Rare event simulation

Adaptive state-dependent importance sampling simulation of Markovian queueing networks

PIETER-TJERK DE BOER
Department of Computer Science, University of Twente, P.O. Box 217, 7500 AE Enschede, The Netherlands.
ptdeboer@cs.utwente.nl

VICTOR F. NICOLA
Department of Electrical Engineering, University of Twente, P.O. Box 217, 7500 AE Enschede, The Netherlands.
nicola@cs.utwente.nl

Abstract. In this paper, a method is presented for the efficient estimation of rare-event (buffer overflow) probabilities in queueing networks using importance sampling. Unlike previously proposed change of measures, the one used here is not static, i.e., it depends on the buffer contents at each of the network nodes. The 'optimal' state-dependent change of measure is determined adaptively during the simulation, using the cross-entropy method. The adaptive state-dependent importance sampling algorithm proposed in this paper yields asymptotically efficient simulation of models for which it is shown (formally or otherwise) that no effective static change of measure exists. Simulation results for queueing models of communication systems are presented to demonstrate the effectiveness of the method.

1 INTRODUCTION

During the last decade, there has been much interest in the estimation of rare-event probabilities in queues and networks of queues, with applications to models of telecommunication networks as well as computer and manufacturing systems. Two methods have gained popularity: importance sampling ([1], [2]), and importance splitting (or RESTART) [3], the former of which is used in this paper.

One simple network that received a lot of attention is a set of two or more queues in tandem. Despite its simplicity, a complete analysis of this system is hard due to the behaviour at the state-space boundaries. As a consequence, no importance sampling change of measure is known that is provably asymptotically efficient for estimating e.g. the overflow probability of the total population in such a system.

In [4], an importance sampling procedure was described for estimating the overflow probability of the total population in tandem queues. A simple and static (i.e., state-independent) change of measure was used: exchange the arrival rate with the service rate (of the bottleneck queue, in case of a tandem system). In [5], the asymptotic efficiency of that method for a single queue was proved. In [6] this heuristic was extended to overflows of the total population in any Jackson network. However, it was shown in [7] that for two or more queues in tandem, this heuristic does not always give an asymptotically efficient simulation, depending on the values of arrival and service rates. It is reasonable to expect that similar problems will occur with this method in other Jackson networks.

Clearly, by allowing the change of measure to depend on the state of the system (i.e., the content of each of the queues), more efficient importance sampling schemes may be obtained. This approach was recently used in [8], where the overflow probability of the second queue in a two-node tandem Jackson network is estimated using a simulation in which the change of measure depends on the content of the first buffer; the functional dependence of the rates on the buffer content is derived from a Markov additive process representation of the system. Furthermore, in [9] a state-dependent change of measure is used for simulating link overloads in a telecommunications network; again, the functional dependence of the importance sampling rates on the system state is derived using a heuristic. And in [10], an efficient exponential change of measure is considered for a certain class of Markov chain problems, where the exponential tilting parameter is a function of the state and is determined using large-deviations theory. The biggest obstacle to the use of a state-dependent change of measure in general is the problem of determining this dependence: rather specific mathematical models are used in the publications mentioned, making the results very specific to those

problems.

As an alternative to avoid the complex mathematical analysis often used to determine a good (state-independent) change of measure, several adaptive methods have been proposed recently; see [11], [12], [13], [14], [15], [16], [17]. All of these either try to iteratively minimize the variance of the estimator involved, or a related quantity like the cross-entropy. However, none of these papers consider a state-dependent change of measure for simulation of queueing models.

In this paper, we present an adaptive method for determining a state-dependent change of measure for rare events in queueing problems. This is a rather versatile method:

- due to the adaptiveness, a complex mathematical analysis of the problem is not necessary.
- since the state-dependent change of measure is less restrictive, problems can be solved for which no effective state-independent change of measure exists.

In particular, with this method the probability of overflow of the total network population in Jackson tandem networks can be asymptotically efficiently estimated, even in those cases for which it is shown in [7] that the heuristic of exchanging the arrival rate with the bottleneck service rate does not work. In addition, the combination of state-dependence and adaptiveness leads to another useful property: the standard deviation of the estimator can decrease faster than proportional to the square-root of the total simulation effort.

Here we restrict our discussion to the estimation of rare-event probabilities in discrete-time Markov chains (DTMCs). However, efforts are underway to expand the method to more general models.

The rest of this paper is organized as follows. Section 2 provides a summary of the most important aspects of the adaptive method from [15]. Section 3 explains the implementation of this algorithm for state-dependent simulation, discusses some of the problems involved and their solutions, and provides a mathematical analysis of the variance. In section 4, empirical results demonstrate the effectiveness of the method. Concluding remarks and directions for further research are given in section 5.

2 PRINCIPLES OF THE CROSS-ENTROPY METHOD

In this section, we briefly review the cross-entropy method for the adaptive optimization of an importance sampling simulation. Only the aspects that are relevant for the rest of this paper are discussed; for more details, the reader is referred to [15] and [16].

2.1 BASICS

Assume that the change of measure (or "tilting") is parameterized by some vector $\boldsymbol{v}$; then an adaptive importance sampling procedure should try to find the value of $\boldsymbol{v}$ which results in minimal variance for the resulting estimator.

Another approach for choosing $\boldsymbol{v}$ was introduced in [16]. It is well known that always an importance sampling distribution (change of measure) exists which results in a zero-variance estimator, and that this distribution is precisely the original distribution conditioned on the occurrence of the rare event. In practice, this distribution may not be within the family of distributions that can be obtained by the change of measure parameterized by $\boldsymbol{v}$. However, if a simulation distribution is used that is in some sense "close" to the unattainable zero-variance distribution, then a low (but non-zero) variance should be expected. So, instead of choosing $\boldsymbol{v}$ such that the variance is minimized explicitly, one could try to devise a procedure that minimizes some distance measure between the distribution under the change of measure given by $\boldsymbol{v}$, and the distribution that would give zero variance. The latter distribution will henceforth be called the "zero-variance distribution".

Before proceeding with details of such a procedure, some more notation needs to be defined. The sample path of one replication of the simulation is denoted by Z. The function $I(Z)$ is the indicator function of the occurrence of the rare event in Z. We already defined $\boldsymbol{v}$ to denote the tilting vector; consistently with this, $f(Z, \boldsymbol{v})$ is the probability (or, for continuous systems, the probability density) of the sample path Z under the tilting $\boldsymbol{v}$, with $\boldsymbol{v} = 0$ corresponding to the original (untilted) system. The likelihood ratio associated with the sample path Z and the tilting vector $\boldsymbol{v}$ is denoted by $L(Z, \boldsymbol{v})$:

$$L(Z, \boldsymbol{v}) = \frac{f(Z, 0)}{f(Z, \boldsymbol{v})}.$$

Finally, $\mathbb{E}_{\boldsymbol{v}}$ denotes the expectation under the tilting $\boldsymbol{v}$.

A suitable "distance" measure for this procedure is the Kullback-Leibler cross-entropy, which is defined, see [19], as

$$CE = \int f(z) \ln \frac{f(z)}{g(z)} dz,$$

where $f(z)$ and $g(z)$ are two density functions whose "distance" is to be calculated. Note that this "distance" measure is not symmetric: in general, exchanging f and g in the above will result in a different value of CE.

We want to apply the Kullback-Leibler cross-entropy to measure the distance between the distribution to be used for the simulation (assumed to be of the form $f(z, \boldsymbol{v})$) and the zero-variance distribution. To do this, substitute $g(z) = f(z, \boldsymbol{v})$ (i.e., the distribution to be optimized by changing $\boldsymbol{v}$) and $f(z) = \rho_0 I(z) f(z, 0)$ with normalization factor $\rho_0^{-1} = \int I(z) f(z, 0) dz$ into the above; note that this $f(z)$ is the original distribution conditioned on the rare

event (i.e., the zero-variance distribution). Then we need to do the following minimization:

$$\begin{aligned}\boldsymbol{v}^\dagger &= \arg\min_{\boldsymbol{v}} \int \rho_0 I(z) f(z,0) \ln \frac{\rho_0 I(z) f(z,0)}{f(z,\boldsymbol{v})} dz \\ &= \arg\max_{\boldsymbol{v}} \int I(z) f(z,0) \ln f(z,\boldsymbol{v}) dz \\ &= \arg\max_{\boldsymbol{v}} \mathbb{E}_0 I(Z) \ln f(Z,\boldsymbol{v}), \end{aligned} \quad (1)$$

where $\boldsymbol{v}^\dagger$ denotes the value of $\boldsymbol{v}$ that minimizes the cross-entropy. In the above form, the equation is not useful, since we do not know $\mathbb{E}_0 I(Z) \ln f(Z,\boldsymbol{v})$. However, we can rewrite it as follows:

$$\boldsymbol{v}^\dagger = \arg\max_{\boldsymbol{v}} \mathbb{E}_{\boldsymbol{v}_j} I(Z) L(Z,\boldsymbol{v}_j) \ln f(Z,\boldsymbol{v}),$$

where $\boldsymbol{v}_j$ is any other tilting vector; we will later interpret it as the tilting vector used during the jth iteration of an iterative procedure. The above form can easily be approximated by a sum (stochastic counterpart of the expectation) over N samples from a simulation performed with tilting $\boldsymbol{v}_j$, thus yielding an approximation to $\boldsymbol{v}^\dagger$ which we call $\boldsymbol{v}_{j+1}$:

$$\boldsymbol{v}_{j+1} = \arg\max_{\boldsymbol{v}} \sum_{i=1}^{N} I(Z_i) L(Z_i,\boldsymbol{v}_j) \ln f(Z_i,\boldsymbol{v}), \quad (2)$$

where the Z_i are sample paths drawn under the tilting $\boldsymbol{v}_j$.

2.2 THE BASIC CROSS-ENTROPY ALGORITHM

Now we have all elements for an iterative procedure to approximate the optimal[1] tilting vector $\boldsymbol{v}^\dagger$:

1. Initialize as follows:
   $j := 1$ (iteration counter)
   $\boldsymbol{v}_1 :=$ initial tilting vector (see below)
2. Simulate N replications with tilting $\boldsymbol{v}_j$, yielding $Z_1 \ldots Z_N$.
3. Find the new tilting vector $\boldsymbol{v}_{j+1}$ from the maximization (2).
4. Increment j and repeat steps 2–4, until the tilting vector has converged (i.e., $\boldsymbol{v}_{j+1} \approx \boldsymbol{v}_j$).

Choosing the initial tilting vector $\boldsymbol{v}_1$ in step 1 is not trivial. The most obvious choice is to set $\boldsymbol{v}_1 = 0$, i.e., use the original transition probabilities. However, with that choice the rare event of interest will typically not be observed, making (2) unusable. In [16], this is solved by introducing an additional step in the algorithm, in which the rare event is temporarily modified into a less rare event (e.g., by choosing a lower overflow level). In the present paper, a different approach is used: we choose $\boldsymbol{v}_1$ on the basis of heuristics like exchanging the arrival rate with the bottleneck service rate. Although such a heuristic by itself does not produce an asymptotically efficient simulation in the cases considered here, it does provide a convenient starting point for the iterative process.

[1] in the sense that it minimizes the cross-entropy; this may not be equal to minimizing the variance of the estimator, although in practice the difference turns out to be small.

3 STATE-DEPENDENT TILTING

One application of the adaptive procedure described above is finding a "static" change of measure for queueing problems, i.e., finding the optimal arrival and service rates for simulation of a buffer overflow. In that case, the vector $\boldsymbol{v}$ just contains one component for every such rate that is allowed to change. Indeed, for many problems this turns out to work well; see [15] and [20] for examples. However, for many other problems no static change of measure seems to exist that gives an efficient simulation; for those systems, a less restrictive change of measure should be used, which can be obtained by allowing the arrival and service rates in the simulation to depend on the state. In this section we will do precisely that for DTMC models of queueing networks.

3.1 PRINCIPLES

A DTMC model is completely described by its initial probability distribution and its set of transition probabilities: the probabilities of going from one state to another. Since many DTMC models (e.g., for queueing systems) are derived from continuous time models with exponential arrival and service time distributions (CTMCs), the transition probabilities are typically calculated from transition rates: the probability of going from state i to state j is given by $\lambda_{ij} / \sum_k \lambda_{ik}$, where λ_{ij} is the transition rate from state i to state j, and k in $\sum_k$ runs over all states. In fact, for the cross-entropy calculations done in this section, it is more convenient to work with rates; the transition probabilities can trivially be calculated from the rates by normalizing their sum to 1. Collectively, all rates λ_{ij} will be referred to as a vector $\boldsymbol{\lambda}$.

In DTMC models, only one type of tilting is possible: changing the transition probabilities. Equivalently, one can change the transition rates of the corresponding CTMC model and calculate the transition probabilities for the DTMC from those, as shown above. It turns out that the latter approach is slightly simpler. So the aim is to find a set of transition rates $\boldsymbol{\lambda}$ which minimizes the cross-entropy.

Before deriving the actual cross-entropy minimization formula, let us first build a mathematical description of one replication Z of a DTMC simulation. Define the sequence z_i, $i = 1, 2, 3, \ldots$, which denotes the state of the system just before the ith transition in this replication Z. Denote by λ_{lm} the rate (or probability) of going from state l to state m. Then obviously the probability of the ith step

is

$$\frac{\lambda_{z_i z_{i+1}}}{\sum_k \lambda_{z_i k}},$$

where k runs over all states (or only those states that can be reached in one step from state z_i, since all other $\lambda_{z_i k}$ are 0). The total probability of the sample path Z is

$$\Pr(Z) = \prod_i \frac{\lambda_{z_i z_{i+1}}}{\sum_k \lambda_{z_i k}},$$

where i runs over all steps in the sample path.

Substitute the above expression for the probability of a sample path into equation (1); then we get the following expression for the optimal transition rate vector $\boldsymbol{\lambda}^\dagger$:

$$\begin{aligned}\boldsymbol{\lambda}^\dagger &= \arg\max_{\boldsymbol{\lambda}} \mathbb{E}_0 I(Z) \ln \prod_i \frac{\lambda_{z_i z_{i+1}}}{\sum_k \lambda_{z_i k}} \\ &= \arg\max_{\boldsymbol{\lambda}} \mathbb{E}_0 I(Z) \sum_i \left(\ln \lambda_{z_i z_{i+1}} - \ln \sum_k \lambda_{z_i k} \right).\end{aligned}$$

To find the maximum in the right-hand side, set the derivative with respect to λ_{lm} to 0, for any two states l and m:

$$0 = \mathbb{E}_0 I(Z) \sum_{i:z_i=l} \left(\frac{1_{(z_{i+1}=m)}}{\lambda^\dagger_{lm}} - \frac{1}{\sum_k \lambda^\dagger_{lk}} \right),$$

or, equivalently:

$$\frac{1}{\lambda^\dagger_{lm}} \mathbb{E}_0 I(Z) \sum_{i:z_i=l} 1_{(z_{i+1}=m)} = \frac{1}{\sum_k \lambda^\dagger_{lk}} \mathbb{E}_0 I(Z) \sum_{i:z_i=l} 1.$$

Thus, we find the following expression for the optimal transition probability q_{lm} from state l to state m:

$$\begin{aligned} q_{lm} &= \frac{\lambda^\dagger_{lm}}{\sum_k \lambda^\dagger_{lk}} \\ &= \frac{\mathbb{E}_0 I(Z) \sum_{i:z_i=l} 1_{(z_{i+1}=m)}}{\mathbb{E}_0 I(Z) \sum_{i:z_i=l} 1}. \end{aligned} \tag{3}$$

Of course, the expectations in the right-hand side are generally not known, but we can approximate them as follows:

$$\begin{aligned} q_{lm} &= \frac{\mathbb{E}_{\boldsymbol{\lambda}_j} I(Z) L(Z, \boldsymbol{\lambda}_j) \sum_{i:z_i=l} 1_{(z_{i+1}=m)}}{\mathbb{E}_{\boldsymbol{\lambda}_j} I(Z) L(Z, \boldsymbol{\lambda}_j) \sum_{i:z_i=l} 1} \\ &\approx \frac{\sum_{Z=Z_1}^{Z_N} I(Z) L(Z, \boldsymbol{\lambda}_j) \sum_{i:z_i=l} 1_{(z_{i+1}=m)}}{\sum_{Z=Z_1}^{Z_N} I(Z) L(Z, \boldsymbol{\lambda}_j) \sum_{i:z_i=l} 1}, \end{aligned} \tag{4}$$

where $\sum_{Z=Z_1}^{Z_N}$ is a sum over the sample paths from N replications, simulated with transition rates $\boldsymbol{\lambda}_j$ (i.e., from the jth iteration). Note that the factor $\sum_{i:z_i=l} 1$ in the denominator is just the number of visits to state l during replication Z, and that $\sum_{i:z_i=l} 1_{(z_{i+1}=m)}$ in the numerator is the number of those visits in which the transition to state m was chosen next. Consequently, the right-hand side of (4) can be interpreted as the *observed* conditional (on the occurrence of the rare event) probability of the transition from state l to state m; this is not surprising, since it is known that using the true conditional distributions for importance sampling yields a zero-variance estimator, as discussed before.

3.2 PRACTICAL PROBLEMS

Using the adaptive importance sampling algorithm from [15] or [16] with state-dependent parameters chosen according to (4) seems very simple. There are, however, practical difficulties. The cause of these is the enormous number of states that a typical queueing network can have. For example, a network with three queues and an overflow level of 50 for the total network population has 23461 states[2]. Doubling the overflow level to 100 multiplies this number of states by almost 8. If the rare event of interest is the overflow of one particular queue, other queues in the network can have an infinite size, thus making the number of states infinite.

One of the consequences of the enormous state space is that a lot of data needs to be stored: this takes a lot of memory capacity; but with present-day computers and the size of the queueing networks studied here, this is typically not a problem (except if the state space is infinite, of course). However, manipulating such a lot of data (e.g., in the smoothing techniques that will be discussed later) can be prohibitively time-consuming.

The accuracy of the estimations in the right-hand side of (4) is more problematic. The only sample paths that give a contribution to the sums in the numerator and denominator are those that reach the rare event (because of the $I(Z)$ factor) and pass through the state l (because of the summation over i for which $z_i = l$). The factor $I(Z)$ will typically not be a problem: the tilting used in the jth iteration is usually such that the event of interest is no longer rare. However, the tilting will not favor visits to states that are away from some optimal path to the rare event of interest. If the state space is multi-dimensional, this means that many states will not be visited often or at all, even under a tilting that makes the event of interest non-rare. States that are not visited at all during the N replications of a simulation yield $0/0$ (undefined) in the right-hand side of (4). And states that are visited only a few times make the quotient of sums a bad approximation of the quotient of expectations.

There is in fact a rather fundamental risk here: suppose the transition from some state l to another state m happens in only 10 % of all visits to state l, and state l is visited only 5 times during the N replications of a simulation. Then it is quite likely that in none of those 5 visits to state l, a transition to state m will be made. Consequently, using (4)

[2] This is the total number of ways to distribute among three distinct queues a total of 1 customer (3 ways), 2 indistinguishable customers (6 ways), 3 indistinguishable customers (10 ways), up to 50 indistinguishable customers.

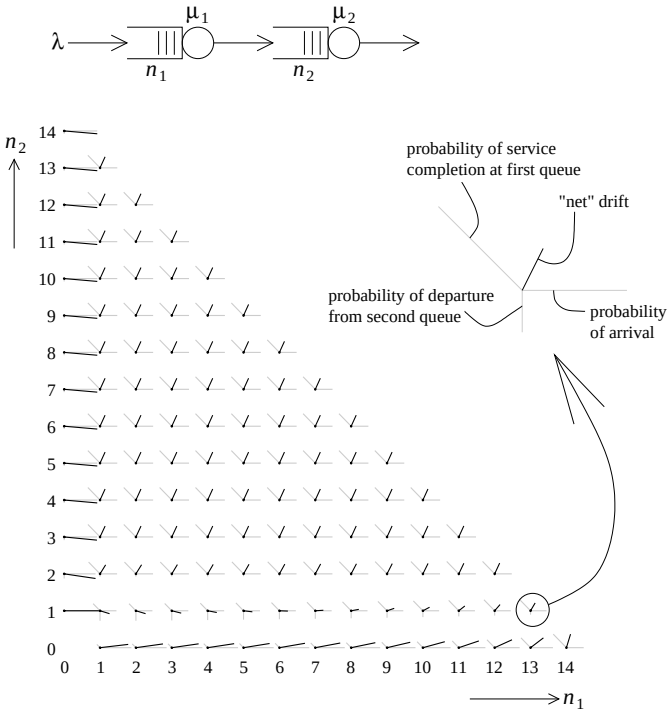

Figure 1: Optimal state-dependent transition probabilities for simulating the overflow of the total network population of two $(M)/M/1$ queues in tandem; $\lambda = 0.04$, $\mu_1 = \mu_2 = 0.48$, overflow level = 15.

to choose the simulation parameters for the next iteration would set the rate (probability) of this transition to 0, thus making the transition impossible. Then in the next simulation, surely no transitions from state l to state m will be observed, so this rate will again be set to 0 for the next iteration: it will remain at 0 forever, even though that is wrong if the transition has a non-zero probability in the untilted system, thus possibly resulting in a biased estimator.

The only case in which the above does not give a biased estimator is when the rare event of interest can no longer be reached after that particular transition has been made. As a matter of fact, all paths Z which contain such a transition necessarily have $I(Z) = 0$; as a consequence, (4) will automatically set the rate of such a transition to zero for the next iteration. Therefore, after the first iteration, *all* sample paths will reach the rare event.

3.3 DEALING WITH A LARGE NUMBER OF STATES

In order to get some ideas for dealing with large state spaces, take a look at figure 1. This figure has been drawn on the basis of numerical evaluation of (6) for every point of the state space of a network of two $(M)/M/1$ queues in tandem. For each state, it shows the optimal transition probabilities to each of its neighbour states as a gray line directed toward that neighbour, with a length proportional to the probability. Furthermore, for each state there is a

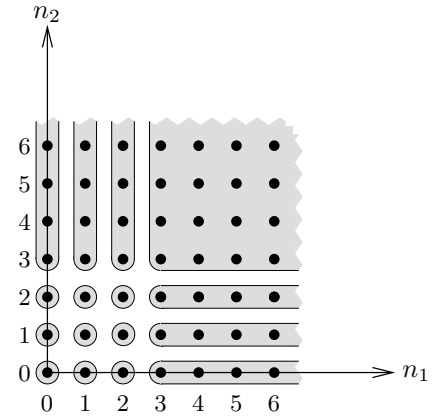

Figure 2: Grouping of states using three boundary layers in the state space of a two-queue system; n_i = level of ith queue.

black line which is just the vector sum of the three gray lines in that state: this can be interpreted as the "average drift" out of that state.

From this figure, it is apparent that the optimal transition probabilities don't just vary arbitrarily from state to state, but follow some patterns:

- States that are close to each other in the state space often have approximately the same transition probabilities.
- On a larger scale, the transition probabilities tend to be monotonous, smooth functions of the state variables (i.e., the number of customers in the queues).
- The strongest variations occur near the boundaries of the state space: only when a queue is nearly empty do the transition probabilities depend strongly on its level.

Each of these observations leads to a way of reducing the large-state-space problems, as explained below.

3.3.1 Boundary layers

The observation that when a queue's content is sufficiently large, the optimal transition probabilities are nearly independent of it, suggests that we can assign a single set of transition probabilities to a group of such states. Graphically, this is illustrated in figure 2, which shows the state space of a system with two queues. Gray shading is used to group states for which one or both queues contain at least 3 customers, which graphically turns into a picture with 3 layers along each boundary. Thus for each gray area, we store one set of transition probabilities, to be used during simulation for each visit to any of the states in that gray area.

In order to choose these transition probabilities on the basis of simulation results from the previous iteration, one

can readily modify (4) to calculate a transition probability q_{lm} on the basis of visits to a range of states (namely those in the same gray group) as follows:

$$q_{lm} \approx \frac{\sum_{l'} \sum_{Z=Z_1}^{Z_N} I(Z) L(Z, \boldsymbol{\lambda}_j) \sum_{i:z_i=l'} 1_{(z_{i+1}=m')}}{\sum_{l'} \sum_{Z=Z_1}^{Z_N} I(Z) L(Z, \boldsymbol{\lambda}_j) \sum_{i:z_i=l'} 1}. \tag{5}$$

Here l' runs over all those states of which the observations should be used for the estimation of the transition probabilities for state l, i.e., all states in the same gray group. State m' is the state whose position relative to state l' is the same as the position of state m relative to state l; e.g., if state m has 1 more customer in queue 2 than state l has, then state m' must also have 1 more customer in queue 2 than state l' has. This sounds complicated, but it just means considering the same "type of event" (e.g., arrival) for each state l' in the same gray area. Note that at state-space boundaries, some types of transitions may be disabled (e.g., departure from a queue that is empty at that boundary); then observations from states where such a transition is enabled cannot be combined with those from states where that transition is disabled, since the m' would not be well-defined. This problem does not arise with the boundary layer technique discussed here, but it does happen in other applications of the formula, discussed later.

Choosing the optimal number of boundary layers seems to be done best by trial and error. Using too few causes the change of measure to be insufficiently state-dependent, which may not only increase the estimator variance, but also cause the same problems that occur with a completely state-independent change of measure, and/or lead to convergence problems. Using unnecessarily many boundary layers mainly increases the amount of computation and storage required. Typical numbers of boundary layers used in practice were between 4 and 10. One could consider using different numbers of boundary layers along different boundaries; for example in figure 1, clearly fewer boundary layers would be needed along the vertical than along the horizontal axis. However, this kind of information is usually not available in advance.

3.3.2 Local average

As noted in section 3.2, it frequently happens that a transition probability estimate is not good due to a lack of observed sample paths involving that particular state. In such cases, we can apply (5) to use observations from surrounding states to try to obtain a more accurate estimate, on the basis of the observation that nearby states often have approximately equal optimal transition probabilities.

The local-average technique uses an adaptive procedure to decide from which states observations should be combined, separately for every state. It starts by using only the observations pertaining to the state of interest; i.e., using just (4). If this already gives good enough estimates of the transition probabilities, we're done for this state. If not, add states for which each queue's content differs at most 1 from the state of interest and apply (5); if the estimates are now good enough, we're done. If not, add states with queue contents differing at most 2, and so on until the estimates are good enough (or no more states are available). Whether estimates are good enough is judged by three criteria: (a) whether the number of observations (visits to the states or group of states) is not too small, e.g. at least 100; (b) whether all transition probabilities are non-zero if they should (cf. section 3.2); and (c) by demanding that the relative error associated with the estimates (5) is smaller than a threshold, e.g. 0.2.

Note that grouping of states as described above has two effects: on one hand, it typically reduces the variance of the transition-probability estimates because it uses more samples for each estimate; on the other hand, the optimal transition probabilities in reality differ from state to state, so combining observations results in a bias (but only the transition probabilities are biased; not the final estimate of the rare-event probability of interest). The adaptive algorithm needs to find a compromise between these two effects, by only grabbing more states until the accuracy seems satisfactory; obviously, parameters such as the upper limit imposed on the relative error determine where this compromise lies. Also the number of replications per simulation run has an influence: the more replications, the less grouping is needed. Fortunately, the process seems to be rather insensitive to these parameters, so trial and error (typically a lack of convergence) is a usable way to choose them.

3.3.3 Smoothing

The estimated transition probabilities on the basis of (4) or (5) tend to be somewhat noisy, and thus not the smooth and monotonic functions of the state variables that they should ideally be. This is natural, since they are simulation results. However, we may improve the accuracy by replacing the noisy data by a smooth function that to some extent fits the noisy data. This is called smoothing.

The smoothing method we used is based on spline-fitting. Only the main points of it will be described here, since many other smoothing methods from the literature would probably work just as well or better; the interested reader can find more details in [20].

Splines are piecewise-polynomial functions. I.e., the domain of the function (the state space in our application) is divided into pieces (we typically used square, cubic, etc., pieces with a size of 3 to 5 on each side), and over each piece a polynomial function is defined. At the boundaries between the pieces, the polynomials are chosen such that they are equal, and possibly also that one or more derivatives are equal, thus providing a smooth connection between the pieces. We chose to use polynomials that are cubic (i.e., of degree three) in each of the state variables. Such polynomials have enough degrees of freedom (parameters) to allow choosing the value of the polynomial itself,

and of its first derivative(s), at each corner point. Thus a natural way to fit such splines to the noisy data, is to estimate the value and the derivatives at each corner point from the noisy data; this can be accomplished by doing a weighted least-squares approximation around every corner point, giving the highest weight to data points closest to the corner point under consideration.

Note that the spline-fitting method is a rather simple method, with two shortcomings: (i) due to the weighted least-squares scheme, not all of the "raw" datapoints contribute equally to the final smooth function (which means information is thrown away), and (ii) the use of cubic polynomials puts a strong restriction on the set of smooth functions which is not supported theoretically. Smoothing methods without these shortcomings are known in the literature, and are expected to perform equally well or better.

However, the above method is still quite effective, as demonstrated in section 4.

3.3.4 Combination

In practice, two or all three of the above methods are combined. Before the simulations are started, the number of boundary layers is chosen; this is very effective to reduce the amount of data to be stored and processed. Next, the simulation is performed. Following this, the transition rates are calculated on the basis of the simulation results, using equation (4) and (5); the local average technique is used to group neighbouring states where necessary to obtain reliable estimates of the transition probabilities. Finally, the spline smoothing can be applied, if needed or desired. If the results after the local average step are already relatively good, spline smoothing may worsen the accuracy by imposing an unsuitable form on the data; on the other hand, if the data is rather noisy, the spline smoothing usually improves its accuracy. We will see examples of both in section 4.

3.4 THE COMPLETE ALGORITHM

To summarize, the adaptive importance sampling procedure looks like this:

1. Choose the number of boundary layers to be used; this should preferably be done in advance, since it significantly reduces memory usage.
   Also, other parameters need to be fixed in advance, such as the number of replications per iteration, maximum acceptable relative error in the local average method, etc.

2. Initialize the iteration counter $j := 1$, and the initial transition probabilities vector $\boldsymbol{\lambda}_1$. In the present paper, we choose $\boldsymbol{v}_1$ on the basis of heuristics like exchanging the arrival rate with the bottleneck service rate. Although such a heuristic by itself does not produce an asymptotically efficient simulation in the cases considered here, it does provide a convenient starting point for the iterative process.

3. Simulate N replications using transition probabilities λ_j. While doing this, keep track of the quantities needed to evaluate (4); i.e.,
   $$\sum_Z I(Z)L(Z,\boldsymbol{\lambda}_j) \sum_{i:z_i=l} 1_{(z_{i+1}=m)}$$
   for all pairs of states l and m; or rather, distinct groups of states resulting from the boundary layer technique. Furthermore, some other quantities needed for the local average algorithm need to be recorded, such as the number of times each state / distinct group of states is visited.

4. Apply the local average algorithm to these data, yielding estimates for the simulation parameters $\boldsymbol{\lambda}_{j+1}$.

5. Optionally, apply the spline-based smoothing to $\boldsymbol{\lambda}_{j+1}$.

6. Increment the iteration counter: $j := j + 1$.

7. Repeat steps 3–6 until convergence has been achieved.

3.5 THE VARIANCE OF THE ESTIMATOR

In this section, we derive some results about the variance of the estimator of the rare-event probability.

We start by showing that in principle, the method tends to converge to precisely that set of transition probabilities which would give a zero variance estimation of the quantity of interest, followed by a discussion of the reason why this zero variance is not actually reached.

Next, the influence of the statistical error in the transition probabilities (due to the fact that they are simulation results themselves) will be investigated. This leads to an explanation for the experimental observation that the variance of the rare-event probability estimator can decrease proportionally to the *square* of the number of replications used.

3.5.1 Zero variance

Start by rewriting the denominator of (3) as follows:

$$\begin{aligned}
\mathbb{E}_0 I(Z) \sum_{i:z_i=l} 1 &= \sum_{i=1}^{\infty} \mathbb{E}_0 I(Z) 1_{(z_i=l)} \\
&= \sum_{i=1}^{\infty} \Pr(I(Z)=1 \ \wedge \ z_i = l) \\
&= \sum_{i=1}^{\infty} \Pr(I(Z)=1 \mid z_i = l) \Pr(z_i = l) \\
&= \sum_{i=1}^{\infty} \pi_l \Pr(z_i = l),
\end{aligned}$$

where π_l is defined as the probability of reaching the rare event before an absorbing state is reached, starting from state l; note that π_{z_1} is the rare-event probability of interest, since z_1 is the initial state. Rewrite the numerator similarly:

$$\begin{aligned}
&\mathbb{E}_0 I(Z) \sum_{i:z_i=l} 1_{(z_{i+1}=m)} \\
&= \sum_{i=1}^{\infty} \mathbb{E}_0 I(Z) 1_{(z_i=l)} 1_{(z_{i+1}=m)} \\
&= \sum_{i=1}^{\infty} \Pr(I(Z)=1 \wedge z_i = l \wedge z_{i+1} = m) \\
&= \sum_{i=1}^{\infty} \Pr(I(Z)=1 \mid z_i = l \wedge z_{i+1} = m) \\
&\qquad \cdot \Pr(z_{i+1} = m \mid z_i = l) \cdot \Pr(z_i = l) \\
&= \sum_{i=1}^{\infty} \pi_m p_{lm} \Pr(z_i = l),
\end{aligned}$$

where p_{lm} is the untilted transition probability from state l to state m. By substituting the above into (3), we find the following expression for the optimal transition probabilities:

$$\begin{aligned}
q_{lm} &= \frac{\lambda^\dagger_{lm}}{\sum_k \lambda^\dagger_{lk}} = \frac{\sum_{i=1}^{\infty} \pi_m p_{lm} \Pr(z_i = l)}{\sum_{i=1}^{\infty} \pi_l \Pr(z_i = l)} \\
&= \frac{\pi_m p_{lm} \sum_{i=1}^{\infty} \Pr(z_i = l)}{\pi_l \sum_{i=1}^{\infty} \Pr(z_i = l)} = \frac{\pi_m p_{lm}}{\pi_l}. \qquad (6)
\end{aligned}$$

One can easily recognise the right-hand side as the conditional probability of going from state l to state m, given that the rare event will be reached.

Now consider a random path Z leading to the rare event, containing n_Z steps. The probability of this path in the tilted simulation is

$$\begin{aligned}
\prod_{i=1}^{n_Z} q_{z_i z_{i+1}} &= \frac{p_{z_1 z_2} \pi_{z_2}}{\pi_{z_1}} \cdot \frac{p_{z_2 z_3} \pi_{z_3}}{\pi_{z_2}} \cdots \frac{p_{z_{n_Z} z_{n_Z+1}} \pi_{z_{n_Z+1}}}{\pi_{z_{n_Z}}} \\
&= \frac{\pi_{z_{n_Z+1}}}{\pi_{z_1}} \prod_{i=1}^{n_Z} p_{z_i z_{i+1}}. \qquad (7)
\end{aligned}$$

The probability of the same path in the untilted system is just $\prod_{i=1}^{n_Z} p_{z_i z_{i+1}}$, so the likelihood ratio is $\pi_{z_1}/\pi_{z_{n_Z+1}}$. Since Z was defined as a path leading to the rare event, its last state z_{n_Z+1} must be the rare event itself; therefore $\pi_{z_{n_Z+1}} = 1$. Consequently, the likelihood ratio of the path is just π_{z_1}, which is (by definition) the rare-event probability of interest, and thus a constant independent of the path. Since *all* sample paths in the tilted system reach the rare event (see[3] the last paragraph of section 3.2), and thus have this same likelihood ratio, the variance of the estimator is zero.

In practice, zero variance is not reached of course. There are two reasons for this:

- The new transition probabilities are estimated from simulation results, which obviously have a statistical error; the effect of this is discussed further below. Note that this error can be made arbitrarily small by increasing the number of replications used in the simulation.
- The techniques for dealing with the large state space discussed in section 3.3 introduce errors. E.g., the boundary-layer method makes all transition probabilities far away from the boundaries equal, and the spline approximation only allows transition probabilities which fit the form of the splines used; these errors do not disappear with increasing number of replications. (Note that the local average technique is different in this respect: with increasing number of replications, fewer states need to be grouped, so the error introduced also decreases.)

3.5.2 *Influence of the statistical error in the transition probability estimates*

As discussed above, the non-zero variance of the rare-event probability estimator is caused by the statistical error in the transition probability estimates. In the following, a model of these statistical errors is constructed, which is subsequently used to calculate the resulting variance of the rare-event probability estimator. This serves to give insight into the relative importance of the iteration which estimates the transition probabilities, and the subsequent iteration which uses these to estimate the rare-event probability.

For clarity, let us refer to the simulation iteration in which the rare-event probability is estimated[4] as the "last" simulation (or iteration). Then this simulation uses transition probabilities which are simulation results from the "second-last" simulation. On its turn, the second-last iteration uses transition probabilities which are simulation results from the third-last simulation, and so on.

As before, define q_{lm} as the optimal transition probability from state l to m. Furthermore, define $\hat{q}_{lm}$ as the simulation estimate of this probability (obtained in the second-last iteration), which contains a statistical error. Assume that this error has a normal distribution[5] with variance σ^2_{lm}

[3] Or observe that the expectation (in the tilted system) of $I(Z)L(Z)$ is the probability of interest π_{z_1}. We have just shown that, given $I(Z) = 1$, L itself equals π_{z_1}, so the expectation of $I(Z)L(Z)$ can be equal to π_{z_1} only if $I(Z) = 1$ for any sample path in the tilted system (remember that $I(Z)$ is an indicator function and thus is either 0 or 1). Thus, every sample path in the tilted system reaches the rare event.

[4] Of course, one can estimate the rare-event probability (as a by-product) in every iteration, and this will in fact be done in the experiments in section 4. However, for the purpose of discussing the estimate's variance, we need to refer clearly to one estimate; that's why this definition is made.

[5] Although a normal distribution is usually an appropriate approximation for a simulation error, it is not completely realistic here. It does not take into account the fact that all transition probabilities from a state sum up to 1. Furthermore, it does not take into account that all transition probabilities must be between 0 and 1: the normal distribution has infinite tails. However, for small variances these effects can be neglected.

as follows:

$$\hat{q}_{lm} = q_{lm} \cdot (1 + f_{lm}) \qquad \text{with} \qquad f_{lm} \sim \mathcal{N}\big(0, \sigma^2_{lm}\big),$$

where $\mathcal{N}(m, s)$ is the normal distribution with mean m and variance s. Using the optimal transition probabilities q_{lm}, the likelihood ratio would be a constant (independent of the sample path), but using the simulation estimates $\hat{q}_{lm}$, this is not the case. The ratio of the ideal (L_Z) and the practical ($\hat{L}_Z$) likelihood ratio on a sample path Z is

$$\begin{aligned}\frac{L_Z}{\hat{L}_Z} &= \prod_{i=1}^{n_Z}(1 + f_{z_i z_{i+1}}) \\ &= 1 + \sum_{i=1}^{n_Z} f_{z_i z_{i+1}} + \sum_{i=1}^{n_Z}\sum_{j=1}^{n_Z} f_{z_i z_{i+1}} f_{z_j z_{j+1}} + \dots \\ &\approx 1 + \sum_{i=1}^{n_Z} f_{z_i z_{i+1}} \sim \mathcal{N}\big(1, \sum_{i=1}^{n_Z} \sigma^2_{z_i z_{i+1}}\big),\end{aligned}$$

where it is assumed that typically $\sigma_{lm} < 1/n_Z$, so the higher order terms can indeed be neglected at the $\approx$ sign. Furthermore, the above assumes that the errors f_{lm} are independent of each other, and that the sample path does not (or rarely) visit a state twice; otherwise, the variance of the resulting normal distribution would be different due to the dependencies.

We see that the variance of the likelihood ratio (which is the variance of the estimator of the rare-event probability, since all sample paths reach the rare event) is proportional to the variance of the individual transition probabilities. Since these individual transition probabilities are simulation results from the second-last simulation, their variance is inversely proportional to the number of replications used in that simulation. Clearly, the variance of the rare-event probability estimator is also inversely proportional to the number of replications used in the last simulation. Thus, we find that the variance of the final estimator is inversely proportional to the *product* of the number of replications used in the last and in the second-last simulation; or, if the same number of replications is used in both of them, the variance is inversely proportional to the *square* of the number of replications. This phenomenon will be demonstrated experimentally in section 4.

One might be tempted to take this reasoning one step further: the accuracy of the transition probabilities used in the last iteration not only depends on the number of replications in the second-last iteration, but also on the accuracy of the transition probabilities used in that second-last simulation; the latter depends on the number of replications of the third-last iteration, so the accuracy of the rare-event probability estimation should also depend on the number of replications in the third-last iteration. This can indeed be the case, but it will not nearly be as strong as the dependence on the number of replications in the second-last iteration. The reason for this is that the estimation of the transition probabilities is inherently *not* a zero-variance simulation: even if the third-last iteration had yielded perfect estimates of the transition probabilities to be used in the second-last simulation, the latter would not give zero-variance estimates of the transition probabilities for the last simulation.

3.5.3 Relationship with optimal splitting (RESTART) simulation

Using a calculation similar to (7), one easily sees that under the zero-variance change of measure, the likelihood ratio on any path from the initial state z_1 to a given state z equals π_z/π_{z_1}. Clearly, this is independent of the path taken to reach state z. (This independence could also be found by first assuming that there are two paths from z_1 to z with different likelihood ratios, from which it follows that there are paths to the rare event with different likelihood ratios, which contradicts the assumption of zero variance.)

According to Theorem 2 in Chapter 5 of [21], the optimal "importance function" for splitting simulation is such that its levelsets are the sets of states from which the rare event is reached with equal probability. Thus, the probability π_z itself is such an optimal importance function. Another optimal importance function is π_z/π_{z_1}, which, as shown above, is the likelihood ratio under importance sampling simulation with the zero-variance change of measure.

Note however that this is only of theoretical interest. Practically, it would not be efficient to first derive the zero-variance change of measure for a system, and then use that to perform a splitting simulation which gives an estimator with non-zero variance.

4 EXPERIMENTAL RESULTS

In sections 4.1–4.3, overflows in a simple Jackson network will be considered. This network consists of four queues in tandem, with arrival and service rates chosen in the region where the standard state-independent change of measure (exchanging the arrival rate with the bottleneck service rate) does not work well according to [7]: the arrival rate is 0.09, the service rates of the first through fourth queue are 0.23, 0.227, 0.227 and 0.226, respectively. The rare event of interest is the total network population reaching a high level, starting from 0 and before returning to 0 again.

For all experiments, the boundary layer technique was used to reduce the enormous state space; 10 boundary layers turned out to work well, but possibly fewer would have been sufficient. Furthermore, the local average technique was used. The spline-smoothing was only used in some cases, as indicated below.

4.1 RESULTS FOR OVERFLOW LEVEL 50

The results for an overflow level of 50 are presented in figure 3, both without and with spline-based smoothing.

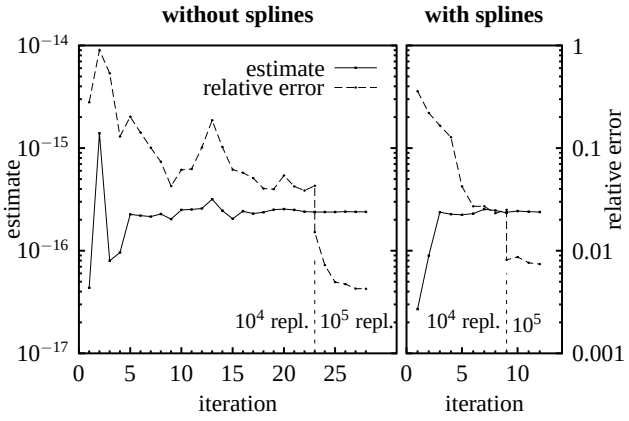

Figure 3: Results for the four-node tandem queue, overflow level = 50.

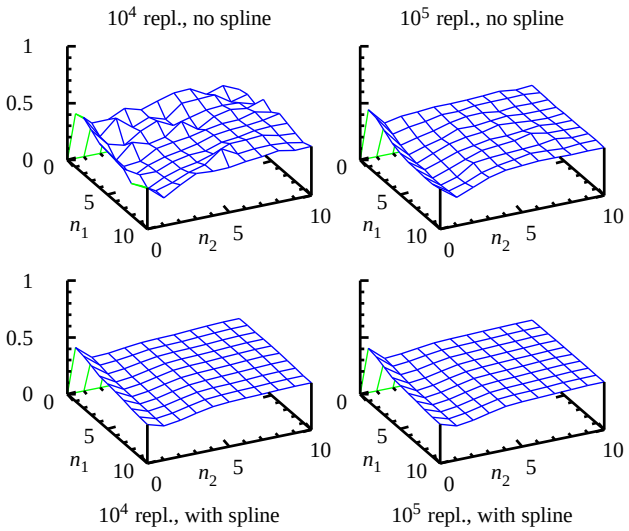

Figure 4: State-dependent transition probabilities.

Along the horizontal axis, the iteration number is indicated. Vertically, the estimate of the overflow probability and its relative error (standard deviation from the simulation, divided by the estimate itself) are shown as two lines in the graph. At the first iteration, a static tilting according to the well-known heuristic of exchanging the arrival rate with the bottleneck service rate was used, to get things started. In the experiments without spline-based smoothing (upper graph), 10^4 replications were used per iteration up to the 23rd iteration; the 23rd iteration was performed twice, once with 10^4 and once with 10^5 replications, and all later iterations used 10^5 replications. With spline-based smoothing (lower graph), the switch from 10^4 to 10^5 replications per iteration was made at the 9th instead of the 23rd iteration.

Obviously, the spline smoothing is quite beneficial to the convergence in this case: without splines, the convergence is rather slow and irregular, with a major excursion around the 13th iteration, whereas with spline smoothing the convergence is quick and monotonic, and the resulting relative error at 10^4 replications is smaller by almost a factor of 2.

Next, note what happens when the number of replications is increased: at the 23rd (without splines) and 9th (with splines) iteration, the same simulation was done with 10^4 and 10^5 replications; the relative error of the latter clearly is about a factor of $\sqrt{10}$ smaller, as it should. However, without splines the relative error continues to decrease in the next iteration: this is a consequence of the fact that these later iterations have better transition probabilities because those have been obtained with 10^5 instead of 10^4 replications, as discussed in section 3.5. In the end, the relative error has decreased by a factor of 10 in total. With splines, this does not happen: the relative error does not significantly decrease further, and in fact is higher than without splines; apparently, the spline form does not fit the optimal state-dependence well enough.

Figure 4 serves to give an idea of how the transition probabilities depend on the state in this particular problem. Of course, since we have up to five transition probabilities and a four-dimensional state space, it is hardly feasible to give a complete picture. Therefore, only the probability of the transition corresponding to a service completion at the first queue is shown, as a function of the contents n_1 and n_2 of the first and second queues, respectively, while the third and fourth queues are empty. Clearly, the splines perform a very effective smoothing: most of the noise disappears. On the other hand, the splines used here are apparently not able to completely follow the true functions: the "dip" at $n_2 = 1$ is much deeper without splines (only sufficiently visible in the 10^5-replications plot) than with splines. This agrees with the experimental observation that at 10^5 replications, the final estimate is more accurate when the transition probabilities are not restricted by applying splines.

4.2 Results for overflow level 200

For the case of an overflow level of 200, figure 5 shows the simulation results. For this problem, all three techniques (local average, 10 boundary layers, and splines) were used initially (up to iteration 16). After convergence had been achieved, the number of replications was increased, resulting in a branch with splines and a branch without splines in the graph.

It seems as if the convergence process can be divided into two phases. During the first phase, the estimate is quite inaccurate (typically too low), but it approaches the correct value; in the present example, this phase comprises iterations 1 through 7. During the second phase, the estimate stays correct, and the relative error decreases to its final value; in the present example this happens during iterations 7 through 11. These phases can also be recognized in the results with overflow level 50 in figure 3.

Note, like before, the strong decrease of the relative error after increasing the number of replications by a factor of 10, and the fact that switching off spline smoothing at

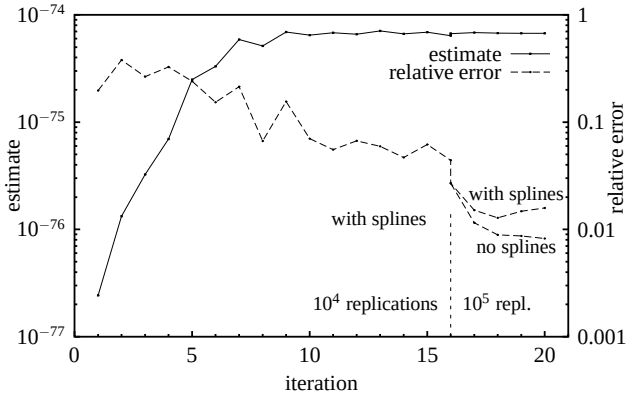

Figure 5: Results for the four-node tandem queue, overflow level = 200.

that point is beneficial.

4.3 ASYMPTOTIC EFFICIENCY

Results from the above experiments, and from repetitions of those experiments at overflow levels 25 and 100, are shown in Table 1. All of these experiments used the same number of replications per iteration (10^5) and no splines in the final iterations. It is clear from the table that the relative error grows with the overflow level, but clearly less than exponentially fast, while the probability of interest does decrease exponentially fast. This demonstrates the asymptotic efficiency of the method for this problem.

Table 1: Test of asymptotic efficiency (four-node tandem queue).

level	exact	estimate	rel.error
25	$3.5283 \cdot 10^{-07}$	$3.504 \cdot 10^{-07}$	0.0026
50	–	$2.396 \cdot 10^{-16}$	0.0042
100	–	$1.422 \cdot 10^{-35}$	0.0044
200	–	$6.722 \cdot 10^{-75}$	0.0082

The table also shows an exact (numerical) calculation of the overflow probability for an overflow level of 25. Comparing this with the simulation estimate shows a good agreement. No exact numbers could be calculated for higher overflow levels due to the large state space involved.

4.4 AN EXAMPLE WITH ROUTING AND FEEDBACK

Figure 6 shows a more complicated network, comprising five queues, three points at which random routing (with probability 0.5) occurs, and two feedback loops. Like in the previous section, we want to estimate the overflow probability of the total network population. The arrival rate λ is 3, and the service rates are μ_1=40, μ_2=20, μ_3=50, μ_4=50 and μ_5=60. Note that this choice makes the load of all queues equal (0.1), which generally causes state-*in*dependent tilting to work badly.

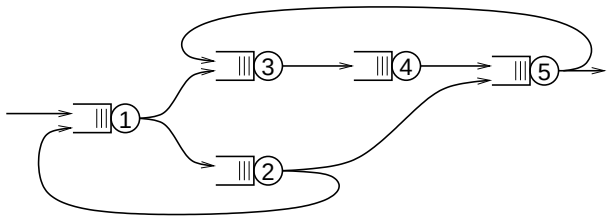

Figure 6: A five-node Jackson network.

Table 2: Results for the five-node Jackson network, showing the convergence for progressively higher overflow levels.

iter.	overfl. level	repl.	bnd. lay.	estimate	relative error
5	20	10^4	4	$7.752 \cdot 10^{-16}$	0.0247
9	50	10^5	4	$2.619 \cdot 10^{-44}$	0.0100
15	50	10^6	4	$2.610 \cdot 10^{-44}$	0.0014
18	100	10^5	7	$3.761 \cdot 10^{-93}$	0.0194
27	100	10^6	7	$3.964 \cdot 10^{-93}$	0.0017

Simulation results using the method described in this paper are given in Table 2. These results were not obtained independently of each other, but by increasing the overflow level in several steps; this was necessary to obtain convergence. We started with an overflow level of 20, and 4 boundary layers. After converging in this case, we increased the overflow level to 50 and continued iterating (i.e., starting with the transition probabilities that were found for the overflow level of 20). Note that starting from scratch with overflow level 50 did not lead to convergence, for both 10^4 and 10^5 replications per iteration. After obtaining convergence for the overflow level of 50, the overflow level was increased to 100, and iterations continued. With 4 boundary layers, no convergence was obtained at this overflow level, but with 7 boundary layers the procedure converged nicely. So we see that in hard cases like this one, convergence can be improved by first starting with a lower overflow level; furthermore, fewer boundary layers can be used at the lower overflow level to reduce the computational effort.

5 CONCLUSIONS AND FURTHER RESEARCH

In this paper, we have proposed an importance sampling simulation method with two important features: the change of measure is completely state-dependent, and a cross-entropy-based adaptive method is used to approximate the optimal change of measure. To show the method's performance, we have applied it to estimate the overflow probability of the total population of two Jackson networks. This simulation has been shown to be asymptotically efficient, at a parameter setting at which asymptotically efficient simulation is not obtained with state-independent tilting. Furthermore, we have demonstrated

experimentally, and explained analytically, that with this method the relative error can decrease faster than proportional to the square root of the total simulation effort. The method has also been applied successfully to other rare-event problems in Jackson-like networks, such as overflow of a non-bottleneck queue in a network with bounded queues; see [20] for more details.

However, all of the systems considered so far are modelled by DTMCs, and the number of queues is not too large to avoid state space explosion. This indicates two obvious directions for future work: extension to non-DTMC systems, and developing more efficient methods for handling large state spaces. Furthermore, it may be possible to improve the method's convergence by combining observations from several iterations.

In the present paper, the good performance of the method has only been demonstrated experimentally. Another direction for further research would therefore be providing more solid mathematical foundations, such as a proof of the convergence of the tilting vector.

ACKNOWLEDGMENT

The authors wish to thank prof. Reuven Rubinstein for useful discussions.

Manuscript received on ...

REFERENCES

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

[2] S. Asmussen and R.Y. Rubinstein. Complexity properties of steady-state rare-events simulation in queueing models. In J. Dshalalow, editor, *Advances in Queueing: Theory, Methods and Open Problems*, pages 429–462. CRC Press, 1995.

[3] M. Villén-Altamirano and J. Villén-Altamirano. RESTART: A straightforward method of fast simulation of rare event. In *1994 Winter Simulation Conference*, pages 282–289, 1994.

[4] S. Parekh and J. Walrand. A quick simulation method for excessive backlogs in networks of queues. *IEEE Transactions on Automatic Control*, Vol. 34, pages 54–66, 1989.

[5] J. S. Sadowsky. Large deviations theory and efficient simulation of excessive backlogs in a $GI/GI/m$ queue. *IEEE Transaction on Automatic Control*, Vol. 36, pages 1383–1394, 1991.

[6] M. R. Frater, T. M. Lennon, and B. D. O. Anderson. Optimally efficient estimation of the statistics of rare events in queueing networks. *IEEE Transactions on Automatic Control*, Vol. 36, pages 1395–1405, 1991.

[7] P. Glasserman and S.-G. Kou. Analysis of an importance sampling estimator for tandem queues. *ACM Transactions on Modeling and Computer Simulation*, Vol. 5, No. 1, pages 22–42, January 1995.

[8] D. P. Kroese and V. F. Nicola. Efficient simulation of a tandem jackson network. In *1999 Winter Simulation Conference*, pages 411–419, 1999.

[9] P. E. Heegaard. A scheme for adaptive biasing in importance sampling. *AEÜ International Journal of Electronics and Communications*, Vol. 52, pages 172–182, 1998.

[10] M. Cottrell, J.-C. Fort, and G. Malgouyres. Large deviations and rare events in the study of stochastic algorithms. *IEEE Transactions on Automatic Control*, Vol. 28, No. 9, pages 907–920, 1983.

[11] M. Devetsikiotis and J. K. Townsend. Statistical optimization of dynamic importance sampling parameters for efficient simulation of communication networks. *IEEE/ACM Transactions on Networking*, Vol. 1, pages 293–305, 1993.

[12] M. Devetsikiotis and J. K. Townsend. An algorithmic approach to the optimization of importance sampling parameters in digital communication system simulation. *IEEE Transactions on Communications*, Vol. 41, pages 1464–1473, 1993.

[13] W. A. Al-Qaq, M. Devetsikiotis, and J. K. Townsend. Stochastic gradient optimization of importance sampling for the efficient simulation of digital communication systems. *IEEE Transactions on Communications*, Vol. 43, pages 2975–2985, 1995.

[14] R. Y. Rubinstein. Optimization of computer simulation models with rare events. *European Journal of Operations Research*, Vol. 99, pages 89–112, 1997.

[15] R. Y. Rubinstein and B. Melamed. *Modern Simulation and Modeling*. Wiley, 1998.

[16] R. Y. Rubinstein. Rare event simulation via cross-entropy and importance sampling. In *Second International Workshop on Rare Event Simulation, RESIM'99*, pages 1–17, 1999.

[17] D. Lieber. *The cross-entropy method for estimating probabilities of rare events*. PhD thesis, William Davidson Faculty of Industrial Engineering and Management, Technion, Israel, 1999.

[18] P. T. de Boer, V. F. Nicola, and R. Y. Rubinstein. Adaptive importance sampling simulation of queueing networks. In *2000 Winter Simulation Conference*, pages 646–655, 2000.

[19] J. N. Kapur and H. K. Kesavan. *Entropy Optimization Principles with Applications*. Academic Press, 1992.

[20] P. T. de Boer. *Analysis and efficient simulation of queueing models of telecommunication systems*. PhD thesis, University of Twente, 2000.

[21] M. J. J. Garvels. *The splitting method in rare event simulation*. PhD thesis, University of Twente, 2000.