mathcal{S} \setminus F \wedge i < k, \\ p_{(i,m),(0,0)} = p, & \text{if } (i, m) \in \mathcal{S} \setminus F \wedge i = k, \\ p_{(i,m),(i,m+1)} = 1 - p, & \text{if } (i, m) \in \mathcal{S} \setminus F \wedge (i, m + 1) \in \mathcal{S}, \\ p_{(i,m),(0,0)} = 1 - p, & \text{if } (i, m) \in \mathcal{S} \setminus F \wedge (i, m + 1) \notin \mathcal{S}. \end{array}$$

As before a model for Bernoulli trials with state and time dependent success probabilities can be constructed. Thus we have shown

Theorem 2 *The optimal importance sampling estimator for the probability of exactly k successes in n independent Bernoulli trials is given by Bernoulli trials with state dependent and time dependent success probabilities, where a state represents the number of already observed successes and time is represented by the number of already performed trials.*

7 CONCLUSIONS

We have described how importance sampling works for Markov chains, thereby focusing on the discrete-time case. The form of optimal importance sampling for general finite horizon and steady state probabilities has been given, from which can be seen that in general optimal importance sampling for Markov chains is not possible with alternative Markov chains as importance sampling distributions. To illustrate that a Markov chain can be useful as importance sampling distribution we have exploited a result on absorbing Markov chains and applied it to sequences of Bernoulli trials resulting in optimal importance sampling estimators given by Markov chains. Several conclusions can be drawn. When simulating general Markovian models, more general (than Markovian) importance sampling distributions should be considered. The specific forms of the optimal estimator suggest that approaches seem to be promising that do not only optimize from one to another simulation run, but also within each single run. Thus, combinations of so-called dynamic and adaptive methods should be useful. For sufficiently simple Markovian models alternative Markov chains as importance sampling distributions are also useful, but even in such cases a generalization of the original model is necessary.

Finally, we want to draw attention to stiff Markov chains, where not only the probability of one single state or a specific set of target states is of interest, but the complete distribution. In such a case, we are concerned with the problem, that importance sampling "needs" some target states to be applicable, whereas often in stiff Markov chains all or at least very many state probabilities are of interest. On the other hand it seems to be reasonable that importance sampling may be useful to circumvent the difficulty in simulating (and numerically solving) stiff Markov chains. It should be a goal of future research to apply importance sampling also to such types of problems.

REFERENCES

Bucklew, J. A. 1990. *Large deviation techniques in decision, simulation, and estimation*. New York: Wiley & Sons.

Glynn, P. W., and D. L. Iglehart. 1989. Importance sampling for stochastic simulations. *Management Science* 35 (11): 1367–1392.

Goyal, A., Shahabuddin, P., Heidelberger, P., Nicola, V. F. and P. W. Glynn. 1992. A unified framework for simulating Markovian models of highly reliable systems. *IEEE Transactions on Computers*, 41:36–51.

Hammersley, J. M. and D. C. Handscomb. 1964. *Monte Carlo Methods*. London: Methuen.

Halton, J. H. 1970. A retrospective and prospective survey of the Monte Carlo method. SIAM Review, 12:1–63.

Heidelberger, P. 1995. Fast simulation of rare events in queueing and reliability models. *ACM Transactions on Modeling and Computer Simulation* 5(1): 43–85.

Juneja, S. 2004. Efficient Rare Event Simulation Using Importance Sampling: An Introduction. *Computational Mathematics Modelling and Algorithms*, pp. 357–396, New Delhi: Narosha Publishers.

Juneja, S., and P. Shahabuddin. 2000. A splitting based importance sampling algorithm for the fast simulation of Markov chains with small transition probabilities. *IEEE Transactions on Reliability*, 50(3):235–245.

Juneja, S., and P. Shahabuddin, P. 2001. Efficient simulation of Markov chains with small transition probabilities. *Management Science*, 47(4): 547–562.

Kuruganti, I. and S. Strickland. 1997. Optimal importance sampling for Markovian systems with applications to tandem queues. *Mathematics and Computers in Simulation* 44:61–79.

Sandmann, W. 2004. Relative Error and Asymptotic Optimality in Estimating Rare Event Probabilities by Importance Sampling. *Proceedings of the 2004 OR Society Simulation Workshop in Cooperation with the ACM SIGSIM, Birmingham, UK, March 23–24 2004, 49–57*: The Operational Research Society.

Sandmann, W. 2005. On optimal importance sampling for discrete-time Markov chains. In *Proceedings of the 2nd International Conference on the Quantitative Evaluation of Systems (QEST) 2005*. Torino, Italy: IEEE Computer Society Press.

AUTHOR BIOGRAPHY

WERNER SANDMANN studied Computer Science and Mathematics at the University of Bonn, Germany, where he received his diploma degree (Dipl.–Inform., 1998) and his Ph.D. (Dr. rer. nat., 2004). Since 2004 he is an assistant professor of Computer Science in the department of Information Systems and Applied Computer Science at the University of Bamberg, Germany. His e-mail address is `<werner.sandmann@wiai.uni-bamberg.de>`, and his web page is `<www.uni-bamberg.de/wiai/ktr/html/mitarbeiter/sandmann/>`.