

REPORT DOCUMENTATION PAGE					Form Approved OMB No. 0704-0188	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to the Department of Defense, Executive Service Directorate (0704-0188). Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.						
1. REPORT DATE (DD-MM-YYYY) 03-10-2012		2. REPORT TYPE Final Technical Report			3. DATES COVERED (From - To) 7/1/09-11/30/11	
4. TITLE AND SUBTITLE Computational Methods for Stochastic Networks				5a. CONTRACT NUMBER		
				5b. GRANT NUMBER AFOSR FA9550-09-1-0378		
				5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) Paul Dupuis and Harold Kushner				5d. PROJECT NUMBER		
				5e. TASK NUMBER		
				5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Division of Applied Mathematics/Brown University 182 George St. Providence, RI 02912				8. PERFORMING ORGANIZATION REPORT NUMBER 3		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Office of Scientific Research 875 North Randolph St., Suite 325, Room 3112 Arlington, VA 22204				10. SPONSOR/MONITOR'S ACRONYM(S)		
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL - OSR - VA - TR - 2012 - 1006		
12. DISTRIBUTION/AVAILABILITY STATEMENT DISTRIBUTION A: APPROVED FOR PUBLIC RELEASE						
13. SUPPLEMENTARY NOTES						
14. ABSTRACT This final report summarizes accomplishments in two areas of computational methods for stochastic systems: 1) Large deviation and computational methods for rare event problems and 2) Numerical methods for controlled stochastic delay systems.						
15. SUBJECT TERMS						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 11	19a. NAME OF RESPONSIBLE PERSON Paul Dupuis	
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (Include area code) 401-863-3258	

Final report on AFOSR FA9550-09-1-0378 entitled
Computational Methods for Stochastic Networks

Period of award
1 July, 2009 – 30 November, 2011

Paul Dupuis
Harold Kushner (co-PI)
Division of Applied Mathematics
Brown University, Box F
Providence, RI 02912

Phone: (401) 863 3238
Fax: (401) 863 1355
Email: dupuis@dam.brown.edu

10 March, 2012

This final report summarizes accomplishments in two areas of computational methods for stochastic systems.

1 Large deviations and computational methods for rare event problems

1.1 Problem formulation and issues

There is significant interest in the application of uncertainty quantification to problems with small probabilities (rare events). For example, it is a research focus for the SAMSI/Sandia Summer School on Uncertainty Quantification in 2011, and has been the focus at other workshops, including workshops at SAMSI and ICERM in 2012. Though these probabilities may be small, they are often critical measures of system performance and one needs reasonably reliable numerical methods. Unfortunately, standard numerical schemes are not at all reliable for problems with rare events.

This part of the research project was concerned with developing efficient Monte Carlo algorithms for rare event simulation and the associated large deviations theory. The analysis and design of such schemes require first an understanding of the qualitative properties of the rare events, and this usually means that a large deviations analysis, which identifies via a variational characterization the quantitative and qualitative properties of the rare events. Once this is available, various methods can be analyzed.

For the problem of estimating the probability of a single rare event there are two methods currently in use. One is based on simulating according to a different distribution and then correcting for any induced bias via the likelihood ratio (importance sampling). The key question here is how to select the new sampling distribution. The second method simulates a branching process, i.e., collection of particles that can split according to certain rules to form new particles, each of which behaves like the original particle or process. The splitting rules are designed to make the rare event likely for at least one of the descendent particles, and the estimator is the ratio of number of particles for which the rare outcome is observed to the total number of descendants. The key question here is what should trigger a split and, given that a split occurs, the number of descendants. Most of the literature on these methods features schemes based on heuristic design, with little or no analysis. However, the design problem with both methods is subtle, and reasonable looking schemes can perform quite poorly. Indeed, simulations based on improperly designed schemes could be highly misleading.

A second class of problems considers the numerical approximation of the

invariant distribution for stochastic systems with multiple metastable states using the occupation measure of a related Markov process. Moving from one metastable state to another is a rare event, and its treatment is the key question in the design of efficient Monte Carlo schemes. There are many ad hoc algorithms available. However, these algorithms do not always work well and have to be applied with some care.

1.2 Results and key findings

The focus of this work has been in the development of the large deviation theory for infinite dimensional systems, and construction and analysis of splitting type schemes for finite dimensional problems.

The papers [1,2] continue a long term project in the development of large deviations theory for infinite dimensional systems. The starting point for much mathematical modeling in fluid dynamics, geophysics, climate science, neurophysics, chemical reaction diffusion systems, and many other areas is often a partial differential equation (PDE) based on fundamental physical constitutive laws. Such PDE are often inadequate—they do not capture stochastic fluctuations and variability resulting from noise processes, randomly varying coefficients, measurement quantization, etc. Stochastic PDEs (SPDEs) are frequently proposed as improved models that systematically account for randomness. A question of fundamental interest is then: How do predictions based on the stochastic model differ from predictions based on the corresponding deterministic PDE? When the noise fluctuations are “small,” a basic mathematical approach to quantify probabilities of divergence between such predictions is through the theory of large deviations. Though well developed for finite dimensional systems, technical difficulties have severely limited the use of large deviations theory for these infinite dimensional problems. Our goal has been to develop techniques capable of handling the wide range of problems of interest.

The paper [2] develops a variational representation for nonnegative functionals of infinite dimensional Brownian motion and Poisson random measures. Using also techniques from the theory of stochastic control and weak convergence, the study of large deviations is reduced to the analysis of basic qualitative properties of controlled analogues of the deterministic PDEs: existence and uniqueness of solutions, stability under bounded perturbations. It gives a technically much simpler and unified framework for treating a broad family of infinite dimensional jump-diffusion models. The approach (also developed in prior publications) has become the method of choice for studying small noise asymptotics for SPDEs, and has been adopted by other

researchers to study many different sorts of problems. With the last several years more than 20 publications have appeared which use the basic theory we have developed to analyze a range of problems in SPDE.

The papers [3,7] are concerned with branching type schemes for rare event estimation. Paper [7] uses ideas from ordinary splitting for the analysis of what are called “counting” problems. These problems are discrete state analogues of the problem of volume estimation. The problem is to estimate the size of a very large but finite population of discrete objects. By embedding the objects into a larger space with a known population, one can interpret the relative size of the populations as a probability, and thereby apply Monte Carlo methods to the estimation problem. In addition to the use of ordinary splitting to obtain a good estimate [7], introduces the novel use of what is known as capture-recapture to bootstrap to much more accurate estimates. Though theoretical analyses have yet to be carried out, the method worked very well for certain very challenging benchmark problems in counting.

An alternative to standard splitting uses an *interacting particle systems*. With this approach a fixed number of particles N is maintained via the following mechanism. Suppose one wants to estimate the probability of hitting a rare set B before some reference event A . The splitting mechanism is defined by a collection of thresholds, which should in some sense help ensure that at least some particles reach B . Starting with N particles in a given threshold, one simulates all N until either they have reached the next threshold or been absorbed into A . Given the locations of the particles that made it to the next threshold, one samples from these according to the uniform distribution to make up the full complement of N particles. The product of the fraction of particles that make it to each of the successive thresholds is then an unbiased estimator.

One possible advantage of this method is that the a priori bound on the number of particles provides some guarantee that the computation will not get out of hand. On the other hand, there has been essentially no analysis of performance of the scheme as the event of interest becomes rare, and in particular there are no known necessary and sufficient conditions for good performance. The reason is of course the interaction, which makes the application of standard techniques such as large deviation methods very difficult. The key time scale in the problem is not that of the underlying process, but rather a time scale that measures progression with respect to threshold levels. These features complicate the analysis greatly.

The paper [3] provides the first rigorous analysis of the performance of this class of algorithms in the small probability limit. Owing to the com-

plexities of the algorithm, the analysis is limited to dimension one. However, the results support some of the claims that have been made concerning these algorithms (but based only on numerical evidence), and in particular suggest that, at least for one dimensional problems, they are less sensitive to the details of the underlying distributions than competing schemes.

The paper [6] considers the problem of approximating stationary distributions of a Markov chain by simulation. Our initial goal was to use large deviation ideas to choose design parameters in an existing scheme known as parallel tempering. Parallel tempering (also known as replica exchange sampling) is a standard method for simulating complex systems, and is used in many commercial software packages. In this algorithm simulations are conducted in parallel for a family of Markov chains indexed by a “temperature” parameter, and the key improvement over standard Monte Carlo is a swap mechanism that exchanges configurations between these parallel simulations at a given rate. The mechanism is designed to allow the low temperature system to escape from deep local energy minima where it might otherwise be trapped, via those swaps with the higher temperature components. Based on large deviation theory, we have argued that the rate of convergence of the empirical measure is a monotone increasing function of the swap rate. This suggests that one should raise the swap frequency in order to improve efficiency, but this is eventually counter-productive since eventually most of the computational effort is directed towards swapping and little towards moving the process dynamics. However, it turns out that one can construct a simulation scheme that is equivalent to the limit of the parallel tempering schemes in the sense of distributions, but which involves no swapping at all. With this scheme, which we call *infinite swapping* (INS), the effect of the swapping is captured by a collection of weights that influence both the dynamics and the empirical measure.

While the infinite swapping scheme optimizes the convergence rate as described above, it has practical limitations in implementation due to the complexity of the weights when the number of temperatures is large (>7). Complex problems often involve scores of temperatures, and so it was critical to overcome this limitation. We have developed an approximation to the full infinite swapping which is based on alternating between *partial infinite swapping* (PINS) algorithms, which can be shown to approximate (theoretically and practically) the INS scheme. The mathematical theory for the INS and PINS is developed in [6]. Numerical studies on fairly complex Lennard-Jones systems (very challenging benchmark problems from chemistry) have been conducted. Improvements of three orders of magnitude in performance over conventional parallel tempering were observed at an increased compu-

tational cost of 5-15%.

2 Numerical methods for controlled stochastic delay systems

2.1 Problem formulation and issues

The Markov chain approximation class of algorithms [A] are effective methods for the numerical approximation of optimal controls and values for general continuous-time nonlinear stochastic systems. First, we approximate the process by a controlled finite-state Markov chain that satisfies certain minimal local consistency properties, and then we solve the Bellman equation for the approximation. This gives the approximating costs and controls. Finally, we prove convergence as the approximation parameters go to zero. These methods were extended to controlled general nonlinear diffusion models with delays in the dynamics in [B,C] and in the monograph [D]. The proofs of convergence are probabilistic, being based on the theory of weak convergence of random processes. They do not depend on the analytical properties of the Bellman equation for the original model, and converge under virtually the weakest possible conditions. In the absence of delays in the dynamics, the probabilistic basis of the proofs is a key to the generality, robustness, and usefulness of the methods. With delays, it is absolutely essential, since then virtually nothing is known about the properties of the Bellman equation for the original model.

The numerical problem is particularly difficult if the control and/or reflection terms in the dynamical model are delayed, since then the memory needs with naive approximations can be enormous, due to the fact that approximations to the memory segments of these processes lead to very high-dimensional numerical models. For such cases, we reformulated the problem in terms of a stochastic “wave” or “transportation” equation [D], whose numerical solution requires much less memory, and which yields the optimal costs and controls. The emphasis in [D] was on the theoretical foundations. Promising classes of algorithms were developed and convergence theorems were proved. However, using these “ideal” algorithms in concrete applications requires many adaptations. There are non-obvious steps in the derivation of the best Markov chain approximation via the “implicit” approximation method, and in the derivation of the Bellman equation, partly due to the need to reduce computation as much as possible and to the fact that the transformations that are used in the transportation equation reformulation introduce non-physical quantities that must be approximated by the observed physical data. Such details as well as data illustrating the use-

fulness of the methods was left to recent works [7,8,9], which provide guides to the coding and applications, and give data illustrating the improvement in performance that is possible when the delays are taken into account., without ad-hoc approximations.

In [10] a main concern was a class of problems that contains simple models of internet regulation, where the physical model turns out to be a reflected diffusion, and both the control and reflection terms are delayed in the dynamics. The variables were the controlled source rates and the queue level, where the queue levels and overflows (lost packets) are observed at the router/queue and the controls are determined there. There are communication delays in getting the information to the sources and the data from the source to the router. This served as a concrete illustration of the methods, and can be taken to be illustrative of the benefits to be gained by the use of numerical methods. It was seen that controls that take the past (over the delay interval) into account can yield much improved performance. The controls to be computed are robust in that they perform well when the system parameters change. These problems would be intractable by any other currently proposed method; e.g., via the use of time discretizations of the memory segments over the delay interval. The source rates and queue levels for the uncontrolled system could oscillate wildly. The controlled system was a vast improvement in terms of the overflow, source rates, and queue levels.

2.2 Results and key findings

Numerical methods for controlled delay systems, more general problems. In [11], we continued the theoretical and algorithmic development for the transportation equation approach for important classes of models not covered by the previous works, and for which their methods of proof are inadequate. One of the motivating examples concerned sources that had files that they wished to be admitted to a network. The file creation process was Poisson. Requests for admission were sent to a queue/router, and received there after a delay. The admission process was controlled, with permission for admission determined at the queue, and received back at the source after another delay. Overflows, total delays and non admissions were penalized.

The admissions control model is covered by the following more general

form, where $\bar{\theta}$ is the maximum delay:

$$\begin{aligned}
dx(t) = & c(x(t), u(t))dt + \sigma(x(t))dw(t) \\
& + dt \int_{-\bar{\theta}}^0 b(x(t+\theta), u(t+\theta), \theta) d\mu_a(\theta) \\
& + \int_{-\bar{\theta}}^0 \int_{\Gamma} g(x((t+\theta)-), u((t+\theta)-), \rho, \theta) N(dt(t+\theta), d\rho) \mu_b(d\theta) \\
& + \int_{-\bar{\theta}}^0 p(x(t+\theta)-, \theta) d_t y(t+\theta) d\mu_c(\theta) + dz(t).
\end{aligned}$$

$N(\cdot)$ is a Poisson random measure and $u(\cdot)$ is the control. The $\mu_i(\cdot)$ are measures over the delay interval. The reflection term $z(\cdot)$ serves to keep the path in a compact polyhedral constraint set G , and $y(\cdot)$ is the vector of its components that are due to reflection from the various faces of G . As an alternative, one could stop the process when it first reaches the boundary of G and add an associated penalty, but the reflecting boundary case models applications in communications systems that are of interest, and it is the more difficult case.

Delayed and controlled jump terms were not previously allowed, and presented particular problems. Further complicating matters, in our model the jump term can be controlled, which requires non-standard methods such as the so-called relaxed Poisson measure approach [A], owing to the problems of characterizing the limit of convergent sequences of approximations to the controlled jump term. Another shortcoming of the prior work is that it did not allow point delays in the reflection terms (and certainly not in the jump terms); e.g., where the delay is concentrated at, say a single point $-\bar{\theta}$. The term $d\mu_c(\theta)$ was replaced by dt , so the delay was “spread out” by a suitable choice of the continuous function $p(\cdot)$. In applications, such as those considered in the data-oriented papers [7,9], a point delay in the original model was replaced by a delay that is slightly distributed. While this did not compromise the usefulness of the data, especially since delays are often distributed, it was still important to develop methods that can handle models with such point delays. The proofs for the simpler case cannot be simply carried over to the current model, and required substantial modifications.

More complicated nonlinear control stochastic systems with delays. Modeling and approximation. There are many important classes of nonlinear control stochastic systems with delays whose analysis has not been previously considered in anyway in the literature. For example systems with dynamics containing terms of the form $b(x(t+\theta_1), x(t+\theta_2), u(t+$

$\theta_3), u(t + \theta_4))$, etc, where $\theta_i < 0$ and they might take different values; or with analogous terms involving reflection and/or jump terms. There are applications in communications systems that originally motivated the work. For such models the issues of admissible controls, admissible approximating controls, or numerical procedures, had not been previously investigated. These issues were addressed in [12,13], which showed that natural but quite non conventional definitions of admissible controls must be used, and constructed natural approximations. The definitions are important since the set of admissible controls and solutions must be closed under weak convergence, and the approximations to the $u_i(t)$ are all selected at time t , but applied at later different times $t + |\theta_i|$. The transportation equation format was extended to cover a large class of such systems as well.

3 People supported by the award

- Harold Kushner, Professor, Brown University
- Paul Dupuis, Professor, Brown University
- Yufei Liu, Graduate Student, Brown University
- Dane Johnson, Graduate Student, Brown University
- Yi Cai, Graduate Student, Brown University
- Kenny Chowdhary, Graduate Student, Brown University
- Markus Fischer, Postdoc, Brown University. Current position: University of Padua

4 Preprints and publications

1. A. Budhiraja, P. Dupuis and M. Fischer. Large deviation properties of weakly interacting processes via weak convergence methods, to appear in *Annals of Probability*.
2. A. Budhiraja, P. Dupuis and V. Maroulas. Variational representations for continuous time processes, *Annales de l'Institut Henri Poincaré*, **47**, (2011), 725–747.
3. Y. Cai and P. Dupuis. Analysis of an interacting particle method for rare event estimation, submitted.

4. K. Chowdhary and P. Dupuis. Distinguishing and integrating aleatoric and epistemic variation in uncertainty quantification, to appear in *ESAIM: Mathematical Modelling and Numerical Analysis*.
5. P. Dupuis and M. Fischer. On the construction of Lyapunov functions for nonlinear Markov processes via relative entropy, submitted.
6. P. Dupuis, Y. Liu, N. Plattner and J.D. Doll. On the infinite swapping limit for parallel tempering, to appear in *SIAM J. Multiscale Modeling and Simulation*.
7. P. Dupuis, B. Kaynar, A. Ridder, R. Rubinstein and R. Vaisman. Counting with combined splitting and capture-recapture methods, to appear in *Stochastic Models*.
8. H. J. Kushner. Numerical methods for optimal controls for nonlinear stochastic systems with delays: Algorithms and data. In *Proceedings of the 8th IFAC conference on delay systems, Sinaia, Romania, Sept, 2009*. IFAC, 2009.
9. H.J. Kushner. Numerical algorithms for optimal controls for nonlinear stochastic systems with delays. *IEEE Trans on Automatic Control*, 55:2170–2176, 2010.
10. H.J. Kushner. Numerical methods for optimal controls for nonlinear stochastic systems with delays, and applications to internet regulation. In *Proceedings of the 9th IFAC conference on delay systems, Prague, June, 2010*. IFAC, 2010.
11. H.J. Kushner. Numerical methods for the optimization of nonlinear stochastic delay systems, and an application to internet regulation. In *Proceedings of the 49th Conference on Decision and Control*. IEEE Press, New York, 2010.
12. H.J. Kushner. Numerical methods for controls for nonlinear stochastic systems with delays and jumps: Applications to admission control. *Stochastics An International Journal of Probability and Statistics*, 83:277–310, 2011.
13. H.J. Kushner. Controlled nonlinear stochastic delay equations: Part I: Modeling and approximations. To appear in *Applied Math and Optimization*.

14. H.J. Kushner. Controlled nonlinear stochastic delay equations: Part II: Approximations and pipe-flow representations. To appear in *Applied Math and Optimization*.

5 Other references

- A. H.J. Kushner and P. Dupuis. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer-Verlag, Berlin and New York, 1992. Second edition, 2001.
- B. H.J. Kushner. Numerical approximations for nonlinear stochastic systems with delays. *Stochastics and Stochastic Reports*, 77:211–240, 2005.
- C. H.J. Kushner. Numerical approximations for stochastic systems with delays in the state and control. *Stochastics and Stochastics Reports*, 78:343–376, 2006.
- D. H.J. Kushner. *Numerical Methods for Controlled Stochastic Delay Systems*. Birkhäuser, Boston, 2008.