Proceedings of the 1999 Winter Simulation Conference
P. A. Farrington, H. B. Nembhard, D. T. Sturrock, and G. W. Evans, eds.

THE PHANTOM SPA METHOD: AN INVENTORY PROBLEM REVISITED

Felisa J. Vázquez-Abad
Manuel Cepeda-Jüneman

Department of Computer Science and Operations Research
University of Montreal
Montreal, H3C 3J7, CANADA

ABSTRACT

It is widely accepted today that the Infinitesimal Perturbation Analysis (IPA) method for estimating sensitivities is the preferred method, when it is applicable. The major problem with IPA is handling certain kinds of discontinuities, such as thresholds. The Smoothed Perturbation Analysis (SPA) method was conceived applying a conditional expectation to a dynamic system, similar to the Filtered Monte Carlo Simulation. Conditioning smoothes out the discontinuities and then IPA can be applied to the conditional estimator. Since this alternative estimator has been partly integrated through the conditioning, some knowledge about the underlying distribution is required. When this is not available, SPA estimators require additional estimation. Traditionally, this has been implemented via off-line simulations that produce independent replications of a difference process. We propose here to bypass this operation by using parallel phantom systems: replicas of the original system that are conditional to the critical events of interest yet use common random numbers instead of independent replications. We show how the efficiency can dramatically improve from the gain in correlation (variance reduction) as well as the gain in computational effort (random variables are generated once and used for all parallel phantoms).

1 INTRODUCTION

When estimating derivatives of averages of discrete event driven systems, the preferred method is the Infinitesimal Perturbation Analysis (IPA) method. A simple inventory model is considered to introduce the problem of estimating the derivative of a stationary average with respect to a threshold parameter. As is typical in these cases, there are path discontinuities when the thresholds are modified infinitesimally. Therefore the stochastic derivative, which considers the derivative of the pathwise cost function with respect to the parameter (for a fixed realization of the trajectory) is biased: we cannot interchange derivative and expectation. IPA uses the stochastic derivative as an estimator of the derivative of the expectation. The Smoothed Perturbation Analysis (SPA) method was conceived applying a conditional expectation to a dynamic system, similar to the Filtered Monte Carlo Simulation. Conditioning smoothes out the discontinuities and then IPA can be applied to the conditional estimator.

The conditioning requires knowledge of the underlying distribution as well as simulations of the so-called difference process. In most cases, as illustrated in the model example, one has to estimate an average probability as well as an expectation of the difference process. These two steps are usually performed separately, simulating the difference process off-line.

We propose the Phantom SPA estimation that uses parallel phantom systems in order to estimate the contribution of the difference process using common random variables with the nominal process. Phantom systems represent replicas of the original system that are conditional to the critical events of interest, yet use common random numbers instead of independent replications. This method carries out the two steps of the estimation simultaneously over one single path.

Section 2 introduces the model, taken from Bashyan and Fu (1991). Section 2 introduces SPA estimation, using an approach which parallels the framework of Rare Perturbation Analysis (RPA) for sensitivities with respect to threshold parameters. Using this approach, the usual SPA estimators of Bashyan and Fu (1991) are inetrpreted in terms of the limits of one-sided finite differences. The two-sided version of the estimator is presented. Section 3 introduces the parallel computation for the difference processes, and Section 4 presents simulation results showing how the Phantom SPA method can outperform conventional SPA with off-line simulations.

2 THE INVENTORY MODEL

We consider here the standard periodic review inventory model of Bashyan and Fu (1991). The optimal policy is parametrized by two threshold parameters and is commonly known as (s, S) policy, described as follows. Let X_n denote the inventory level at the beginning of the n-th period, during which a total demand $D_n \geq 0$ occurs. If $X_n - D_n \geq s$ then there is no reordering at the end of the period, otherwise an order is placed to fill up the inventory to level S. That is:

$$X_{n+1} = \begin{cases} X_n - D_n & \text{if } X_n - D_n \geq s \\ S & \text{otherwise} \end{cases} \tag{1}$$

where $\{D_n\}$ is a sequence of consecutive i.i.d. demands. For discrete inventory models, X_n counts the number of items in stock, and D_n has a discrete distribution. The case considered in Bashyan and Fu (1991) is that of a continuous level of inventory, where consecutive demands are assumed to have a bounded density $g(\cdot)$, on $\mathbb{R}^+$. Let the underlying probability space be denoted by (Ω, P) and let $\mathcal{F}_n$ be the σ-algebra generated by $\{D_1, \ldots, D_n\}$. It is not difficult to see that the inventory level X_n is measurable w.r.t. $\mathcal{F}_n$ and that $\{\mathcal{F}_n, n \geq 0\}$ is the natural filtration of the process $\{X_n\}$.

The inventory cost during period n is calculated as h dollars per unit of inventory at the end of the period, if this quantity is non-negative. Backlog is assumed, so that inventory levels can become negative, modeling situations when the demand is fulfilled at a later time, but we assume that in these cases a reduction in price has to be made. The penalty cost is p dollars per unit of unsatisfied demand. Therefore the cost is a function of the Markov process $\{X_n\}$ and satisfies:

$$C_n = h(X_n - D_n)\mathbf{1}_{\{X_n > D_n\}} - p(X_n - D_n)\mathbf{1}_{\{X_n \leq D_n\}} + K\, \mathbf{1}_{\{X_n - D_n \leq s\}}.$$

The process $\{X_n\}$ above is a Markov chain in a bounded interval (s, S) and it possesses an ergodic stationary measure. The average long term cost per period is defined by the stationary expected value of C_n, or:

$$C(s, S) = \lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} E(C_n).$$

Figure 1 shows a typical trajectory of the process (continuously interpolated in time).

The actual optimal values of s and S are not available in closed form for general demand distributions. Finding the optimal values of the thresholds s and S is harder for the continuous than for the discrete model. Bashyan and Fu propose to implement a gradient search method to find the optimal values by simulation, which leads to the problem of estimating the derivatives of the cost function with respect to s and S.

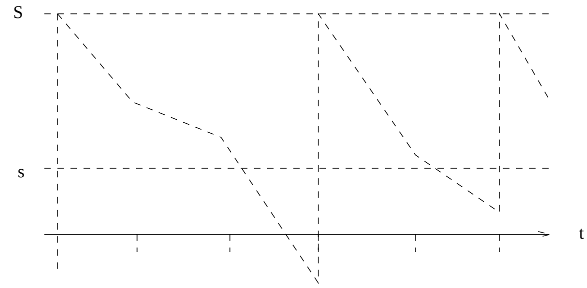

Figure 1: A trajectory of $\{X_n\}$.

In order to estimate the cost and its derivatives, a finite horizon simulation of length N is used, assuming that N is "large enough", thus approximating

$$C(s, S) \approx C_N(s, S) = \frac{1}{N} \sum_{n=1}^{N} C_n.$$

Alternatively, regeneration of the inventory level process at the ordering periods can be used to apply the Renewal Theorem (see Taylor and Karlin, 1994) and compute $C(s, S)$.

It should be clear from the dynamics of the process that estimating the derivative with respect to S yields indirectly the derivative w.r.t. s. To see this, remark first that if we translate our coordinate system by shifting the zero level, this transformation does not alter at all the evolution of the process, which will be identical path by path, except for the relative location of what we call the zero level. Therefore the process only depends on the relative difference $S - s$. However, the cost incurred does depend on the absolute location of the zero level. As explained in Bashyan and Fu (1991), the pathwise derivatives contain a conditional term (Called the SPA term) that accounts for the effect in the iventory levels of an infinitesimal change in $S - s$, plus a stochastic derivative of the cost function (called the IPA term) which contributes only to the derivative w.r.t. S. For a more detailed description see Vázquez-Abad and Zubieta (1999).

3 THE SPA DERIVATIVES

3.1 One-Sided SPA

We shall first briefly reproduce the SPA derivative of Bashyan and Fu (1991). Consider the *perturbed* inventory process when S is replaced by $S + \Delta$, $\Delta > 0$. The same sequence of demands D_n is used for both *nominal* and perturbed paths, as shown in Figure 2. This is used to evaluate the right-sided finite difference:

$$D_{\Delta}^{+}(S) = \frac{1}{N}\sum_{n=1}^{N} E\left[\frac{C_n(S+\Delta) - C_n(S)}{\Delta}\right]$$

Recall that $\{\mathcal{F}_n, n \geq 0\}$ denotes the natural filtration of the process $\{X_n\}$ on the underlying probability space $(\Omega, \mathcal{F}, P)$. The *critical sets* $\Omega_i^+ \in \mathcal{F}_i$ are defined by:

$$\Omega_i^+ = \{\omega : D_i > Z_i\}, \text{ with } Z_n = X_n - s$$

where $Z_n \geq 0$ a.s., from the ordering policy.

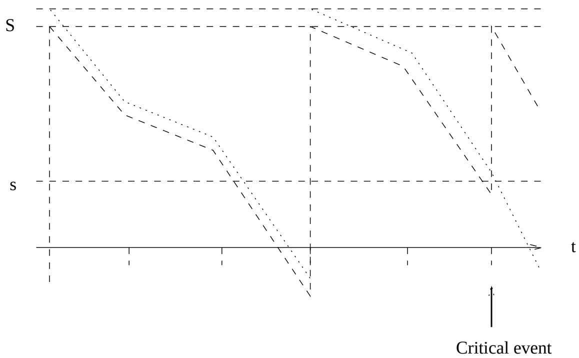

Figure 2: Perturbed and nominal paths.

A difference in the ordering decision can take place for the first time at the end of period i for the nominal and perturbed paths only if an order is placed at the end of the nominal period but not at the end of the period in the perturbed path. In other words, if Ω_i denotes the event that the first order change occurs at the end of period i, then:

$$P[\Omega_i | \Omega_i^+] = P[D_i - Z_i \leq \Delta | \Omega_i^+] = \frac{G(Z_i + \Delta) - G(Z_i)}{1 - G(Z_i)} \qquad (2)$$

while, if $D_i < Z_i$ then with certainty, there will be no order placed at the end of the period, both for the nominal and the perturbed paths, which implies that $P[\Omega_i | (\Omega_i^+)^c] = 0$ and A^c denotes the complement of set A. Let $\Omega^* \subset \Omega$ denote the set of paths where no ordering decisions change under the prescribed perturbation. Since $P[\Omega_i] = \mathcal{O}(\Delta)$, then $P[\Omega^*] = 1 - \mathcal{O}(\Delta)$, and:

$$D_{\Delta}^{+}(S) = E\left[\frac{(h\Delta)N^+ - (p\Delta)N^-}{\Delta N} P[\Omega^*]\right] + \sum_{n=1}^{N} \frac{1}{N}\sum_{i=n}^{N} E\left[\frac{C_i(S+\Delta) - C_i(S)}{\Delta} \;\Big|\; \Omega_n\right] P[\Omega_n]\Bigg)$$

where N^+ denotes the number of periods where the end level of inventory $X_n - D_n$ is positive and $N^- = N - N^+$. In the limit as $\Delta \to 0$, Bashyan and Fu (1991) obtain:

$$\frac{\partial}{\partial S} C_N(s, S) = E[\frac{hN^+ - pN^-}{N}] + \lim_{\Delta \to 0} \frac{1}{N} \sum_{n=1}^{N} \left(\frac{E[\Delta C_n^+ \mid \Omega_n]}{\Delta} P[\Omega_n]\right) \qquad (3)$$

where $\Delta C_n^+ = \sum_{i=n}^{N} [C_i(S+\Delta) - C_i(S)]$. The first term in (3) is called the IPA term and it is the stochastic derivative of the cost process. The second term is called the SPA term and it is the one that considers the effects of sample path discontinuities.

Finally, in the SPA term above, the conditional expectation of the difference process $\{C_i(S+\Delta) - C_i(S), i \geq n\}$ is independent of i, n and Δ: indeed, in the limit as $\Delta \to 0$, both processes coincide before time n, by definition of Ω_n. Next, at period n the nominal process starts by ordering (end level is within $s + \Delta$) while the perturbed process starts at level s. Given this initial condition, which is obviously independent of n, and using the fact that demands are i.i.d., it follows that the difference process has a distribution which is independent of $\mathcal{F}_n$. Use now regeneration of the difference process (see Bashyan and Fu, 1991) to obtain the limit expectation (as $N \to \infty$):

$$E[\Delta C_n^+] = E\left\{\sum_{i=n}^{\phi(n)} [C_i^+(S) - C_i(S)]\right\} \qquad (4)$$

with $X_n^+ = S, X_n = s$ as the initial inventory levels. The random index $\phi(n)$ denotes the first time when both processes reorder: from this time onwards the two evolve identically and their difference vanishes. Since $E[\Delta C_n^+ | \Omega_n]$ is independent of n and $\mathcal{F}_n$, we can drop the index n and use the notation $E[\Delta C^+]$ to obtain the SPA term of (3):

$$\left(\frac{\partial C(s, S)}{\partial S}\right)_{SPA}^{+} = E[\Delta C^+] \lim_{\Delta \to 0} \frac{1}{N} \sum_{n=1}^{N} \frac{P[\Omega_n]}{\Delta}. \qquad (5)$$

Using (2), the expression of Bashyan and Fu (1991) can be obtained:

$$\left(\frac{\partial C(s, S)}{\partial S}\right)_{SPA}^{+} = E[\Delta C^+] \times E \quad \frac{1}{N} \sum_{n=1}^{N} \frac{g(Z_n)}{1 - G(Z_n)} \mathbf{1}_{\{D_n > Z_n\}}\Bigg).$$

Analogously, their expression for the left-sided derivative can be obtained:

$$\left(\frac{\partial C(s,S)}{\partial S}\right)^{-}_{SPA} = E[\Delta C^{-}] \times E\left(\frac{1}{N}\sum_{n=1}^{N}\frac{g(Z_n)}{G(Z_n)}\mathbf{1}_{\{D_n<Z_n\}}\right),$$

where now $\Omega_n^- = \{\omega : D_n < Z_n\}$ gives the conditioning set for potential order changes. The difference process in this case would be $\{C_i(S) - C_i(S-\Delta), i \geq n\}$, which is now conditioned to Ω_n^-. It is not hard to see that the initial corresponding inventory levels are again $X_n = S$, $X_n^- = s$, so that $E[\Delta C^+] = E[\Delta C^-]$.

Remark: The independence of the limit difference process and the history $\mathcal{F}_n$ allows us to separate the estimation of the SPA term in (3) as the product of two estimators. This follows from the fact that the difference processes only differ in their initial states. Most problems where the control parameters are thresholds will present this property, making them suitable for our phantom SPA method, as explained shortly.

3.2 Two-Sided SPA

Our formulation of the SPA estimation in (3) above follows closely the RPA framework in Vázquez-Abad and Davis (1995) that generalizes the original formulation of the Phantom RPA method of Brémaud and Vázquez-Abad (1992). The difference between the Phantom SPA and the formulation in Vázquez-Abad and Davis (1995) for threshold parameters is that the decision of whether to order or not depends on the state X_n of the process, while the usual RPA framework assumes independence of the decision variables with respect to the process $\{X_n\}$. RPA estimators can be implemented using one-sided formulas as in Bashyan and Fu (1991) or two-sided formulas.

Applying a two-sided formula to this problem, consider now the finite difference:

$$D_\Delta(S) = \frac{1}{N}\sum_{n=1}^{N} E\left[\frac{C_n(S+\Delta/2) - C_n(S-\Delta/2)}{\Delta}\right].$$

While (5) continues to hold true and the limit difference process is the same as before, now the potential change in the ordering decision can happen at every period, regardless of whether D_n is larger or smaller than Z_n. It is no longer necessary to condition on $\Omega_i^{\pm}$, which means that the contribution of $P[\Omega_i]$ can be estimated by simply observing X_i at the beginning of the period, before generating the demand. In this case we obtain:

$$\begin{aligned} P[\Omega_i] &= P\left[Z_i - \frac{\Delta}{2} \leq D_i < Z_i + \frac{\Delta}{2}\right] \\ &= G(Z_i + \tfrac{\Delta}{2}) - G(Z_i - \tfrac{\Delta}{2}), \end{aligned}$$

and in the limit, we have:

$$\left(\frac{\partial C(s,S)}{\partial S}\right)_{SPA} = E[\Delta C]\, E\left(\frac{1}{N}\sum_{n=1}^{N} g(Z_n)\right). \qquad (6)$$

The efficiency of an estimator is the inverse of the mean square error multiplied by the mean CPU time. The obvious advantage of this estimator over either of the one-sided ones proposed by Bashyan and Fu (1991) is the improvement in efficiency. This improvement results from three facts. First, the full utilization of all samples yields more observations and thus more precision in the estimation.

Second, dividing by the random observations $1-G(Z_i)$ and $G(Z_i)$ may introduce yet more variability in the estimation. Given the correlation between the numerator and denominator in the one-sided formulas, it is possibe *a priori* that the variance may even be reduced, but our experiments confirm a variance reduction when using the two-sided estimator.

Finally, the two-sided version avoids evaluating the distribution function at every period. In many common models the evaluation of the distribution has to be donde numerically for each value of Z_n, such as for gamma, beta or normal densities, and it is typically much slower than evaluating the density. The two-sided version (6) can thus operate much faster in these situations.

4 PHANTOM SPA ESTIMATORS

4.1 Off-Line Estimation

As is common in SPA, off-line simulations are used to estimate $E[\Delta C]$ given in (4). Notice that the limit difference process is defined by starting two inventory processes at levels S and s and using common random numbers for the demands, until both processes reach level S, thus completing a regenerative cycle. This limit process is therefore independent of whether a right-sided ($\Delta > 0$), a left-sided ($\Delta < 0$) or a two-sided formula is used.

As explained in Bashyan and Fu (1991), the usual implementation of SPA for this example would require generating the inventory process (1) for N periods to estimate, say $\sum g(Z_n)$ –if the two-sided version is used. Next, an off-line simulation is performed where the difference process is generated as follows. Two initial levels are prescribed, namely $X_0^+ = S$ and $X_0 = s$. Demands are generated and used in (1) to evaluate the ensuing "nominal" and "phantom" processes, until their levels coincide, which in the case of a continuous demand distribution, hapens only when they both order and reach level S. This constitutes a cycle and, naturally, may contain several periods. It is also possible that the nominal system reaches level S several times before the two processes reach S simultaneously. The off-line simulation replicates M such cycles to obtain an estimate of the difference in cost.

4.2 On-Line Estimation via Parallel Phantoms

Consider the two-sided SPA formula (6) (parallel phantoms can also be implemented for the one-sided versions of SPA). We propose here to implement parallel phantom systems to simultaneously estimate $E\left(\sum g(Z_n)\right)$ as well as $E[\Delta C]$ using a single path. In accordance with our description in the previous subsection, the latter estimator requires starting two processes with initial levels $X_0^+ = S, X_0 = s$.

First, suppose that we only simulate a single long path of length N. We start the nominal process at level S and use the path information to compute $\sum g(Z_n)$. In addition, we also define a phantom process that starts at level s and uses the same sequence of demands. Their difference process is obtained using the sequence of demands from the (unique) trajectory of the process being simulated. Once a cycle of the difference process has finished, another cycle can be computed by resetting the initial conditions: the next time that the nominal process reaches level S a phantom is started in parallel at level s. At the end of the simulation, the estimators of $E\left(\sum g(Z_n)\right)$ and of $E[\Delta C]$ are obtained from the observation of the system. Although the estimators are correlated, they are both consistent and their product is also consistent for the SPA derivative.

As mentioned before, however, the estimation of $E[\Delta C]$ uses cycles of the difference process. Within the first cycle, it is possible that the nominal process reaches S. If this happens before the end of the cycle, we create a second phantom system that starts at s and will contribute a second difference process, using again the same demands as the nominal and the first phantom systems. Continuing in this manner, every time that the iventory level in the nominal system is filled up to S we start a new parallel phantom system initiated at level s. By using common random numbers, we expect to keep the variance at reasonable levels, while increasing the number of observations of cycles of the difference process.

5 COMPUTER SIMULATIONS

Our simulations were performed on a PC using PASCAL. We used $K = 10, h = 10, p = 50, s = 10, S = 20$ to compare our Phantom SPA estimtion with Fu's SPA. The simulation horizon was set at $N = 5000$ and in the case of off-line simulations, $M = 5000$ cycles were used. Table 1 shows the result of the estimation for an approximate 95% confidence interval using the CLT. We compare the one and two sided versions of SPA (denoted SPA$^{\pm}$ and SPA* respectively), using in all cases off-line simulations for estimating $E[\Delta C]$. The first two columns give Fu's method, as explained in Bashyan and Fu (1991). Notice from this table that the reduction in variance obtained with the two-sided version (6) is not very high and it seems to be problem dependent. In any case, we believe these results indicate that this is the more consistent and robust estimator of the three.

In Table 2 we present the results of the estimation using the two-sided SPA estimator in (6), in one column we give the results using off-line simulations for $E[\Delta C]$ and in the other, using on-line parallel phantoms.

Table 1: Results with Off-Line Simulations

Distribution	Theoretical	SPA$^-$	SPA$^+$	SPA*
Exp(1/3)	−4.531	-4.680 ± 0.26	-4.570 ± 0.10	-4.572 ± 0.08
Exp(1/5)	−2.862	-2.779 ± 0.145	-2.843 ± 0.088	-2.837 ± 0.082
Exp(1/10)	1.516	1.586 ± 0.114	1.569 ± 0.094	1.576 ± 0.101
$U(0, 6)$	–	-5.185 ± 0.171	-5.225 ± 0.171	-5.261 ± 0.122
$U(0, 12)$	–	-5.359 ± 0.211	-5.312 ± 0.088	-5.259 ± 0.044
$U(0, 20)$	–	-2.181 ± 0.100	-2.212 ± 0.084	-2.200 ± 0.079

Table 2: Results of Two-Sided SPA

	Off-Line		Parallel Phantoms	
Distribution	SPA*	CPU (secs)	SPA*	CPU (secs)
Exp(1/3)	-4.509 ± 0.081	25.54	-4.530 ± 0.027	11.91
Exp(1/5)	-2.816 ± 0.097	14.99	-2.829 ± 0.056	8.40
Exp(1/10)	-1.488 ± 0.088	10.5	-1.497 ± 0.078	6.75
$U(0, 6)$	-5.180 ± 0.085	45.91	-5.281 ± 0.025	21.69
$U(0, 12)$	-5.234 ± 0.036	11.31	-5.232 ± 0.018	7.30
$U(0, 20)$	2.242 ± 0.130	7.19	2.224 ± 0.084	5.32

It should be clear from these results that the greatest contribution to the improvement in efficiency of the Phantom SPA comes from the parallel on-line phantom systems rather than from the implementation of the two-sided version of SPA. The gain in precision is noticeable and the gain in speed is remarkable. We would like to point out that this gain is even more dramatic when normal or gamma distributions are used, where the distribution function is calculated numerically at each period for the one-sided formula.

6 CONCLUDING REMARKS

We have presented a new approach to implement SPA, which often requires off-line simulations of a difference process. We bypass this stage by evaluating parallel systems with common random variables, called the phantom systems. Our method uses one single path for the estimation and improves considerably the efficiency of the estimation both in variance reduction as well as computational time. Besides an improvement in the efficiency of the estimator, we believe this technique is amenable to on-line optimization as well, since it performs the calculations in parallel and using only the observations from the nominal system. We have chosen to focus on a simple example to introduce the mehtod, but it is general for SPA estimators and can readily be implemented for many situations. We are currenlty working on the general setting and applying the method to other problems in finance and queueing. We are also developing a general approach to study the statistical properties of the parallel phantom estimators.

ACKNOWLEDGMENTS

Supported in part by NSERC-Canada grant # WFA0184198.

REFERENCES

Bashyan, S., and M. Fu. 1991. Application of perturbation analysis to (s, S) inventory systems. In *Proc of the Winter Simulation Conference*, 542–550. B. Nelson, D. Kelton, G. Clark Eds.

Brémaud, P., and F. Vázquez-Abad. 1992. On the pathwise computation of derivatives with respect to the rate of a point process: the phantom RPA method. *Queueing Systems, Theory and Appl*, 10:249 – 270.

Taylor, H., and S. Karlin. 1994. *Introduction to Stochastic Processes*. Revised edition ed. Boston: Academic Press, Inc.

Vázquez-Abad, F., and K. Davis. 1995. Efficient implementation of the phantom RPA method, with an application to a priority queueing system. Technical Report 989, Université de Montréal.

Vázquez-Abad, F., and L. Zubieta. 1999. Ersätz estimation for sensitivity analysis: a methodology. To be submitted.

AUTHOR BIOGRAPHIES

FELISA J. VÁZQUEZ-ABAD is Professor at the Department of Computer Science and Operations Research at the University of Montreal. She is also Research Fellow at the Department of Electrical and Electronic Engineering at the University of Melbourne. She received her B.Sc. degree in Physics in 1983 and her M.Sc. degree in Statistics and OR in 1984 from Universidad Nacional Autónoma de México, and her Ph.D. degree in Applied Mathematics in 1989 from Brown University. Her research interests include stochastic control and simulation with applications in telecommunications, manufacturing, transportation, insurance and finance. She is a member of SIAM, INFORMS and IEEE.

MANUEL CEPEDA-JÜNEMAN received his M.Sc. degree in Computer Science in 1998 from the University of Montreal. He is currently working towards his PhD degree.