Optimal Importance Sampling in Markov Process Simulation

Esa Hyytiä and Jorma Virtamo

Helsinki University of Technology
Networking Laboratory
P.O.Box 3000, FIN-02150 HUT, Finland
E-mail: {esa.hyytia, jorma.virtamo}@hut.fi

ABSTRACT
In this paper we present an adaptive algorithm to estimate the transient blocking probability of a communication system, described by a Markov process, during a finite time interval starting from a given state. The method uses importance sampling for variance reduction and adjusts the parameters of the twisted distribution based on earlier samples. The method can be effectively applied to a decision making problem where future revenues are estimated with extensive simulations, in order to find an improved policy by so called first policy iteration.

KEY WORDS
Importance Sampling, Markov Process Simulation

1. Introduction

The theory of Markov decision processes (MDP) is an important tool which can be applied to many problems emerging in different areas of business and technology. For example, in a modern all-optical network using WDM, a well-known problem is to assign a route and wavelength to each arriving connection so that there is no wavelength conflicts and, at the same time, trying to minimize the long term blocking probability. This is commonly referred to as the routing and wavelength assignment problem (RWA) [1, 2]. In recent papers [3, 4] the authors have presented a robust idea to improve any given RWA algorithm by so called first policy iteration. The method relies on having good estimates of the future costs during time $(0, T)$ given that the system starts from state i. This paper is a continuation of that work in a more general framework. In this paper we present a method to effectively estimate by simulations the incurred costs of a Markov process starting from certain state i during a finite time interval.

Importance sampling is a well-known method to reduce the variance of the estimators [5]. Briefly the idea is to take samples using another pdf instead of the original one to make interesting events more probable (and uninteresting events less likely at the same time). In particular, events having no contribution to the estimated quantity should be made to occur with zero probability. In this paper we present how the importance sampling parameters can be adjusted based on the earlier outcomes.

The rest of the paper is organized as follows. First, in Section 2 we review the idea behind the importance sampling. Then in Section 3 the variance of IS estimator is studied and a relation between the original random variable and variance of the IS estimator is presented. In Section 4 the IS is formulated for the case of process simulation with Poisson arrival process. In Section 5 some simulation results are shown for the Erlang loss system and, finally Section 6 contains conclusions.

2. Importance Sampling

In importance sampling (IS) one tries to reduce the variance of the estimator by taking samples of interesting quantity using another pdf instead of the original. IS is especially effective in the rare event simulation [6, 7].

In this section a brief introduction to IS is given. For further reference about IS see e.g. [5]. In Section 3 the treatment of IS is extended to case where the statistics of the studied system are unknown and the optimal IS parameters are estimated as the simulations proceed.

Later in this paper IS is formulated for a finite time process simulation where customers arrive to the system according to a Poisson process and incurred costs are to be estimated.

Definition 1 (Twisted distribution)
Let X be a random variable with pdf $p(x)$. Let $p^*(x)$ be another pdf for which it holds that,

$$p^*(x) = 0 \qquad \Rightarrow \qquad p(x) = 0. \tag{1}$$

Then $p^*(x)$ defines a twisted distribution of $p(x)$, and corresponding random variable X^* is a twisted random variable of X.

For the rest of the paper term twisted random variable is used quite freely. Basically by a twisted random variable we mean another random variable X^* having a pdf $p^*(x)$ such that requirement (1) is satisfied. The requirement guarantees that possible events for X are also possible for X^*.

Proposition 1
Let X be an arbitrary random variable and $h(x)$ a mapping to $\mathbb{R}$. Denote the expectation of $h(X)$ by $\theta = \mathrm{E}[h(X)]$. Then, for any twisted random variable X^* it holds that,

$$\theta = \mathrm{E}[h(X)] = \mathrm{E}[q(X^*)h(X^*)], \tag{2}$$

where the ratio $q(x) = p(x)/p^*(x)$ is called the likelihood ratio.

Proof:

$$\begin{aligned} \mathrm{E}\left[h(X)\right] &= \int_{\mathbb{R}} p(x)h(x)\,dx \\ &= \int_{\mathbb{R}} p^*(x)\left(\frac{p(x)}{p^*(x)}h(x)\right)\,dx \\ &= \mathrm{E}\left[\frac{p(X^*)}{p^*(X^*)}h(X^*)\right]. \end{aligned}$$

□

Definition 2 (Importance Sampling estimator)
Assume that m independent samples X_i^*, $i = 1,\ldots,m$, have been drawn, each from a separate twisted distribution having the pdf $p_i^*(x)$. Then the obvious estimator for the expectation θ is,

$$\begin{aligned} \hat{\theta} &= \frac{1}{m}\sum_{i=1}^{m}\frac{p(X_i^*)}{p_i^*(X_i^*)}\cdot h(X_i^*) \qquad (3) \\ &= \frac{1}{m}\sum_{i=1}^{m} q_i(X_i^*)\cdot h(X_i^*) \overset{def}{=} \frac{1}{m}\sum_{i=1}^{m} Y_i^*, \end{aligned}$$

The random variable Y_i^* we refer to as the observed random variable. For the estimator $\hat{\theta}$ it holds that

$$\begin{aligned} \mathrm{E}\left[\hat{\theta}\right] &= \frac{1}{m}\sum_i \mathrm{E}\left[Y_i^*\right] = \theta, \\ \mathrm{V}\left[\hat{\theta}\right] &= \mathrm{V}\left[\frac{1}{m}\left(Y_1^* + \ldots + Y_m^*\right)\right] = \frac{1}{m^2}\sum_i \mathrm{V}\left[Y_i^*\right]. \end{aligned}$$

If the twisted distributions are identical, i.e. $p_i^*(x) = p^*(x)$, then $Y_i^* \sim Y^*$ and,

$$\mathrm{V}\left[\hat{\theta}\right] = \frac{\mathrm{V}\left[Y^*\right]}{m} = \frac{\mathrm{E}\left[{Y^*}^2\right]}{m} - \frac{\theta^2}{m}.$$

The basic requirement for an estimator is that it is unbiased, i.e. the expectations of the estimator and the estimated quantity are equal[1]. But also, as stated before, the variance of the estimator should be as small as possible (variance reduction). The smaller the variance, the less samples are needed to obtain an estimate with a required confidence interval. Hence, as a criterion for the goodness, the variance of the estimator should be studied. Thus, we are led to the problem,

$$\text{minimize } \mathrm{V}\left[Y^*\right] \qquad \Leftrightarrow \qquad \text{minimize } \mathrm{E}\left[{Y^*}^2\right].$$

3. Variance in Importance Sampling

In this section a formula to express the variance of the observed random variable in terms of the original random variable is presented. This relation between the two random variables will be later used to estimate the variance of the IS estimator based on samples drawn from the original distribution. In particular, it is possible to use the previous samples to find a twisting that minimizes the variance of the estimator.

[1] However, sometimes even biased estimators are acceptable.

Proposition 2
Let X and X^* be random variables, where the latter is obtained by twisting the original random variable X in the standard way. Then the variance of the observed random variable $Y^* = q(X^*)\cdot h(X^*)$ can be expressed in terms of the original random variable X, and is,

$$\mathrm{V}\left[Y^*\right] = \mathrm{E}\left[q(X)h^2(X)\right] - \mathrm{E}\left[h(X)\right]^2. \qquad (4)$$

Proof:

$$\begin{aligned} \mathrm{V}\left[Y^*\right] &= \mathrm{E}\left[(Y^*)^2\right] - \mathrm{E}\left[Y^*\right]^2 \\ &= \mathrm{E}\left[q^2(X^*)h^2(X^*)\right] - \mathrm{E}\left[h(X)\right]^2 \\ &= \mathrm{E}\left[q(X)h^2(X)\right] - \mathrm{E}\left[h(X)\right]^2. \end{aligned}$$

□

The variance of Y^* is the quantity to be minimized. The smaller the variance is, the less samples are needed to obtain a good estimate.

It is straightforward to obtain an unbiased estimator for the variance of the observed random variable. Here, however, we do not need it since we can as well minimize the "second moment" in Eq. (4) as the latter term does not depend on the twisting. For the first term we have an obvious estimate,

$$\hat{r} = \frac{1}{m}\sum_i \frac{p(X_i^*)}{p^*(X_i^*)}h^2(X_i^*). \qquad (5)$$

Above it is assumed that all the samples X_i^* are drawn from the same distribution. This, however, does not have to be the case. To this end, Proposition 2 can be rewritten in a more general form.

Proposition 3
Let Y_i, $i = 1,\ldots,m$, be m observed random variables, $Y_i = \frac{p(X_i^*)}{p_i^*(X_i^*)}h(X_i^*)$, where each X_i^* is a twisted random variable with a separate twisted pdf $p_i^*(x)$. Let X^* be another twisted random variable with pdf $p^*(x)$, and $Y^* = \frac{p(X^*)}{p^*(X^*)}h(X^*)$ respectively. Then,

$$\mathrm{V}\left[Y^*\right] = \mathrm{E}\left[\frac{p^2(X_i^*)}{p_i^*(X_i^*)\cdot p^*(X_i^*)}\cdot h^2(X_i^*)\right] - \mathrm{E}\left[h(X)\right]^2.$$

Proof:

$$\begin{aligned} \mathrm{V}\left[Y^*\right] &= \mathrm{V}\left[\frac{p(X^*)}{p^*(X^*)}\cdot h(X^*)\right] \\ &= \mathrm{V}\left[\frac{p_i^*(X^*)}{p^*(X^*)}\cdot\left(\frac{p(X^*)}{p_i^*(X^*)}h(X^*)\right)\right], \end{aligned}$$

which, using Proposition 2, gives,

$$\begin{aligned}\mathrm{V}[Y^*] &= \mathrm{E}\left[\frac{p_i^*(X_i^*)}{p^*(X_i^*)}\cdot\left(\frac{p(X_i^*)}{p_i^*(X_i^*)}h(X_i^*)\right)^2\right] \\ &\quad - \mathrm{E}\left[\frac{p(X_i^*)}{p_i^*(X_i^*)}h(X_i^*)\right]^2 \\ &= \mathrm{E}\left[\frac{p^2(X_i^*)}{p_i^*(X_i^*)\cdot p^*(X_i^*)}\cdot h^2(X_i^*)\right] - \mathrm{E}\left[h(X)\right]^2.\end{aligned}$$

□

Proposition 3 can be used to estimate the variance of Y^* based on the previous samples in the same way as was done in Eq. (5). Then our task reduces to finding the twisted pdf $p^*(x)$ such that the estimator,

$$\begin{aligned}\hat{r} &= \frac{1}{m}\sum_i \frac{p^2(X_i^*)}{p_i^*(X_i^*)\cdot p^*(X_i^*)}\cdot h^2(X_i^*) \qquad (6)\\ &= \frac{1}{m}\sum_i q(X_i^*)q_i(X_i^*)h^2(X_i^*),\end{aligned}$$

is minimized.

4. Application: Costs with Poisson Arrivals

In this section we define some random variables for a stochastic system S with Poissonian arrivals (see e.g. [8–10]). Then some well-known results for Poisson processes are presented and finally applied to formulate the IS for S.

Example 1 (Stochastic System with Poisson Arrivals)
Suppose a system S with Poissonian arrivals, e.g. connections offered to a data communication network. At time 0 the system is in some initial state and the aim is to estimate the future costs during a certain time interval $(0, T)$. Let,

$$\begin{cases} X &= \text{system path in } (0,T)\text{, defined by arrivals and departures,} \\ C &= \text{incurred costs, } C = c(X), \\ N &= \text{number of arrivals, } N = n(X). \\ \tilde{p}(k) &= \mathrm{P}\{n(X) = k\} \end{cases}$$

Note that X is a random variable in the path space.

In a simulation, samples of X are drawn and the expectation $\mathrm{E}[C]$ is estimated on the basis of them. The obvious estimator for the mean costs is,

$$\hat{C} = \frac{1}{m}\sum_i C_i. \qquad (7)$$

Next, we briefly review some well-known results of Poisson processes which will be needed when formulating IS for the stochastic system S defined above.

Proposition 4
Given the number of arrivals, n, from a Poisson process in a time interval $(0, T)$, these n arrivals are uniformly distributed over this time interval.

The proof is simple and given in many textbooks. Proposition 4 is very useful when characterizing the IS with Poissonian arrivals, namely we have the following corollary.

1 Corollary
When sampling a system S with system paths X^* resulting from a twisted Poisson arrival process with the arrival intensity λ^*, instead of the original intensity λ, the likelihood ratio $q(x)$ needed in Eq. (2) depends only on the number of arrivals $n = n(x)$, i.e.,

$$q(x) = \frac{\tilde{p}(n)}{\tilde{p}^*(n)} \stackrel{\text{def}}{=} \tilde{q}(n),$$

where $\tilde{p}^*(k) = \mathrm{P}\{n(X^*) = k\}$.

Example 2 (IS with Poisson Arrivals)
We continue with Example 1. For mean costs $\mathrm{E}[C]$ it holds that,

$$\mathrm{E}[C] = \mathrm{E}[\mathrm{E}[C|N]] = \mathrm{E}[\tilde{c}(N)],$$

where $\tilde{c}(n) = \mathrm{E}[C|N = n]$. Let X^* be the sample path resulting from the twisted Poissonian arrival process with arrival intensity λ^*, $N^* = n(X^*)$ be the number of arrivals and $C^* = c(X^*)$ the incurred costs. The importance sampling becomes,

$$\begin{aligned}\mathrm{E}[C] &= \sum_n \tilde{p}(n)\cdot\tilde{c}(n) = \sum_n \tilde{p}^*(n)\left[\frac{\tilde{p}(n)}{\tilde{p}^*(n)}\cdot\tilde{c}(n)\right] \\ &= \mathrm{E}\left[\frac{\tilde{p}(N^*)}{\tilde{p}^*(N^*)}\cdot\tilde{c}(N^*)\right].\end{aligned}$$

and when taking m twisted samples of X we get the IS estimator,

$$\hat{C} = \frac{1}{m}\sum_i \frac{\tilde{p}(N_i^*)}{\tilde{p}^*(N_i^*)}C_i^* = \frac{1}{m}\sum_i G_i^*, \qquad (8)$$

where $G_i^* = \frac{\tilde{p}(N_i^*)}{\tilde{p}^*(N_i^*)}C_i^*$ is the real valued observed random variable.

Note that N_i^* and C_i^* are not independent but functions of the random variable X_i^* in the path space, defining the actual realisation, i.e. arrivals and departures from the system.

There is no reason why the path samples X_i^* should be generated using the same twisted arrival rate λ^*. Assume that the i^{th} sample (or replication) X_i^* is obtained by using a twisted Poisson arrival process with arrival intensity λ_i^*. Then, the IS estimator for the mean cost $\hat{C}$ generalizes to,

$$\hat{C} = \frac{1}{m}\sum_i \frac{\tilde{p}(N_i^*)}{\tilde{p}_i^*(N_i^*)}C_i^*,$$

which differs from (8) only in that now the denominator $\tilde{p}_i^*(n)$ is also a function of i. The question is what kind of distribution should be used to obtain the next sample $i+1$?

In dynamic programming, the twisting can be adjusted based on the earlier samples, i.e. in this case after each replication the next arrival rate λ_i^* can be chosen based on earlier samples. This will be studied in the next section.

4.1 Variance of Observed Random Variable

In this section we try to minimize the variance of the observed random variable. Propositions 2 and 3 presented in Section 3 will be applied to the case where the twisting concerns only the Poisson arrival process. The variance of the IS estimator can be approximated based on the previous samples and the optimal arrival intensity λ^* can be estimated by minimizing the estimated variance .

Example 3 (Estimating the variance of an IS estimator) Recall the stochastic system defined in Example 1. As before, let G^* be the observed random variable. Using Proposition (2) the variance of G^* can be expressed in terms of the original random variable C,

$$\mathrm{V}\left[G^*\right] = \mathrm{E}\left[\frac{\tilde{p}(N)}{\tilde{p}^*(N)} C^2\right] - \mathrm{E}\left[C\right]^2.$$

The latter term is constant and a change in the arrival process only affects the first term. Hence, minimizing the variance is equivalent to minimizing the first term, for which we have the obvious unbiased estimate (c.f. 5),

$$\hat{r} = \frac{1}{m} \sum_i \frac{\tilde{p}(N_i)}{\tilde{p}^*(N_i)} C_i^2. \tag{9}$$

The next step is to choose p^* such that (9) is minimized.

Example 4 (Constant increase in Arrival Rates) We continue with the previous example. Let $p^* = p^*(\alpha)$, where $\alpha > 0$ is a constant multiplier for the arrival intensity, i.e. $\lambda^* = \alpha\lambda$. Then the number arrivals $N^* \sim \text{Poisson}(\alpha\lambda T)$, and (9) becomes,

$$\hat{r}(\alpha) = \frac{e^{(\alpha-1)\lambda T}}{m} \sum_i \frac{1}{\alpha^{N_i}} C_i^2.$$

After enough samples have been obtained the above equation can be minimized with respect to the parameter α. Thus, omitting the constant divisor m, the function f to be minimized is,

$$f(\alpha) = e^{(\alpha-1)\lambda T} \sum_i C_i^2 / \alpha^{N_i},$$

where clearly the factor $e^{(\alpha-1)\lambda T}$ is a strictly increasing function of α and the sum strictly decreasing function of α. Also both parts are always positive. Taking the first two derivatives one gets,

$$f'(\alpha) = e^{(\alpha-1)\lambda T} \sum_i C_i^2 \alpha^{-N_i-1} (\lambda T\alpha - N_i),$$

$$f''(\alpha) = e^{(\alpha-1)\lambda T} \sum_i C_i^2 \alpha^{-N_i-2} \left[(\alpha\lambda T - N_i)^2 + N_i\right].$$

The second derivative is always positive and hence the function is convex and the minimum is reached exactly at one point. At the minimum the first derivate is equal to zero, which happens iff,

$$g(\alpha) \stackrel{\text{def}}{=} \sum_i C_i^2 (\lambda T \alpha^{-N_i} - N_i \alpha^{-N_i-1}) = 0,$$

The root can be found, e.g., by using the Newton-Raphson method,

$$\alpha_{k+1} = \alpha_k - \frac{g(\alpha_k)}{g'(\alpha_k)},$$

where

$$g'(\alpha) \stackrel{\text{def}}{=} \sum_i C_i^2 (-\lambda T N_i \alpha^{-N_i-1} + N_i(N_i+1)\alpha^{-N_i-2}).$$

The Newton-Raphson method is known to converge at least quadratically for a simple root if y'' is continuous [11, p. 366]. Let $M_i = M - N_i$ where $M = \max_i N_i$. Then the method gets the form,

$$\alpha_{k+1} = \alpha_k - \frac{\alpha_k \sum_i C_i^2 \alpha_k^{M_i} (\lambda T \alpha_k - N_i)}{\sum_i C_i^2 \alpha_k^{M_i} (-\lambda T N_i \alpha_k + N_i(N_i+1))}.$$

In the previous example the optimal twisting parameter α was estimated based on the samples obtained by using the original arrival intensity λ. However, it is also possible to adjust the arrival rate after each replication of the system path.

Example 5 (Adaptive IS with Poissonian Arrivals) Return again to example where a stochastic system with Poissonian Arrivals is simulated for a finite time period and the twisting concerns the arrival intensity. Using (6) the estimate for the quantity r to be minimized (which also minimizes the variance of the next sample $G^* = \tilde{q}(N^*) \cdot C^*$) becomes,

$$\hat{r}(\alpha) = \frac{1}{m} \sum_i \tilde{q}(N_i^*) \cdot \tilde{q}_i(N_i^*) \cdot {C_i^*}^2.$$

Thus, by storing the triple

$$\{n_i, s_i, g_i\} \stackrel{\text{def}}{=} \{N_i^*,\ \tilde{p}_i(N_i^*) \cdot (G_i^*)^2,\ G_i^*\},$$

for each sample gives,

$$\boxed{\begin{aligned} \hat{\theta} &= \frac{1}{m} \sum_i g_i, \\ \hat{r} &= \frac{1}{m} \sum_i \frac{s_i}{p^*(n_i)}. \end{aligned}}$$

However, the adaptive approach can make it difficult to estimate the variance of the estimator as the samples are no longer independent (the distribution of each sample depends on the previous samples). An easy solution to this is to freeze the adaptive twisting for a certain number of samples and estimate the variance based on them.

5. Simulation Results: Erlang's Loss System

The Erlang loss system is a simple stochastic system often used to model a link to which calls are offered. The user population is assumed to be infinite, i.e. the arrival process is modelled as a Poisson process with some parameter λ (calls / time unit). Furthermore, the call holding times are assumed to obey exponential distribution with mean $1/\mu$. The capacity of the link is finite, C, and calls arriving when there is no free capacity are blocked and lost.

As an example we consider a system where $\lambda = 10$, $\mu = 1$ and $C = 16$ and initially the system starts from state 10, i.e. there are 10 customers in service at time $t = 0$. Our problem is to estimate the average number of blocked customers during the interval $(0, 1)$? For the reference, in steady state this system has a blocking probability of 2.23%, and on average 10 customers arrive during the given time interval.

In Fig. 1 the variance estimate, with an added common constant $\hat{C}^2$, is depicted as a function of α based on 4000 simulation replications using the original arrival distribution. The minimum is obtained with $\alpha \approx 1.4$ and the variance reduction in this case is around $1 : 3$.

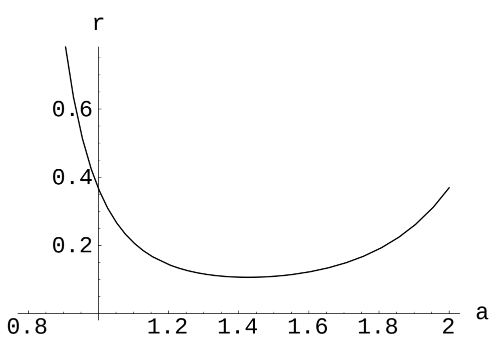

Figure 1. Estimated estimator variance as a function of $\boldsymbol{\alpha}$ after $\mathbf{4000}$ samples using the original arrival distribution ($\boldsymbol{\lambda} = \mathbf{10}$).

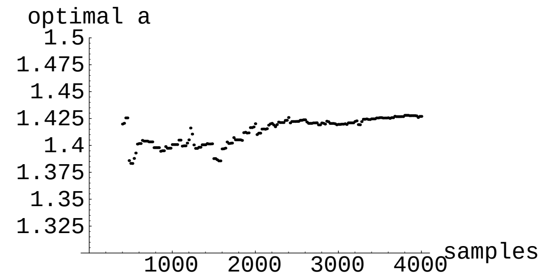

Figure 2. Development of the estimated optimal $\boldsymbol{\alpha}$ as a function of the number of the samples. For $\boldsymbol{N > 400}$ the suggested $\boldsymbol{\alpha}$ stays close to $\mathbf{1.4}$, which is the optimal twisting in the light of these simulations.

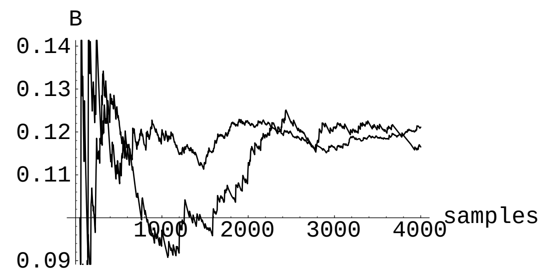

Figure 3. Convergence of the normal and IS($\boldsymbol{\alpha} = \mathbf{1.4}$) estimators as a function of the number of replications. The system is initially in state $\boldsymbol{X(0) = 10}$.

Fig. 2 shows how the optimal α, estimated on the basis of the generated samples, evolves as the number of samples grows. It can be noted that after around 400 samples the behaviour seems to be quite stable, actually all dots past that point lie close to 1.4. This suggests that a reasonably good α can be obtained after a few thousand replications, or even less.

Fig. 1 suggests that the variance of the IS estimator is about 3 times smaller than the original estimator. Hence, the IS estimator should also converge 3 times quicker. In Fig. 3 the convergence of both the original and IS estimators are depicted as a function of the number of replications in one series of simulations. The IS estimator converges after around 500 samples to the neighbourhood of 0.12, while from the original estimator it takes about 2000 samples to converge to that level.

In the previous example the system started from a state with a high occupancy. It can be expected that the performance improvement is even greater when the interesting event is rare, e.g. if the system occupation is initially smaller or the arrival intensity λ is lower. To demonstrate this another experiment was made where the initial occupancy was chosen to be 5. Based on 4000 sample paths the optimal α was estimated to be $\alpha \approx 1.6$. Among those 4000 samples only about 20 caused blocking events. In Fig. 4 the convergence of the original and IS estimator is again depicted. The IS estimator gives a reasonably good estimate after a few hundred samples, while the original estimator seems to be far from reliable region. The expected number of blocking events is 10 times less than in the previous case $(0.12 \rightarrow 0.011)$.

5.1 Application: Call Admission Control

In this Section we present an example of a possible application of the adaptive IS technique presented earlier. Assume that calls arriving to an Erlang loss system originate from two independent processes with arrival intensities λ_1 and λ_2, respectively. The higher priority class 1 yields expected revenue r_1 and the lower priority class 2 yields ex-

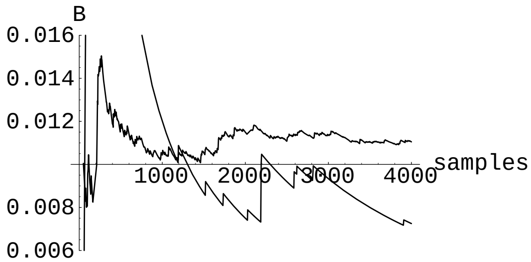

Figure 4. Convergence of the normal and IS($\alpha = 1.6$) estimators as a function of the number of replications. The system is initially in the state $X(0) = 5$.

pected revenue r_2 per call. It is clear that the optimal call admission control (CAC) policy rejects a class 2 call if the expected future costs, with the call admitted, exceed the expected future costs, with the call rejected, by the amount r_2 or more.

In the MDP theory the difference in the infinite time horizon costs between a system starting from a given state instead of equilibrium is called the relative value of the state [12, 13]. If the difference in relative values is greater than r_2 the class 2 call should be rejected. Instead of trying to solve the relative values exactly, which usually is not feasible due to the prohibitive size of the state space, they can be approximated with process simulations on the fly [3, 4]. The adaptive IS approach, presented earlier, can be applied to reduce the time needed to run these simulations.

One possible implementation is as follows. For each decision, n, the state costs are estimated by a number of simulation replications using a fixed twisting parameter α_n. After the simulations and the CAC decision based on these have been made, we estimate what would have been the optimal twisting parameter α_n^*. This estimate is used as the twisting parameter for the next decision, i.e., $\alpha_{n+1} = \alpha_n^*$. This kind of adaptive updating of the twisting parameter can be expected to work well when the system changes relatively little between the decision epochs.

6. Conclusions

In this paper we presented a method to evaluate the optimal twisting for IS, based on samples taken from the original pdf. This can be easily applied to process simulation where the system is studied for a finite time interval. The adaptive IS method is especially interesting when applied to the first policy iteration where the expected future costs are estimated by making repeated simulations of the system. Often the interesting quantity, e.g. expected blocking probability, is a rare event and IS can improve the performance of the algorithm considerably. In the light of preliminary simulation study of the Erlang loss system the approach looks promising. The real testbed, however, will be an application to more complex systems such as the RWA problem in a WDM network.

References

[1] R. Ramaswami and K. N. Sivarajan, *Optical Networks, A Practical Perspective*. Morgan Kaufmann Series in Networking, Morgan Kaufmann Publishers, 1998.

[2] T. E. Stern and K. Bala, *Multiwavelength Optical Networks: a Layered Approach*. Addison Wesley, 1999.

[3] E. Hyytiä and J. Virtamo, Dynamic Routing and Wavelength Assignment Using First Policy Iteration, *ISCC'2000, the Fifth IEEE Symposium on Computers and Communications, Antibes, Juan les Pins, France*, pp. 146–151, IEEE, July 2000.

[4] E. Hyytiä and J. Virtamo, Dynamic Routing and Wavelength Assignment Using First Policy Iteration, Inhomogeneous Traffic Case, *P&QNet2000, the International Conference on Performance and QoS of Next Generation Networking, Nagoya, Japan*, pp. 301–316, Nov. 2000.

[5] S. M. Ross, *Introduction to Probability Models*. Academic Press, 7th ed., 2000.

[6] P. Heidelberger, Fast simulation of rare events in queueing and reliability models, *ACM Transactions on Modeling and Computer Simulation*, vol. 5, pp. 43–85, Jan. 1995.

[7] P. E. Heegaard, A survey of Speedup simulation techniques. Workshop tutorial on Rare Event Simulation, Aachen, Germany, Aug. 1997.

[8] R. B. Cooper, *Introduction to Queueing Theory*. Elsevier Science Publishing, 1981.

[9] S. Karlin and H. M. Taylor, *A Second Course in Stochastic Processes*. Academic Press, 1981.

[10] Z. Brzeźniak and T. Zastawniak, *Basic Stochastic Processes*. Springer Undergraduate Mathematics Series, Springer-Verlag, 2000.

[11] L. Råde and B. Westergren, *Mathematics Handbook for Science and Engineering*. Studentlitteratur, third ed., 1995.

[12] H. C. Tijms, *Stochastic Models, An Algorithmic Approach*. John Wiley & Sons Ltd, 1994.

[13] Z. Dziong, *ATM Network Resource Management.* McGraw-Hill, 1997.