

Fast Simulation of Rare Events Using Reverse-Time Models

Michael R. Frater
Rodney A. Kennedy

Robert R. Bitmead
Brian D.O. Anderson

Dept Systems Engineering,
Australian National University,
G.P.O. Box 4,
Canberra, ACT 2601
AUSTRALIA

Work supported by the Australian Telecommunications and Electronics Research Board and the A.N.U. Centre for Information Science Research.

Abstract *We consider the estimation by simulation of probabilities of rare but potentially damaging events in Markovian systems. The problem is that the rarity of the event mitigates against its frequent occurrence in simulation. Here we show how a reverse-time model may be simulated in place of the forward-time system in order to provide significantly faster estimates of the rare probabilities. An example of an ALOHA system will be used and simulation speed up compared with alternative methods based on Large Deviations theory.*

1 Introduction

A number of works, e.g. [1], have studied the problem of finding the expected time for the state of an asymptotically stable process with a noise input to leave some specified bounded region containing the origin. In [2] these results were extended in relation to diffusion processes, while [3] applied them to discrete-time processes, and also to ALOHA-type processes. Similar results relating to queueing networks are derived in [4], and extended in [5]. In communications systems, the exiting of the neighbourhood of the origin by the state may be related to critical events such as buffer overflows. Such events should be rare, but still quantitative estimates of their probabilities are needed and must be obtained by simulation. However, the rarity of the event mitigates against its direct simulation, and here we consider the alternative of simulation of a related system which yields estimates of the original probabilities much faster.

In [3], it is shown that the exponential change of measure associated with Large Deviations provides an asymptotically optimal simulation. *Asymptotic* here refers to the behaviour with increasing rarity of the event, and *optimal* refers to the estimate variance amongst the class of simulation systems under an exponential change of measure. In this paper, we show by way of an example using uncontrolled ALOHA systems [3] that significantly greater speed up can be obtained by using the reverse-time model of the original process in cases of practical interest.

Section 2 describes the uncontrolled ALOHA system. Section 3 poses the probability changing problem for the simulation of rare events, and summarizes the Large Deviations results of [3]. This is followed in Section 4 by a discussion of the reverse-time model and how it is used to speed up simulation. In Section 5 a formula for the speed up factor is derived. In Section 6, the speed up obtained by using the reverse-time model is compared to that obtained using the model suggested by [3].

2 ALOHA System

ALOHA is a packet radio system, shared between many users. It is assumed that packets are of constant duration, which is approximately 10^{-4} seconds in practice. The channel time is divided into slots whose length is the same as the packet length, and it is assumed that all packets start at the beginning of a slot. If more than one user attempts to transmit a packet in the same slot, there is a clash, and both users are said to be blocked. They wait for a random amount of time before attempting to retransmit. In the uncontrolled ALOHA system examined here, a blocked user tosses a weighted coin¹ to decide whether or not to retransmit in a given slot. The probability of retransmission is p . The global arrival process is assumed to be Poisson with parameter α , and the state of the system at time t , denoted $x(t)$, is the number of blocked users.

The transition probabilities for the ALOHA system starting at state i to go to state j are [3]:

$$p_{ij} = P(x(t+1) = j \mid x(t) = i) \quad (1)$$

$$= \begin{cases} 0 & \text{for } j \leq i - 2 \\ ip(1-p)^{i-1}e^{-\alpha} & \text{for } j = i - 1 \\ (1-p)^i\alpha e^{-\alpha} \\ \quad + [1 - ip(1-p)^{i-1}]e^{-\alpha} & \text{for } j = i \\ \alpha e^{-\alpha}[1 - (1-p)^i] & \text{for } j = i + 1 \\ \frac{\alpha^{j-i}}{(j-i)!}e^{-\alpha} & \text{for } j \geq i + 2 \end{cases}$$

The uncontrolled ALOHA system has two equilibrium states.² The lower, n_0 , is stable, while the upper, n_c , is unstable. Thus, when the number of blocked users exceeds n_c , the expectation of the next state is even greater, and all users rapidly become blocked. This event is certain to happen at some time, and the problem treated by [3] is to estimate via simulation the expected time for an ALOHA system to become unstable, starting with n_0 users blocked.

To estimate the transition time to the (effectively) absorbing state $i > n_c$, we construct a recurrent Markov chain by forcing states $i > n_c$ to transit to n_0 deterministically. When this is done the states $i < n_c$ are no longer transient, and the expected time to exit is very nearly the recurrence time of n_c . In future, when we speak of an ALOHA process, we will be referring to this recurrent modification of the process.

3 Large Deviations Based Results

Here we shall describe the results of Cottrell *et al.* [3] and Parekh and Walrand [4] which deal with the use of the theory of Large Deviations to achieve speed up in simulations designed to provide empirical estimates of probabilities of rare events. Specifically, Cottrell *et al.* deal with the ALOHA example and derive estimates of overflow probabilities, as defined below.

Consider the empirical, i.e. simulation, estimation of the probability of an event A from a sequence of n independent trials. If the estimate is derived as

$$\widehat{P}(A) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}_A(\omega_i), \quad (2)$$

¹Unweighted coins were tried, but failed to return from the toss.

²An equilibrium state occurs where the expected value of the next state given the current state equals the current state.

where $\mathbf{1}_A(\omega_i)$ is the indicator function that the event A occurs in the i^{th} trial, then the variance of this estimator is $P(A)(1 - P(A))/n$. If A is rare, then the number of trials, n , required to produce an estimate of $P(A)$ with high confidence is very large, because the 'strike rate' of the event in the trial is so low. The method used to speed up the simulation evaluation of $P(A)$ is to replace the underlying probability distribution in the trials, $P(\omega)$, by a distribution $P^*(\omega)$, where $P(\omega)$ is absolutely continuous with respect to $P^*(\omega)$, under which the event A is much more likely. Using this P^* as the basis for the simulation we may introduce the estimator,

$$\widehat{P}(A) = \frac{1}{n} \sum_{i=1}^n \frac{dP}{dP^*}(\omega_i) \mathbf{1}_A^*(\omega_i), \quad (3)$$

where now dP/dP^* , essentially the Radon-Nikodym derivative of the two distributions, is used to map the probabilities between the simulation and the original system. The variance of this estimator is [3]

$$\frac{1}{n} \left(\int_A \left(\frac{dP}{dP^*} \right)^2 dP^*(\omega) - P^2(A) \right), \quad (4)$$

which we, via careful selection of P^* , would like to be smaller than that of (2).

This procedure of substitution of probability distributions in simulation problems for rare events is known as *Importance Sampling* in the Communications literature. Recently, Cottrell et al. have proposed an explicit procedure for the generation of substitute probabilities, P^* , for the speed up of simulations of rare events with finite-state Markov processes, which we shall now explain.

Let $x(t)$ be the state of a discrete-time markov chain at time t with 'jump probability' distribution μ_x , i.e. the distribution of the increment $V(x(t)) = x(t+1) - x(t)$, where $x(t) = x$. Thus μ_x plays the role here of P . The transformed chain is defined by the new jump distribution μ_x^*

$$d\mu_x^*(y) = \frac{e^{v_x y}}{\hat{\mu}_x(v_x)} d\mu_x(y), \quad (5)$$

where $\hat{\mu}_x(\cdot)$ is the moment generating function of the distribution μ_x , viz.

$$\hat{\mu}_x(z) = \int e^{zy} d\mu_x(y) \quad (6)$$

and v_x (a state-dependent parameterization) takes the positive real solution of

$$\hat{\mu}_x(v_x) = 1. \quad (7)$$

P^* is the probability measure associated with this μ^* distribution.

This change of probability measure is called 'exponential' and is closely related to the solution of Large Deviations problems with the original P system using the Freidlin-Wentzell theory [1] and Cramér transforms. The important feature of this change of measure is that it becomes asymptotically optimal amongst the class of exponential changes as the jump distribution alters to make the simulated event, A , more rare (see [3] for more details.) This optimality is in terms of minimizing the variance of $\widehat{P}(A)$ given by (4). No optimality can be claimed, however, when either non-exponential changes are included or the simulated event is non-vanishingly rare.

For the ALOHA example of the previous section Cottrell et al. [3] establish considerable simulation speeds up using this technique and compare this to diffusion approximation methods

(which offer another approach again to the problem of determining the expected time to instability.) To apply their method, one considers replacing the transition probabilities from state i to state j , p_{ij} of (1), by

$$p_{ij}^* = \begin{cases} e^{\lambda_i(j-i)} & , n_o < i < n_c \\ p_{ij} & , i \leq n_o, i \geq n_c \end{cases} \quad (8)$$

where for each $n_o < i < n_c$, λ_i is chosen as the positive solution of $\hat{\mu}_i(\lambda_i) = 1$, or, after a little arithmetic, as the positive solution of $(1-p)^{i-1} e^{\lambda_i i - \alpha} \left\{ (1 - e^{\lambda_i}) [e^{\lambda_i i} p + \alpha(1-p)] + \exp(\alpha e^{\lambda_i}) \right\} = 1$. The recurrent modification of this distribution can be used to enable simulation.

Here we shall present an alternative approach via reverse-time models which indicates that still further considerable improvement is possible by not restricting to exponential changes of distribution; further this improvement may be achieved in cases where the event probability is not infinitesimal. The issue here is that analytical methods, such as the theory of Large Deviations, which are based strongly upon tail properties of the distribution functions involved may better be replaced by non tail-dependent procedures for nonasymptotic situations. Unfortunately, however, and as will be clarified in the sequel, the use of reverse-time models frequently may obviate the need to perform simulation at all. Nevertheless, we believe that this example will help to delineate potential shortcomings of and improvements over Large Deviations methods.

4 Simulation Using Reverse-Time Models

4.1 Calculation of Reverse-Time Model

Given a stationary, irreducible, finite-state Markov process defined by the transition probabilities $p_{ij} = P(x(t+1) = j \mid x(t) = i)$, its reverse-time model is defined by the transition probabilities $\tilde{p}_{ji} = P(x(t) = i \mid x(t+1) = j)$ ([6] p 28), where:

$$\tilde{p}_{ji} = \frac{\pi(i)}{\pi(j)} p_{ij} \quad (9)$$

and $\pi(\cdot)$ is the invariant probability of a state. (Some readers may simply recognize (9) as Bayes Rule, given stationarity.) The natural direction of time flow for the reverse-time model is backwards (i.e. with the time index t decrementing), under which it has the same invariant probability as the forwards-time system (whose time index t increases.) In Section 4.3, we shall use the reverse-time model, defined by the \tilde{p}_{ji} , as the basis for defining the probabilities P^* , and thus an alternative to the system defined by an exponential change of measure (5) subject to (7) in Section 3.

4.2 Invariant Probabilities and the Reverse-Time Model

A useful property of the reverse-time model of a stationary finite-state Markov process is that the ratio of the probability of a fixed but arbitrary trajectory in the forwards-time system conditioned on the initial point $x(0) = a$, say, to the probability of the same trajectory in the reverse-time system, conditioned on the final point $x(T) = b$, say, depends only on the end points of the trajectory; i.e. if q is the product of the transition probabilities for a trajectory of length T , starting at a and ending at b , in the forward-time system, and \tilde{q} the product of the transition probabilities of the same trajectory in the reverse-time system, the value of (q/\tilde{q}) for the trajectory

does not depend on the path taken, but is equal to $\frac{\pi(b)}{\pi(a)}$. For all of x_1, x_2, \dots, x_{T-1} and T we have, using (9):

$$\begin{aligned}
q &= P(x(1) = x_1 | x(0) = a)P(x(2) = x_2 | x(1) = x_1) \dots \\
&\quad P(x(T) = b | x(T-1) = x_{T-1}) \\
&= \frac{\pi(b)}{\pi(a)} P(x(0) = a | x(1) = x_1)P(x(1) = x_1 | x(2) = x_2) \dots \\
&\quad P(x(T-1) = x_{T-1} | x(T) = b) \\
&= \frac{\pi(b)}{\pi(a)} \tilde{q}
\end{aligned} \tag{10}$$

In general, in order to find the reverse-time model of a stationary finite-state Markov process, we first find the invariant probabilities of the states. It is easy to show that the recurrence time of n_c is $\frac{1}{\pi(n_c)}$ [7], and hence simulation is redundant with known $\pi(n_c)$. However, other methods for finding the reverse-time model based on optimal control, *not requiring explicit knowledge of the invariant probabilities*, exist for some types of systems [8]. Also, there exists a class of system known as reversible systems for which the construction is trivial [6]. If such methods, or even approximate methods, are used to find the reverse-time model of a certain finite-state Markov process, then the reverse-time model could be of use in finding recurrence times of states. Also, as will be shown in the next section, there are exit problems where the quantity of interest is *not* simply expressible as a recurrence time and for these problems simulation using reverse-time models (whether or not invariant probabilities of the states are known) can be used to provide significant acceleration in the simulation.

In the ALOHA example, if we know the reverse-time model, then we can find exactly the value of the ratio $\frac{\pi(n_c)}{\pi(n_0)}$ by tracing out any one trajectory between the two states n_0 and n_c in the two systems, (and using the one-step transition probabilities). Now, $\pi(n_0)$ is easily found by simulation if the invariant probabilities are not known, (for example, as the reciprocal of the expected recurrence time of n_0 , which is not large if n_0 is stable, and easily obtained by simulation,) and hence $\pi(n_c)$ can be found. The recurrence time of the state n_c is merely $\frac{1}{\pi(n_c)}$.

4.3 Simulation Using the Reverse-Time Model

The problem examined in [3] is to find the expected time for the ALOHA system to become unstable (i.e. have more than n_c users blocked) starting with n_0 users blocked. As was stated in Section 2, if we replace all transitions starting in unstable states (i.e. those greater than n_c) with deterministic transitions to n_0 , then the expected exit time becomes very nearly the recurrence time of n_c .

The second exit-time problem is to find the expected time for the system to become unstable, starting in some set F , not identified with a single state, with the relative probabilities of states within this set the same as those defined by the invariant distribution, i.e. for any state y in F , the probability of starting in state y is $\frac{\pi(y)}{\pi(F)}$. In the example used here, $F = \{x | x < 2n_0\}$. The expected exit time does not directly correspond to the recurrence time of a state, in contrast to the problem examined in [3], and therefore is not trivially found from the invariant probabilities. Simulation using the reverse-time model can be carried out as follows.

Consider the event that the system follows a trajectory with $x(0) \in F$, $x(T) = n_c$, and $x(t) \notin F \cup \{n_c\}$ for $0 < t < T$. Let A_i be the event that we follow this same trajectory, given

$x(0) = i$ for some particular $i \in F$. Then we have

$$A = \cup\{A_1, A_2, A_3, \dots, A_{2n_0-1}\} \quad (11)$$

The reverse-time model is run, starting at n_c , until it hits the set F . The reverse-time probabilities of the trajectory $P^*(A)$ is determined by the “strike rate” in this system, $\mathbf{1}^*(\omega_j)$, and the Radon-Nikodym derivative is given by (10):

$$\left(\frac{dP}{dP^*}\right)(A_i) = \frac{\pi(n_c)}{\pi(i)}, \quad \forall i \in F. \quad (12)$$

Either $\left(\frac{dP}{dP^*}\right)$ for the trajectory is found from (12), or, by noting that with the transition probabilities in the reverse and forward time models known, the Radon-Nikodym derivative is computable as a ratio of products of these transition probabilities. (Hence the invariant probabilities do not have to be used for its evaluation.) Finally, $P(A)$ is estimated from (3).

As for the *fast simulation method* of [3], we will find separately $P(A)$, as defined above, and the recurrence time ($\hat{\tau}_F$) of the starting set F (through simulation of the forward-time model). Let β_R be the estimate of $P(A)$. (β is the estimate used in [3], and is determined with a single initial state n_0 , in contrast to the determination of β_R .) Then, with identical argument to [3], the expected exit time $\hat{\tau}$ (where we begin in F and end in n_c) is given by:

$$\hat{\tau} = \frac{\hat{\tau}_F}{\beta_R} \quad (13)$$

provided that $\hat{\tau}$ is large. The variance of the estimator β_R is found from:

$$\begin{aligned} \hat{\sigma}^2 &= E[(\hat{P}(A) - P(A))^2] \\ &= \frac{1}{N^2} \sum_{i=0}^{2 \times n_0} \left(\frac{dP}{dP^*}\right)_i^2 N_i - P^2(A) \end{aligned} \quad (14)$$

where N is the number of trajectories run, and N_i is the number of these trajectories that ended at state i .

Using the reverse-time model for simulation allows us to examine exactly the trajectories of the original system that exit, while *preserving the relative frequency of these trajectories*.

5 Speed Up Factor of Simulation

Given an event A and an unbiased estimator of A as defined in (2), the normalized standard deviation of the simulation is given by:

$$\bar{\sigma} = \frac{\sqrt{1 - P(A)}}{\sqrt{nP(A)}} \quad (15)$$

where n is the number of tests performed. For example, to have 95 % confidence that the error in $\hat{P}(A)$ is less than 20 % requires $\bar{\sigma} \leq 0.1$.

Now, n can be expressed as $n = t_A/\ell_A$, where t_A is the number of simulation steps run for n tests, and ℓ_A the average number of simulation steps per test.

Therefore, we can write:

$$t_A = \frac{1}{\sigma^2} \frac{\ell_A(1 - P(A))}{P(A)}. \quad (16)$$

In calculating the speed up factor afforded by simulating using the reverse-time model, we are interested in the ratio of the number of simulation steps required in a direct simulation to give a specified variance to that required in the reverse-time model. This is calculated by running a simulation using the reverse-time model. Equation 14 is used to calculate the variance, and the simulation is stopped when this has reached some predefined value. Let $t_{\bar{A}}$ be the number of steps required to yield the required variance in the simulation using the reverse-time model, then the speed up factor s is given by: $s = t_A/t_{\bar{A}}$, where t_A is calculated from (16); the values of σ and $P(A)$ in (16) are those obtained using the reverse-time model, and ℓ_A , as already described, is obtained by simulation (of the forward model).

6 Results

The reverse-time model was used to estimate the probability of overflow (β_R of the uncontrolled ALOHA system for a number of ALOHA systems (parameterized by α and p , see (1)). All simulations were run until the normalized standard deviation was less than 5 %. The speed up over simulation of the original ALOHA model obtained for each system was then averaged over 10 such runs. The results are shown in Table 1. To find the expected time to reach instability from these results, we would simulate the forward system to find $\hat{\tau}_F$, and then apply (13). Table 2 compares the results obtained here with those obtained in [3]. It should be noted that we do not expect β and β_R to take the same values.

Table 1: Speeds Up obtained for a number of different ALOHA systems.

Parameters		Reverse-Time	
α	p	β_R	Speed Up
0.3	0.049	$8.8 \cdot 10^{-4}$	171
0.3	0.039	$7.1 \cdot 10^{-5}$	322
0.3	0.032	$2.8 \cdot 10^{-5}$	482
0.3	0.028	$5.7 \cdot 10^{-6}$	1232
0.3	0.024	$5.9 \cdot 10^{-6}$	2190
0.3	0.019	$2.7 \cdot 10^{-7}$	10643
0.27	0.049	$4.8 \cdot 10^{-6}$	7913
0.27	0.032	$3.5 \cdot 10^{-8}$	58834

The reason that the reverse-time model is able to provide greater acceleration of simulations than the Large Deviations method appears to be twofold. Firstly, because it allows us to start at the rare event in which we are interested, we only simulate exactly those trajectories that end at this event. In the method of [3], many trajectories return to the initial state before reaching the exit state. Secondly, in deriving the Large Deviations results, it is assumed that in the simulation system, we follow trajectories in the same time direction as in the original system, but alter the relative probabilities of the trajectories. The reverse time model selects only those trajectories whose final state is that of interest, but leaves the relative probabilities of these trajectories unchanged.

Table 2: Comparison of speed up factors obtained for a number of different ALOHA systems between reverse-time and [3].

Parameters		Reverse-Time		Cottrell <i>et al.</i>	
α	p	β_R	Speed Up	β	Speed Up
0.3	0.049	$8.8 \cdot 10^{-4}$	171	$1.8 \cdot 10^{-4}$	18
0.3	0.039	$7.1 \cdot 10^{-5}$	322	$6.0 \cdot 10^{-5}$	48
0.3	0.028	$5.7 \cdot 10^{-6}$	2232	$6.0 \cdot 10^{-6}$	403
0.27	0.049	$4.8 \cdot 10^{-6}$	7913	$4.3 \cdot 10^{-6}$	744

7 Conclusion

It has been shown that for a number of uncontrolled ALOHA systems, a much faster simulation can be performed using the reverse-time model than is possible using the asymptotically optimal (but constrained) techniques of [3]. For a general stationary finite-state Markov process, the only presently available method of finding the reverse-time model involves first finding the invariant probabilities of the states of the process, which immediately gives us the recurrence times of states, and therefore makes simulation unnecessary for certain exit time problems. However, this new speed up technique can potentially be applied to situations where the expected exit time cannot be approximated as a recurrence time.

In [8], it is shown that there is a connection between Large Deviations and the construction of reverse-time models for diffusion processes. Such a connection should also be explored for finite-state Markov processes.

References

- [1] M. I. Freidlin and A. D. Wentzell, *Random Perturbations of Dynamical Systems*. New York: Springer-Verlag, 1984.
- [2] J. Zabczyk, "Exit problem and control theory," *Systems and Control Letters*, vol. 6, pp. 165–172, August 1985.
- [3] M. Cottrell, J. C. Fort, and G. Malgouyres, "Large deviations and rare events in the study of stochastic algorithms," *I.E.E.E. Trans. Automatic Control*, vol. AC-28, pp. 907 – 918, September 1983.
- [4] S. Parekh and J. Walrand, "A quick simulation of excessive backlogs in networks of queues," *I.E.E.E. Trans. Automatic Control*, vol. 34, pp. 54 – 66, January 1989.
- [5] M. R. Frater and B. D. O. Anderson, "Estimation of rare event statistics in data communications networks," in *3rd Australian Teletraffic Research Seminar*, November 1988.
- [6] F. P. Kelly, *Reversibility and Stochastic Processes*. New York: John Wiley and Sons, 1979.
- [7] D. R. Cox and H. D. Miller, *The Theory of Stochastic Processes*. London: Chapman and Hall, 1965.
- [8] M. R. Frater, R. A. Kennedy, and B. D. O. Anderson, "Reverse-time modeling, optimal control and large deviations," *Systems and Control Letters*, vol. 12, pp. 351 – 6, May 1989.