UNIVERSITEIT
AMSTERDAM

VU Research Portal

Gradient estimation for multicomponent maintenance systems with age-replacement policy

Heidergott, B.F.; Farenhorst - Yuan, T.

document version
Publisher's PDF, also known as Version of record

Link to publication in VU Research Portal

citation for published version (APA)
Heidergott, B. F., & Farenhorst - Yuan, T. (2010). Gradient estimation for multicomponent maintenance systems with age-replacement policy. *Operations Research*, *58*, 706-718. https://doi.org/10.1287/opre.1090.0760

Gradient Estimation for Multicomponent Maintenance Systems with Age-Replacement Policy

Bernd Heidergott

Department of Econometrics and Operations Research and Tinbergen Institute, Vrije Universiteit Amsterdam,
Amsterdam, The Netherlands, bheidergott@feweb.vu.nl

Taoying Farenhorst-Yuan

Department of Econometrics and Operations Research, Vrije Universiteit Amsterdam,
Amsterdam, The Netherlands, tyuan@feweb.vu.nl

We consider multicomponent maintenance systems with an F-failure group age-replacement policy: it keeps failed components idling until F components are failed and then replaces all failed components together with the nonfailed components whose age has passed the critical threshold age θ_n for components of type n. With each maintenance action, costs are associated. We derive various unbiased gradient estimators based on the measure-valued differentiation approach for the gradient of the average cost. Each estimator has its own domain of applicability. We also compare the performance of our gradient estimators when applied to stochastic optimization with other general gradient-free methods.

1. Introduction

Nowadays individuals and companies depend more and more on complex systems for their daily functioning. The potential costs of breakdowns of these systems are quite high, so the maintenance of systems is crucial to increase their availability. Therefore, during the past decades, there has been a growing interest in the modeling and optimization of the maintenance of systems consisting of multiple components. For overviews on multicomponent maintenance optimization, we refer to Nicolai and Dekker (2006) and the references therein.

Maintenance activities usually involve an intervention cost and a component replacement cost. Typically, the intervention cost (e.g., set-up, repair, instrument or system down cost, etc.) is relatively high. To avoid immediate failure replacements, there is usually some redundancy, so that one can postpone the replacement of a failed component until the number of failed components reaches a certain predefined number. Moreover, the intervention cost can be reduced when maintenance activities on different components are executed simultaneously. Therefore, it is worthwhile to replace components preventively during maintenance of failed components.

In this paper, we consider the following modified F-group failure replacement rule. There are J interchangeable components that may have different lifetime distributions. We say that two components are of the same type if their lifetime distribution is identical. We assume that there

are N types of components in the system and an age threshold is assigned to each type of component. Failed components will be kept idling for a certain time until there are in total F failed components. Then, all failed components are replaced together with "old" nonfailed components. More precisely, if the age of a component of type n upon the maintenance action has passed the critical threshold age θ_n, then this component is preventively replaced. The age policy for the system is completely described by the vector $\theta = (\theta_1, \ldots, \theta_N)$. This is an opportunistic maintenance rule; see Dekker et al. (1998). According to the terminology in maintenance theory, we study an N-type $(J - F + 1)$-out-of-J system, which means that at least $J - F + 1$ out of J components must operate. If $F > 1$, then there is redundancy in the system and it fails less often than its individual components. This way a certain reliability can be guaranteed.

Upon each F-group failure, a maintenance action is taken. With each maintenance action, costs are associated. The optimal age policy is given by the value of θ that yields minimal average cost $C(T, \theta)$ over a given time period $T > 0$, which, for example, models the length of the contract period for maintenance. Note that under suitable stability conditions, choosing T appropriately, $C(T, \theta)$ approximates the long-run average cost.

A typical application of the age-replacement policy to an $(N = 1)$-type $(J - F + 1)$-out-of-J system is given in Dekker et al. (1998). Therein, the authors deal with maintenance of light standards illuminating a container terminal

at Rotterdam harbor. Such a light standard consists of J identical lights. To guarantee a minimum luminance, the lamps are replaced if the number of failed lamps reaches a predefined number F. To replace the lamps, the assembly has to be lowered. This set-up activity is an opportunity to combine corrective and preventive maintenance. An example of applying the age-replacement policy to an $(N = 2)$-type $(J - F + 1)$-out-of-J system is given in Van der Duyn Schouten et al. (1998), where the problem of replacing light bulbs in traffic control signals is studied. More specifically, each installation consists of three compartments in which light bulbs serve the green, red, and yellow lights. The components (light bulbs) seem to be identical, but the yellow lights burn less often, and their failure rate differs from that of the other lights. Hence, two types of components can be defined in such a system.

In cost minimization, one is searching for the optimal solution $\theta^\star$ such that

$$C(T, \theta^\star) = \min_{0 \leqslant \theta_n < \infty,\, 1 \leqslant n \leqslant N} C(T, \theta_1, \ldots, \theta_N). \qquad (1)$$

However, most real-world systems violate the rather restrictive conditions for obtaining an explicit expression for $C(T, \theta)$ in terms of θ or suffer from intractability when the number of components grows. Thus, one has to resort to simulation for obtaining the optimal threshold value. A standard method for finding the optimal threshold value in an iterative procedure is stochastic approximation (SA). The general form of SA is

$$\theta(k+1) = \Pi_{[0, \infty)}(\theta(k) - a_k \nabla C_k), \qquad (2)$$

where $\theta(k)$ is the parameter vector at the beginning of iteration k, ∇C_k is an estimate of $\nabla C(T, \theta(k))$ (the gradient of $C(T, \theta(k))$), a_k is a (positive) sequence of step sizes, and $\Pi_{[0, \infty)}$ is the projection onto $[0, \infty)$. It can be shown that under suitable conditions $\theta(k) \to \theta^*$ for k toward ∞ with probability one. When $\nabla C(T, \theta(k))$ is an unbiased estimator for $\nabla C(T, \theta(k))$, (2) is called a Robbins-Monro algorithm, but when a finite difference estimator is used, it is called a Kiefer-Wolfowitz algorithm; see Kushner and Clark (1978) for more details. It is well known that the Robbins-Monro algorithm has faster convergence rates than the Kiefer-Wolfowitz algorithm. This motivates our research on unbiased, sample-path-based gradient estimation methods.

It is worth noting that, in general, state-independent policies are not optimal; see L'Ecuyer (1983) and Haurie and L'Ecuyer (1986). However, due to the complexity of the problem, computing an optimal state-dependent policy is usually computationally infeasible. Therefore, we restrict our analysis to the class of parameterized age-replacement policies described above.

We establish three unbiased gradient estimators based on *measure-valued differentiation* (MVD) for the derivative of $C(T, \theta)$ with respect to θ for a multicomponent system

with general lifetime distribution. The key step for applying MVD is to derive a Markov chain description, so that the associated transition kernel is differentiable with respect to the components of the parameter vector. For a multicomponent maintenance system we first show how one can derive a Markov description with a differentiable kernel. Based on this result, we will introduce unbiased phantom estimators for $\nabla C(T, \theta)$.

To evaluate our phantom estimators, we perform a number of numerical experiments. For comparison, we choose finite differences (FD), two-sided FD, and the simultaneous perturbation (SP) method because they are the only estimators available for our problem; see §3 for a detailed discussion of other gradient estimators. We perform a numerical study with respect to two criteria:

(1) the "work-normalized variance," which balances the computational effort and variance of the estimator, and is given by the product of the variance and the expected work per run, see Glynn and Whitt (1992);

(2) the number of iterations and the total computation time required for finding the optimal value for θ via SA.

This paper is organized as follows. In §2, we give a detailed description of the system model. A detailed discussion of the literature is provided in §3. In §4, we provide a brief discussion on the MVD method and the phantom processes, and we present the phantom estimator and the randomized phantom estimator. In §5, we introduce a new estimator, called the combined phantom estimator, and we provide sufficient conditions for unbiasedness of the phantom estimators. In §6, we first compare the performance of the phantom estimator with the FD, two-sided FD, and SP estimator according to the above criterion (1) with the help of a series of numerical experiments. We conclude §6 by comparing the estimators with respect to criterion (2) in an optimization.

2. The System Model

We consider a general maintenance system consisting of $N \geqslant 1$ types of components. There are $J_n \geqslant 1$ components of type n and there are in total $\sum_{n=1}^N J_n = J$ components in the system. We denote by $\mathbb{J} = \{1, \ldots, J\}$ the set of components, and we denote by $\mathbb{J}_n$ the set of components belonging to type n. Without loss of generality, we assume that components are numbered in ascending order; i.e., component 1 to J_1 are of type 1 and $J_1 + 1$ to $J_1 + J_2$ are of type 2, and so on. Furthermore, for any $\mathscr{A} \subset \mathbb{J}$, we write $\mathscr{A}^c := \mathbb{J} \backslash \mathscr{A}$.

The lifetime of the components of type n are independent and identically distributed with cumulative distribution function $F_n(\cdot)$ for $1 \leqslant n \leqslant N$, with density f_n. For reference, we introduce the following assumption.

ASSUMPTION 1 (A1). *For each n, with $1 \leqslant n \leqslant N$, we assume that $f_n(\cdot)$ is continuous on its support and that the support of $f_n(\cdot)$ is a subset of $[0, \infty)$ containing at least $[0, T]$.*

Let $[j]$ denote the type of component j; i.e., $[j] = n$ if and only if $j \in \mathbb{J}_n$. Then, the residual lifetime of the component j in state $s = (\tau, u, a_1, \ldots, a_J) \in S$, denoted by $R_j(s)$, has density

$$h_j(t \mid a_j) = \frac{f_{[j]}(t + a_j)}{1 - F_{[j]}(a_j)}, \quad t \geqslant 0. \tag{3}$$

Note that (A1) implies that $h_j(\cdot \mid a_j)$ in (3) is continuous on $[0, \infty)$. To avoid inflation of the subscript $\theta \in (0, \infty)^N$ in the notation, we will suppress denoting the dependence on θ when this leads to no confusion.

Let $(\Omega, \mathscr{A}, \mathbb{P})$ denote a probability space and assume that all random variables introduced are defined on this probability space. We say that every F-failure causes a *maintenance action* and the F-failure epochs are called *maintenance action epochs*. We introduce the following variables:

- $A(j, m)$, the age of component j right after the mth maintenance action epoch;
- $\tau(m)$, the time at which the mth maintenance action takes place; with $\tau(0) = 0$;
- $T(m)$, the time elapsed between the $(m-1)$st and the mth maintenance action; more specifically, $T(m) = \tau(m) - \tau(m-1)$ for $m \geqslant 1$;
- $\mathscr{J}(m)$, the set of components that have failed during time $T(m)$;
- $\mathscr{R}(m)$, the set of components that have been preventively replaced at the mth maintenance epoch; note that $\mathscr{R}(m) \subseteq (\mathscr{J}(m))^c$.

The system evolves as follows. The system starts at time zero with J components new. After $\tau(k)$ time-units, the kth maintenance action occurs. Upon the kth maintenance action, the F failed components, given by set $\mathscr{J}(k)$, are replaced and, due to our replacement policy, we preventively replace all components $j \in \mathscr{R}(k)$. More specifically, component j of type $[j]$ is preventively replaced if $A(j, k-1) + T(k) > \theta_{[j]}$. All replaced components are of age $A(j, k) = 0$, and the nonreplaced components are of age $A(j, k) = A(j, k-1) + T(k)$.

We describe the system by a Markov chain

$$X_\theta \triangleq \{X_\theta(m) \triangleq (\tau(m), T(m), A(1, m), \ldots, A(J, m)): m \in \mathbb{N}\},$$

and denote its state space by $S \triangleq [0, \infty)^{J+2}$. We assume that the system starts at time 0 with all components new and set $X_\theta(0) = 0^{J+2}$, where 0^{J+2} denotes the zero vector of length $J + 2$. A given state $s \in S$ is denoted by $s = (\tau, t, a_1, \ldots, a_J)$, where τ denotes the current time, t denotes the holding time, and a_j denotes the age of component j. The transition kernel of this Markov chain is denoted by Q_θ; we will give the explicit form of Q_θ in §5.1.

With each maintenance action, there is cost associated. Let $g_c: S \mapsto [0, \infty)$ be the cost function. For example, if the replacement of one component causes cost $C_R > 0$ and the

cost incurred for an intervention is $C_I > 0$, then we obtain the classical cost function for $(\tau, t, a_1, \ldots, a_J) = s \in S$,

$$g_c(s) \triangleq C_I + n(s)C_R, \tag{4}$$

where $n(s) = \sum_{j=1}^{J} \mathbb{1}\{a_j = 0\}$ is the number of components replaced.

For $s \in S$, we introduce the stopping time $M(T, \theta) = \inf\{k \geqslant 1: \tau(k) \geqslant T\}$ as the index of the first maintenance action after T provided that all components are initially new. The performance measure over a fixed time horizon $C(T, \theta)$ can be written as

$$C(T, \theta) = \frac{1}{T} \mathbb{E}\left[\sum_{i=1}^{M(T, \theta)-1} g_c(X_\theta(i)) \right]. \tag{5}$$

Differentiating the expression in (5) with respect to θ_n yields

$$\frac{\partial}{\partial \theta_n} C(T, \theta) = \frac{1}{T} \frac{\partial}{\partial \theta_n} \mathbb{E}\left[\sum_{i=1}^{M(T, \theta)-1} g_c(X_\theta(i)) \right]. \tag{6}$$

In what follows we will use MVD to establish unbiased estimators for the above expression.

3. Review of Literature

Optimization of a multicomponent maintenance system with age-replacement policy has been studied by L'Ecuyer et al. (1999) using a "split-and-merge" approach. This approach has been developed by Vázquez-Abad et al. (1996) and L'Ecuyer and Vázquez-Abad (1997). The key idea of this method is that different values of θ only lead to finitely many different outcomes of the cycle costs; i.e., the sample cost is a piecewise constant function in θ. Simulating all possible alternatives in parallel, the method is capable of estimating the entire performance function within one experiment. For a single type J-out-of-J system (i.e., $N = 1$, $F = 1$), the authors proved that the expected growth rate of the tree is bounded linearly in the total number of simulated regenerative cycles and quadratic in the number of the components in the system. It is worth noting that the actual optimization based on the split-and-merge approach for threshold-type parameters (such as θ defining our age-replacement policy) is still an open question; see L'Ecuyer (1993).

An importance-sampling-based approach to gradient estimation is the *score-function* (SF) method; see Rubinstein (1998). SF yields an unbiased gradient estimator through rescaling the observed sample-path performance by appropriate weights. SF is designed for sensitivity analysis w.r.t. distributional parameters and cannot be applied to our maintenance problem in a straightforward way. As pointed out in Rubinstein (1992), it is sometimes possible to translate the dependence of the sample performance on θ into a

θ-dependent transformation of a random variable and thus turn θ into a distributional parameter. This is called the "push out" SF in Rubinstein (1992). However, applying the push out SF requires that the dependence of the sample cost on θ is given in a simple explicit form, which is not the case for the maintenance problem under study.

Perturbation analysis (PA) provides derivative estimators based on sample path analysis; see Glasserman (1988; 1991a, b, c), Ho and Cao (1991), Suri (1987, 1989), and Heidelberger et al. (1988). Infinitesimal PA (IPA) is the technique of exploiting the commutativity of the differential and expectation, and it is known as the most efficient means of sensitivity estimation from a single sample path. However, a condition for unbiasedness of IPA is Lipschitz continuity of the sample performance function in the parameter of interest. Lipschitz continuity fails for our maintenance problem because θ is a threshold variable and a small change in θ may yield a discrete change in the sample path dynamic.

Smoothed perturbation analysis (SPA), as introduced by Gong and Ho (1987) and further developed by Glasserman and Gong (1990), overcomes the restriction of IPA to continuous sample cost functions. SPA is a technique that takes advantage of the conditional expectation to smooth out the discontinuities of sample performance functions and give the unbiased estimator for general models. For details on SPA, we refer to the monograph by Fu and Hu (1997). An SPA estimator for $dC(T, \theta)/d\theta$ is derived by Fu et al. (1993). Although the analysis of the SPA estimator is carried out for a single-type multicomponent maintenance system (i.e., a single-type J-out-of-J system), the actual estimator is, due to its complexity, only computationally efficient for the case of $J = 2$ under the assumption that the lifetime of the components are uniformly distributed on $[0, \eta]$ for $\eta > 0$. As the study in Heidergott and Farenhorst-Yuan (2007) shows, the phantom estimator has a considerably lower variance than the SPA estimator and requires less computational effort for a two-component maintenance system with uniformly distributed lifetime.

A weak derivative approach to optimization of the threshold parameters in a multicomponent maintenance system is given in Heidergott (2001). Therein, a maintenance system with identical components (i.e., a single-type J-out-of-J system) is analyzed. The estimator presented in Heidergott (2001) involves the computation of a polynomial that can only be solved explicitly in special cases. The estimator presented in our paper will not require computation of the polynomial as described in Heidergott (2001).

The above sophisticated methods either do not apply to the cost optimization problem considered in this paper or are numerically inefficient. For this reason, we will resort to the general purpose methods such as FD, two-sided FD, and SP as benchmarks in our numerical experiments.

4. The Measure-Valued Differentiation Approach

4.1. Basic Definition

MVD is an extension of the concept of weak differentiation. For details on weak differentiation see Pflug (1996) and Heidergott et al. (2010b). The key step for the MVD approach is computing $\partial Q_\theta / \partial \theta_n$ for $1 \leqslant n \leqslant N$, i.e., the nth partial derivative of the Markov kernel of X_θ. We refer to Heidergott and Vázquez-Abad (2008) for introductory examples of the MVD method. It is easily shown that $\partial Q_\theta / \partial \theta_n$ fails to be a Markov kernel. Fortunately, one can represent $\partial Q_\theta / \partial \theta_n$ by the rescaled difference between two Markov kernels Q_θ^{n+} and Q_θ^{n-}. More specifically, let Q_θ be a Markov kernel on state-space S and let C^b denote the set of continuous bounded mappings from S to $\mathbb{R}$. If for n, with $1 \leqslant n \leqslant N$, and $\forall s \in S$,

$$\forall g \in C^b : \frac{\partial}{\partial \theta_n} \int g(u) Q_\theta(s; du)$$

$$= c_\theta^n(s) \left(\int g(u) Q_\theta^{n+}(s; du) - \int g(u) Q_\theta^{n-}(s; du) \right), \quad (7)$$

then $(c_\theta^n, Q_\theta^{n+}, Q_\theta^{n-})$ is called a *partial weak derivative of Q_θ with respect to θ_n*; see Heidergott and Vázquez-Abad (2008). In shorthand notation, $\partial Q_\theta / \partial \theta_n = c_\theta^n (Q_\theta^{n+} - Q_\theta^{n-})$. It can be shown that the fact that (7) holds for all $g \in C^b$ implies that $\partial Q_\theta / \partial \theta_n$ is uniquely defined; see Billingsley (1966). Under quite general conditions, the existence of weak derivatives follows from a Hahn-Jordan-type decomposition for Markov kernels; see Heidergott et al. (2008). Note that g_c defined in (4) is an element of C^b.

For the reader less acquainted with MVD, it is worth noting that (7) often can be checked in a simple way. Provided that θ is a distributional parameter (e.g., the mean of the distribution), MVD enjoys a modular structure: weak differentiability of the input distributions of a discrete event system imply, under quite general conditions, weak differentiability of the Markov kernel modeling the system process. For example, if the service time distribution in a $G/G/1$ queue is weakly differentiable with respect to θ, then, under some mild integrability conditions, weak differentiability of the Markov kernel of the waiting times with respect to θ follows; see Heidergott et al. (2010a). Once weak differentiability of the Markov kernel of the system process is established, under appropriate stability conditions, weak differentiability of the Markov kernel of the system dynamic implies weak differentiability of the performance over a cycle or the stationary performance; see Heidergott and Vázquez-Abad (2006) and Heidergott et al. (2006). In case θ is a threshold parameter, as in the maintenance model under consideration, the analysis is more challenging because for this problem the Markov kernel has to be analyzed directly (i.e., the modular structure is lost). However, as we will show in this paper, for our maintenance problem the validity of (7) follows from two simple

assumptions, (A1) and (A2) (the latter being introduced in §5).

MVD for Markov chains is an operator approach to calculating derivatives of k-fold convolutions of Markov kernels such as Q_θ. If Q_θ has partial weak derivative $(c_\theta^n, Q_\theta^{n+}, Q_\theta^{n-})$ with respect to θ_n, then in Heidergott and Vázquez-Abad (2008) it is shown that under quite general conditions

$$\frac{\partial Q_\theta^k}{\partial \theta_n} = \sum_{j=1}^{k} c_\theta^n (Q_\theta^{k-j} Q_\theta^{n+} Q_\theta^{j-1} - Q_\theta^{k-j} Q_\theta^{n-} Q_\theta^{j-1}). \quad (8)$$

Hence, $\partial Q_\theta / \partial \theta_n$ and consequently the derivative of $\mathbb{E}[g(X_\theta(k+1)) \mid X_\theta(1) = 0^{J+2}]$ with respect to θ_n, for $g \in C^b$, can be expressed by the sum over k differences evaluated for appropriately defined Markov chains. More specifically, the transition kernel of these Markov chains is Q_θ except for one transition (indexed by j in (8)) for which the transition kernel is replaced by either Q_θ^{n+} or Q_θ^{n-}, respectively. These versions of the nominal Markov chain X_θ are called *phantoms* and will be formally introduced in the next section (for an early reference on the use phantoms in gradient estimation, see Brémaud and Vázquez-Abad 1992). The formula for the derivative of cost accumulated over a fixed number of transitions can be easily derived from (8).

4.2. Phantom Processes for the nth Partial Derivative

The definition in (7) can be expressed in terms of random variables $X_\theta^{n\pm}(s)$ as

$$\forall g \in C^b \colon \frac{\partial}{\partial \theta_n} \int g(u) Q_\theta(s; du)$$
$$= \mathbb{E}_\theta[c_\theta^n(s)(g(X_\theta^{n+}(s)) - g(X_\theta^{n-}(s)))] \quad (9)$$

for all $s \in S$, where $X_\theta^{\pm}(s)$ are appropriately defined random variables such that for all $g \in C^b$ and all $s \in S$,

$$\mathbb{E}_\theta[g(X_\theta^{n\pm}(s))] = \int g(u) Q_\theta^{n\pm}(s; du). \quad (10)$$

We introduce the "plus" and "minus" processes $\{X_\theta^{n\pm}(s, k) \colon k \geqslant 1\}$, called *phantoms for the nth partial derivative*, as follows. At a particular state s, the nominal process "splits" in three different trajectories. Suppose that the phantoms are generated from the kth state of X_θ; that is, $s = X_\theta(k)$. The transition from s to $X_\theta^{n+}(s, 1)$ is governed by Q_θ^{n+} and that from s to $X_\theta^{n-}(s, 1)$ by Q_θ^{n-}, respectively. That is, $X_\theta^{n\pm}(s, 1) = X_\theta^{n\pm}(s)$, where $X_\theta^{n\pm}(s)$ satisfies expression (10). For $k > 1$, the transition from $X_\theta^{n\pm}(s, k-1)$ to $X_\theta^{n\pm}(s, k)$ is governed by Q_θ. That is, after the splitting, the transition kernels are equal for the three processes. Consequently, for $k \geqslant 1$ it holds that

$$\mathbb{P}(X_\theta^{n\pm}(s, k) \in A \mid X_\theta(k) = s) = (Q_\theta^{k-1} Q_\theta^{n\pm})(s; A)$$

for any $s \in S$ and any measurable $A \subset S$. In §5.2, we will show how to obtain $X_\theta^{n\pm}(s, 1)$ for our maintenance system. With the help of the phantom processes, operator Equation (8) now reads

$$\frac{\partial}{\partial \theta_n} \mathbb{E}[g(X_\theta(k+1)) \mid X_\theta(0) = s]$$

$$= \sum_{j=1}^{k} \mathbb{E}[c_\theta^n(X_\theta(j))(g(X_\theta^{n+}(X_\theta(j), k-j+1))$$
$$- g(X_\theta^{n-}(X_\theta(j), k-j+1))) \mid X_\theta(0) = s] \quad (11)$$

for all $g \in C^b$ and all $s \in S$.

Gradient estimation for the case that X_θ is evaluated over a random horizon is discussed in the following section.

4.3. Phantom Estimators

For $g \in C^b$ and $X_\theta(m) = s \in S$, $m \geqslant 1$, the *scaled difference* between the cumulative costs over the plus and minus phantoms for the nth partial derivative is given by

$$D^n(g, s) \triangleq c_\theta^n(s) \Bigg(\sum_{k=1}^{M(X_\theta^{n+}(s,1), T, \theta)-1} g(X_\theta^{n+}(s, k))$$
$$- \sum_{k=1}^{M(X_\theta^{n-}(s,1), T, \theta)-1} g(X_\theta^{n-}(s, k)) \Bigg), \quad (12)$$

where $M(s, T, \theta) = \inf\{k \geqslant 1 \colon X_\theta(k) \in [T, \infty] \times [0, \infty)^{J+1}$ and $X_\theta(0) = s\}$. Note that with this definition it holds that $M(T, \theta) = M(0^{J+2}, T, \theta)$. In §5, we will derive the initial state of the plus and minus processes, i.e., $X_\theta^{n\pm}(s, 1)$, for maintenance systems. The computational burden for evaluating (12) can be reduced by coupling the phantoms; for details see §EC.1 in the electronic companion.

An electronic companion to this paper is available as part of the online version that can be found at http://or.journal.informs.org/.

For $g \in C^b$, the *phantom estimator* (PhE) for the nth partial derivative is given by

$$\text{PhE}^n(g) \triangleq \sum_{k=1}^{M(T, \theta)-1} D^n(g, X_\theta(k)). \quad (13)$$

$\text{PhE}^n(g)$ can be phrased as follows: at each state of the nominal processes a plus and a minus phantom for the nth partial derivative is generated and the derivative is estimated by the scaled difference between the cumulative costs over the plus and minus phantoms.

Alternatively, we may obtain the derivative by the so-called *randomized phantom estimator* (RPhE) for the nth derivative. It works as follows: (i) we first choose an integer σ independent of everything else and uniformly distributed over $\{1, \ldots, M(T, \theta)-1\}$; (ii) we generate a plus and minus phantom at state $X_\theta(\sigma)$; and (iii) we evaluate the difference between the cumulative costs of the plus and the minus

Figure 1.   The phantom estimator and the randomized phantom estimator for the nth partial derivative.
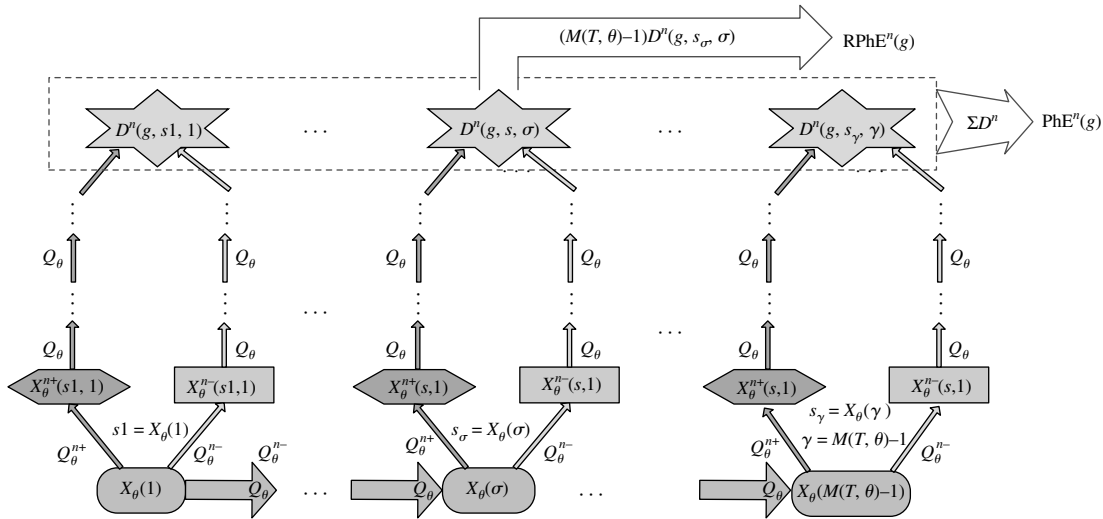
phantom and rescale this difference by $M(T, \theta) - 1$. For all $g \in C^b$, the resulting estimator becomes

$$\text{RPhE}^n(g) \triangleq \big(M(T, \theta) - 1\big)D^n(g, X_\theta(\sigma)). \tag{14}$$

To the best of our knowledge, the idea of replacing the sum of $M(T, \theta) - 1$ by uniformly choosing one particular transition between 1 and $M(T, \theta) - 1$ for phantom generation first appeared in Vázquez-Abad and Davis (1995).

We conclude this section with a discussion of RPhE and PhE. The essential difference between these two estimators is that PhE generates, from each state of a simulation run, two phantoms (namely, the plus and minus phantoms, respectively), whereas RPhE generates phantoms only at one randomly chosen state. This essential difference is illustrated in Figure 1, which depicts the situation for a simulation run. As can be seen in Figure 1, PhE requires simulating $2(M(T, \theta) - 1)$ phantoms (those generated at states $X_\theta(1)$ to $X_\theta(M(T, \theta) - 1)$), whereas RPhE only requires two phantoms per cycle to be simulated (those generated at $X_\theta(\sigma)$, with σ uniformly distributed over $\{1, \ldots, M(T, \theta) - 1\}$). Numerical examples will illustrate the performance of these estimators in §6.

5. The Phantom Estimator for Multicomponent Maintenance Systems

5.1. The Markov Kernel

In this section, we provide a discrete-time Markov chain description of an N-type F-group failure multicomponent maintenance system such that its associate transition kernel is differentiable as a function in θ. Note that the previously mentioned kernel Q_θ is being explicitly defined here. For $m \geqslant 0$, let

$$X_\theta(m+1) = \phi(\mathcal{J}(m+1), \mathcal{R}(m+1), T(m+1), X_\theta(m))$$
$$= (\tau(m) + T(m+1), T(m+1),$$
$$A(1, m+1), \ldots, A(J, m+1)) \tag{15}$$

denote the state mapping function. The transition from $X_\theta(m)$ to $X_\theta(m+1)$ as formalized in (15) can be phrased as follows: the transition from $X_\theta(m)$ to $X_\theta(m+1)$ is triggered by the F-group failure of components $\mathcal{J}(m+1)$, the holding time of $X_\theta(m)$ is $T(m+1)$, and the components out of $\mathcal{R}(m+1)$ are preventively replaced upon this failure. Moreover, the age of a component is updated as follows:

$$A(j, m+1) = \begin{cases} 0, & j \in \mathcal{J}(m+1) \cup \mathcal{R}(m+1), \\ A(j, m) + T(m+1), & \\ & j \in (\mathcal{J}(m+1) \cup \mathcal{R}(m+1))^c. \end{cases} \tag{16}$$

The key idea for deriving the transition kernel Q_θ is as follows. Given the components which constitute the next F-group failure (i.e., $\mathcal{J}(m+1)$), the time to the next failure (i.e., $T(m+1)$) can be determined. On the basis of the F-group failure and the holding time, the components which are preventively replaced (i.e., $\mathcal{R}(m+1)$) can be determined.

From (3), the transition out of $X_\theta(m) = s = (\tau, t, a_1, \ldots, a_J) \in S$ is triggered by a failure of component set $\mathcal{J}(m+1)$ with probability

$$\mathbb{P}(\mathcal{J}(m+1) = \mathcal{J} \mid X_\theta(m) = s)$$
$$= \sum_{j \in \mathcal{J}} \int_0^\infty \left(\prod_{i \in \mathcal{J} \setminus \{j\}} \int_0^u h_i(r \mid a_i)\, dr \right.$$
$$\left. \cdot \prod_{i \in \mathcal{J}^c} \int_u^\infty h_i(r \mid a_i)\, dr \right) h_j(u \mid a_j)\, du,$$

$s \in S$. Provided that components in $\mathcal{J}(m+1) = \mathcal{J}$ trigger the transition out of state $X_\theta(m)$, the holding time

$T(m+1) = T(\mathcal{J})$ has distribution

$$\mathbb{P}(T(m+1) \leqslant t, \mathcal{J}(m+1) = \mathcal{J} \mid X_\theta(m) = s)$$

$$= \sum_{j \in \mathcal{J}} \int_0^t \left(\prod_{i \in \mathcal{J} \setminus \{j\}} \int_0^u h_i(r \mid a_i) \, dr \right.$$

$$\left. \cdot \prod_{i \in \mathcal{J}^c} \int_u^\infty h_i(r \mid a_i) \, dr \right) h_j(u \mid a_j) \, du, \quad t \geqslant 0, \quad (17)$$

for $s = (\tau, t, a_1, \ldots, a_J) \in S$. Thus, the density of $T(\mathcal{J})$, denoted by $h(\mathcal{J}, t)$, can be calculated as follows:

$$h(\mathcal{J}, t) = \sum_{j \in \mathcal{J}} h_j(t \mid a_j) \prod_{i \in \mathcal{J} \setminus \{j\}} \int_0^t h_i(r \mid a_i) \, dr \prod_{i \in \mathcal{J}^c} \int_t^\infty h_i(r \mid a_i) \, dr$$

$$= \sum_{j \in \mathcal{J}} \frac{f_{[j]}(t + a_j)}{\prod_{i \in \mathbb{J}} (1 - F_{[i]}(a_i))} \prod_{i \in \mathcal{J} \setminus \{j\}} (F_{[i]}(t + a_i) - F_{[i]}(a_i))$$

$$\cdot \prod_{i \in \mathcal{J}^c} (1 - F_{[i]}(t + a_i)). \quad (18)$$

Then, Equation (17) can be written as

$$\mathbb{P}(T(m+1) \leqslant t, \mathcal{J}(m+1) = \mathcal{J} \mid X_\theta(m) = s)$$

$$= \int_0^t h(\mathcal{J}, u) \, du, \quad t \geqslant 0. \quad (19)$$

We have explained at the beginning of this section that conditioning the holding time on the triggering F-group failure and the preventive replacements group is crucial for our analysis. However, in Equation (17), we only use the conditioning on the F-group failure. In what follows, we will present the conditioning on the preventive replacements group.

Given state $s = (\tau, t, a_1, \ldots, a_J)$ and the set of components $\mathcal{J}$ that triggers the transition out of state s, the set of preventively replaced components is determined through the holding time $T(\mathcal{J})$. Indeed, component $j \in \mathcal{J}^c$ is preventively replaced if and only if $a_j + T(\mathcal{J}) > \theta_{[j]} \Leftrightarrow T(\mathcal{J}) > \theta_{[j]} - a_j$. Hence, there exists a partitioning of $[0, \infty]$ into a finite number of disjoint intervals $\mathcal{I}_0, \ldots, \mathcal{I}_{N_\theta(\mathcal{J})}$, along with a corresponding collection of distinct subsets $\mathcal{R}_0, \ldots, \mathcal{R}_{N_\theta(\mathcal{J})} \subseteq \mathcal{J}^c$ such that the set of preventively replaced components is $\mathcal{R}_i(\mathcal{J})$ if and only if $T(\mathcal{J}) \in \mathcal{I}_i$. In the following, we will construct this partitioning. This construction will depend on s. However, to avoid a notational overflow, we will suppress the dependence on s in the notation.

For $1 \leqslant n \leqslant N$, we introduce the set of the critical holding times

$$V_\theta(n, \mathcal{J}) = \{\theta_n - a_i, i \in \mathcal{J}^c \cap \mathbb{J}_n\}, \quad (20)$$

which is the set of mutually different critical holding times of type n components. Let $V_\theta(\mathcal{J}) := \bigcup_{n=1}^N V_\theta(n, \mathcal{J})$ denote the set of all critical holding times, which is the basis for generating the partition of interval $[0, \infty]$. More specifically,

if $N_\theta(\mathcal{J})$ denotes the number of elements of $V_\theta(\mathcal{J})$, then $N_\theta(\mathcal{J}) + 1$ is the number of partitions.

We now construct an ascending ordered vector of critical holding times $U(\mathcal{J})$ with $U_0(\mathcal{J}) = 0$, $U_i(\mathcal{J}) < U_{i+1}(\mathcal{J})$, and $U_{N_\theta(\mathcal{J})+1}(\mathcal{J}) = \infty$ such that $\mathcal{I}_i = (U_i(\mathcal{J}), U_{i+1}(\mathcal{J})]$, where $U_i(\mathcal{J})$ is defined as

$$U_i(\mathcal{J}) = \min\{t \in V_\theta(\mathcal{J}) : t > U_{i-1}(\mathcal{J})\} \quad (21)$$

for $1 \leqslant i \leqslant N_\theta(\mathcal{J})$. Consequently, if $T(\mathcal{J}) \in \mathcal{I}_i = (U_i(\mathcal{J}), U_{i+1}(\mathcal{J})]$, then component j is preventively replaced if $\theta_{[j]} - a_j < U_{i+1}(\mathcal{J})$. Note that $\mathcal{R}_i(\mathcal{J})$ as a mapping of $T(\mathcal{J})$ is constant on $\mathcal{I}_i$, and it holds that

$$\mathcal{R}_i(\mathcal{J}) = \bigcup_{n=1}^N \{j \in (\mathcal{J}^c \cap \mathbb{J}_n) : \theta_n - a_j < U_{i+1}(\mathcal{J})\}. \quad (22)$$

Denote $R(\mathcal{J}) = \bigcup_{i=0}^{N_\theta(\mathcal{J})} \mathcal{R}_i(\mathcal{J})$ the set of all components that are potentially preventively replaced upon failure of the components in $\mathcal{J}$.

To illustrate the relationship between the holding time $T(\mathcal{J})$ and the preventive replacement set $R(\mathcal{J})$, we give an example in §EC.2 in the electronic companion.

Based on Equations (19) and (22), we obtain

$$\mathbb{P}(U_i(\mathcal{J}) < T(m+1) \leqslant U_{i+1}(\mathcal{J}), \mathcal{R}(m+1) = \mathcal{R}_i(\mathcal{J}),$$

$$\mathcal{J}(m+1) = \mathcal{J} \mid X_\theta(m) = s) = \int_{U_i(\mathcal{J})}^{U_{i+1}(\mathcal{J})} h(\mathcal{J}, t) \, dt. \quad (23)$$

The transition dynamic of X_θ is completely described by (15) together with (23), and we denote the corresponding transition kernel of X_θ by Q_θ: $Q_\theta(s; A) \triangleq \mathbb{P}(X_\theta(m+1) \in A \mid X_\theta(m) = s)$ for all $s \in S$ and any measurable $A \subset S$. Let $\mathcal{J}^f = \bigcup_{\substack{\mathcal{A} \subset \mathbb{J} \\ |\mathcal{A}| = F}} \mathcal{A}$. For any $g \in C^b$ and any $s \in S$, we obtain

$$\int g(u) Q_\theta(s; du) = \int g(u) \mathbb{P}(X_\theta(m+1) \in du \mid X_\theta(m) = s)$$

$$= \sum_{\mathcal{J} \in \mathcal{J}^f} \sum_{i=0}^{N_\theta(\mathcal{J})} \int_{U_i(\mathcal{J})}^{U_{i+1}(\mathcal{J})} g(\phi(\mathcal{J}, \mathcal{R}_i(\mathcal{J}), t, s))$$

$$\cdot h(\mathcal{J}, t)(dt). \quad (24)$$

Equation (24) concludes the construction of the Markov kernel Q_θ. We introduce the following assumption.

ASSUMPTION 2 (A2). *The components of $(\theta_n : 1 \leqslant n \leqslant N)$ are mutually distinct, that is, $\theta_i \neq \theta_j$ for $i, j \in \{1, N\}$ with $i \neq j$.*

Note that (A2) imposes no severe restriction to our maintenance system because different types of components have different distributions, i.e., different failure rates which leads to different age thresholds per type.

LEMMA 1. *Assume that* (A2) *holds for* $(\theta_n : 1 \leqslant n \leqslant N)$. *Then, there exists an open neighborhood Θ_0 of θ, such that $N_\theta(\mathcal{J})$ and $R_i(\mathcal{J}), 0 \leqslant i \leqslant N_\theta(\mathcal{J})$, are constant on Θ_0 as a mapping of θ. Moreover, $U_i(\mathcal{J})$ is differentiable for $0 \leqslant i \leqslant N_\theta(\mathcal{J})$ and it holds:*

$$\frac{\partial}{\partial \theta_n} U_i(\mathcal{J}) = \begin{cases} 0, & U_i(\mathcal{J}) \notin V_\theta(n, \mathcal{J}), \\ 1, & U_i(\mathcal{J}) \in V_\theta(n, \mathcal{J}). \end{cases} \quad (25)$$

The proof of this lemma is presented in §EC.4 in the electronic companion. Note that in case (A2) fails, a result similar to that in Lemma 1 holds for one-sided derivatives.

5.2. The Phantom Estimator

So far, we have developed a transition kernel for X_θ such that the kernel is differentiable with respect to θ_n for $1 \leqslant n \leqslant N$; see Lemma 1. In what follows, we will calculate the partial derivatives and represent them as the difference of two suitable transition kernels. Let $\mathcal{I}(n, \mathcal{J})$ denote the *essential-time* index set, which indicates where the partial derivative with respect to θ_n has nonzero value,

$$\mathcal{I}(n, \mathcal{J}) = \{1 \leqslant i \leqslant N_\theta(\mathcal{J}) \colon U_i(\mathcal{J}) \in V(n, \mathcal{J})\}. \tag{26}$$

An example illustrating the construction of $\mathcal{I}(n, \mathcal{J})$ is given in §EC.3 in the electronic companion. For $s = (\tau, t, a_1, \ldots, a_J) \in S$, set

$$c^n(s) = \sum_{\mathcal{J} \in \mathcal{J}^f} \sum_{i \in \mathcal{I}(n, \mathcal{J})} h(\mathcal{J}, U_i(\mathcal{J})). \tag{27}$$

By construction, $\mathcal{I}(n, \mathcal{J})$ and $U_i(\mathcal{J})$ are measurable mappings in s (see (26), and (20) together with (21)), and $c^n(\cdot)$ is thus measurable. We now introduce discrete probabilities on $\bigcup_{\mathcal{J} \in \mathcal{J}^f} \bigcup_{i \in \mathcal{I}(n, \mathcal{J})} \{(\mathcal{J}, i)\}$, by

$$p^n(\mathcal{J}, i) = \frac{h(\mathcal{J}, U_i(\mathcal{J}))}{c^n(s)}. \tag{28}$$

We define $X_\theta^{n+}(s)$ and $X_\theta^{n-}(s)$ as follows. With probability $p^n(\mathcal{J}, i)$, we let

$$X_\theta^{n+}(s) = \phi(\mathcal{J}, \mathcal{R}_{i-1}(\mathcal{J}), U_i(\mathcal{J}), s) \quad \text{and}$$

$$X_\theta^{n-}(s) = \phi(\mathcal{J}, \mathcal{R}_i(\mathcal{J}), U_i(\mathcal{J}), s). \tag{29}$$

The phantom processes are now constructed as described in §4.2, where the initial state of the phantom is given by $X_\theta^{n\pm}(s, 1) = X_\theta^{n\pm}(s)$. With this notation at hand (as shown in §EC.5 in the electronic companion), it holds for any $s \in S$ and $g \in C^b$ that

$$\frac{\partial}{\partial \theta_n} \mathbb{E}_\theta[g(X_\theta(k+1)) \mid X_\theta(k) = s]$$

$$= \mathbb{E}[c_\theta^n(s)(g(X_\theta^{n+}(s)) - g(X_\theta^{n-}(s)))]. \tag{30}$$

Thus, we have obtained an explicit representation of $(\partial/\partial \theta_n) \int g(u) Q_\theta(s; du)$; compare (30) with (9). Introducing the phantom processes for the nth partial derivative $\{X_\theta^{n\pm}(s, k) \colon k \geqslant 1\}$, as described in §4.2, the phantom estimators for $\partial C(T, \theta)/\partial \theta_n$ can be obtained from Equation (13).

We summarize the analysis in the following theorem (for a proof, see §EC.6 in the electronic companion).

THEOREM 1. *Let conditions* (A1) *and* (A2) *be satisfied. Then, it holds for all* $g \in C^b$ *that*

$$\frac{1}{T} \frac{\partial}{\partial \theta_n} \mathbb{E}\left[\sum_{i=1}^{M(T, \theta)-1} g(X_\theta(i))\right] = \frac{1}{T} \mathbb{E}_\theta[PhE^n(g)]$$

$$= \frac{1}{T} \mathbb{E}_\theta[RPhE^n(g)].$$

5.3. Extension of the Phantom Estimator

As explained in §4, we have proposed two gradient estimators: the phantom estimator (PhE) and the randomized phantom estimator (RPhE). PhE uses formula (6) together with (13) and RPhE uses formula (6) together with (14) to calculate the gradient of $C(T, \theta)$. For these estimators, the gradient information is obtained from the scaled difference between the cumulative costs over the plus and minus phantoms. Therefore, for each sample path, PhE needs $2(M(T, \theta) - 1)$ phantoms per simulation run. Simulating the phantoms can be stopped either when the phantoms couple or when the time horizon T is reached. The coupling is achieved if the plus and minus phantom simultaneously hit the set $[0, \infty)^2 \times \{0^J\}$ and the time to this event is denoted by $\gamma^n(s)$; see §EC.1 in the electronic companion for details. If the probability that the plus and minus phantom simultaneously hit the set $[0, \infty)^2 \times \{0^J\}$ decreases, $\gamma^n(s)$ will increase, which implies that many phantoms need to be simultaneously simulated. In such a situation PhE will have a "phantom overflow" problem, leading to a heavy storage and computational burden. In contrast, RPhE only generates two phantoms each run, but at the price of higher variance.

To control the number of phantoms that have to be simulated simultaneously, we combine PhE and RPhE in the following way. For each sample path, we only generate N_{Ph} phantoms, where the number N_{Ph} is predefined by the user. That means we generate two phantoms (plus and minus phantom) per batch of $\lceil (M(T, \theta) - 1)/N_{Ph} \rceil$ failures. Within each batch of failures, we apply the randomized phantom method. Therefore, we will simulate in total $2N_{Ph}$ phantoms instead of $2(M(T, \theta) - 1)$. This allows control of the computational burden and the variance through N_{Ph}.

We call this the *combined phantom estimator* (CPhE), which can be given as follows:

$$\text{CPhE}^n(g) \triangleq \left(\sum_{k=1}^{N_{Ph}-1} C_1 D^n(g, X_\theta(\sigma_1(k)))\right)$$

$$+ C_2 D^n(g, X_\theta(\sigma_2)), \tag{31}$$

where $C_1 = \lceil (M(T, \theta) - 1)/N_{Ph} \rceil$, $C_2 = M(T, \theta) - 1 - C_1(N_{Ph} - 1)$, $\sigma_1(k)$ is an integer uniformly distributed on $\{1 + (k-1)C_1, \ldots, kC_1\}$, and σ_2 is an integer uniformly distributed on $\{1 + (N_{Ph} - 1)C_1, \ldots, C_2 + (N_{Ph} - 1)C_1\}$. The following theorem provides sufficient conditions for unbiasedness of $\text{CPhE}^n(g)$ (for a proof, see §EC.7 in the electronic companion).

THEOREM 2. *Let conditions* (A1) *and* (A2) *be satisfied. Then, it holds for all* $g \in C^b$ *that*

$$\frac{1}{T} \frac{\partial}{\partial \theta_n} \mathbb{E}\left[\sum_{i=1}^{M(T, \theta)-1} g(X_\theta(i))\right] = \frac{1}{T} \mathbb{E}_\theta[CPhE^n(g)].$$

6. Numerical Results

6.1. General Experiment Settings

In the previous section, we have analyzed a general $(J - F + 1)$-out-of-J maintenance system with N-types of components. In this section, we will perform experiments on a system with two types of components. We assume that both types of components have a Weibull-(α, β) distribution, i.e., $F_{\alpha,\beta}(x) = 1 - \exp(-(\alpha x)^\beta)$. We choose the Weibull parameter for both types as $\alpha_1 = 0.8930, \beta_1 = 3$ and $\alpha_2 = 0.8862, \beta_2 = 2$. That is, the lifetime distributions have a long right tail and first moment equal to one. However, type 1 components have higher failure rate than type 2 components. Because the cost incurred for failure intervention C_I is usually higher than the replacement cost per component C_R, we set $C_I = 4$ and $C_R = 1$.

As discussed in §3, we apply the finite difference (FD) method. This leads to the "twin numerical evils" of approximation error (approximating the slope of a curve by its secant) and numerical noise (dividing by small $\Delta\theta$); see Strickland (1993). Any application of the FD method involves a trade-off between these two effects. Thus, we first perform a series of numerical experiments to identify a good setting for the FD estimator.

For $N > 1$, the control threshold parameter θ is an N-dimensional vector, and we need to solve a multivariate optimization problem. In this case, the standard FD approach perturbs each component θ_n of θ one at a time and the two-sided FD approximation of $\nabla C(T, \theta)$ is given by

$$\nabla C(T, \theta) \approx \left(\frac{C(T, \theta + ce_1) - C(T, \theta - ce_1)}{2c} \right.$$
$$\left. \ldots \frac{C(T, \theta + ce_N) - C(T, \theta - ce_N)}{2c} \right), \quad (32)$$

where e_i denotes a vector with a one in the ith place and zeros elsewhere. FD requires $2N$ measurements, which is usually not very efficient. The SP method, as introduced by Spall (1987) and more fully analyzed in Spall (1992), overcomes this drawback. The essential feature of SP is its underlying gradient approximation that requires only two measurements per iteration independent of the dimension N. SP has all elements of θ randomly perturbed together (simultaneously) to obtain two measurements $C(T, \cdot)$, but each component of $\nabla C(T, \theta)$ is formed from a ratio involving the difference in the two corresponding measurements and the individual components in the perturbation vector. For two-sided SP, we have

$$\frac{\partial}{\partial \theta_i} C(T, \theta) \approx \frac{C(T, \theta + c\Delta) - C(T, \theta - c\Delta)}{2c\Delta_i}$$
$$\text{for } 1 \leqslant i \leqslant n, \quad (33)$$

where Δ_i is the ith component of the Δ vector (which may be ± 1 random variables). Note that the common numerator in all N components of $\nabla C(T, \theta)$ reflects the simultaneous perturbation of all components in θ. This provides the potential for large savings in the overall cost of optimization.

6.2. Gradient (Sensitivity) Estimation

As presented in §4.3, if both phantom processes simultaneously hit set $[0, \infty)^2 \times 0^J$, we can terminate simulating the phantoms. That is, simulating the phantoms can be stopped either after $\gamma^n(s)$ transitions or when the time horizon T is reached. We mentioned in §5.3 that PhE will have a "phantom overflow" problem. In the following example, we show how this effects the simulation time of PhE.

EXAMPLE 1. We perform an experiment on a small system, with $N = 2, F = 2, J_1 = 3$, and $J_2 = 3$, and we let $T = 14,000$ time units. We fix $\theta_2 = 0.6$, and the number of replications is 20. The computation time of PhE as a function of θ_1 is depicted in Figure 2. Note that the computation time increases exponentially as θ increases. The main reason for this increase is that the probability decreases that, upon the failure of in total two components $(F = 2)$, the ages of all $J - F = 4$ operational components are larger than θ. The phantoms, therefore, couple too slowly. The phantom estimator requires simulating approximately $\sum_{k=1}^{M(T, \theta)} \min(\gamma(X_\theta(k)), M(T, \theta) - k + 1)$ additional failures; see Figure 1. As a result, if θ increases, the computational burden (i.e., the number of phantoms) also increases which leads to higher simulation time.

As presented in §5.3, PhE has large computational burden while RPhE has high variance. Because these two estimators require different computational efforts, which leads to different variance, it is important to reflect effort appropriately. A standard measure, dating at least to Hammersley and Handscomb (1964) and formalized in a general framework by Glynn and Whitt (1992), compares estimators based on the product of variance and expected effort per run. This is the *work-normalized variance* (WNV).

For the set-up in Example 1, we compare the half confidence interval width, the computation time, the variance, and the WNV of these two estimators in Table 1. PhE has much smaller variance than RPhE, but it takes PhE much more time than RPhE. Nevertheless, PhE has smaller work-normalized variance than RPhE. This indicates that, if PhE is applicable, PhE has better performance than RPhE.

Figure 2. The computation time of PhE for a 2-type 5-out-of-6 system.
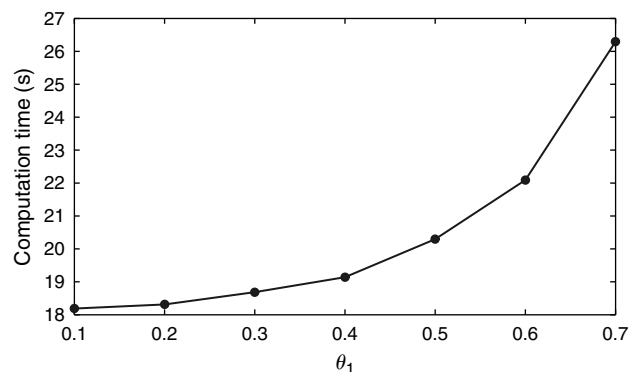
Table 1. Comparison of two different phantom estimators: PhE and RPhE.

Method	$(\partial/\partial\theta_1)C$	Half conf. interval	Total time (s)	Variance	WNV
PhE	−0.447598	0.005477	696.035	0.001561	0.005435
RPhE	−0.458010	0.586107	11.984	17.8843	1.071626

Note. $J_1 = 3, J_2 = 3, \theta_1 = 0.4, \theta_2 = 0.6, T = 14{,}000$, number of runs $= 200$.

For the model in Example 1, RPhE suffers from a high variance and the number of simulation runs has to be increased to obtain a relatively accurate gradient estimation. PhE, on the other hand, has a high computational burden due to the fact that there are too many parallel phantoms needed to be simulated. To control the number of phantoms that have to be simulated simultaneously, we use CPhE; see §5.3. We illustrate the influence of N_{Ph} on the performance of this estimator in the following example.

EXAMPLE 2. In this experiment, a two-type component maintenance system is studied, with $J_1 = 4, J_2 = 6, F = 2$, and the time horizon is set to $T = 14{,}000$ time units. Numerical experiments for different values of N_{Ph} for CPhE show that with increasing N_{Ph} the WNV of the combined phantom estimator is decreasing. (The detailed results are presented in tabular form in §EC.8 of the electronic companion.) Therefore, in practice there is a trade-off between low variance and short computation time for large maintenance systems. The choice of the value of N_{Ph} is up to the user.

Example 2 suggests the following rule of thumb. If the computation time underlies no restriction, one can choose a large value for N_{Ph} to obtain a small variance of the gradient estimator. On the other hand, if the computational burden becomes an issue, one can decrease the value of N_{Ph} to obtain the best gradient estimation within a given restriction of computation time and storage limitation. Determining a priori the optimal setting of N_{Ph} is a topic of further research.

For the setup of Example 2, we compare the WNV of CPhE, SP, and FD in Figure 3. To apply FD and SP, we have first performed a series of experiments and determined $c = 0.001$ as an appropriate choice in (32) and (33). For CPhE, we choose $N_{Ph} = 200$, which has relatively low computational burden while the WNV is relatively high. As can be seen from the figure, even for such a setting, the CPhE estimator outperforms the other methods. On average, the WNV of FD, two-sided FD, and SP is, respectively, 10, 6, and 7 times that of CPhE with $N_{Ph} = 200$.

It is worth noting that for estimating a single partial derivative, the phantom estimator is more efficient than the SP method, but as the dimension of θ increases, this advantage will be compromised. The reason is that the phantom estimator requires N estimates for the gradient, whereas SP only requires one simulation for all the gradient

Figure 3. Comparison of the WNV of four different estimators for a 2-type 5-out-of-6 maintenance system.
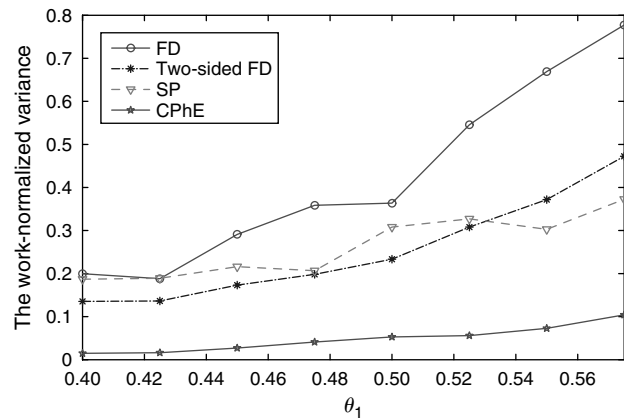

estimations. However, phantom estimators provide the possibility for parallel computing because the plus and minus phantoms can be simulated independently from the nominal phantom. To explore the speed-up factor by using parallel computing in gradient estimation is a topic of further research.

6.3. Optimization

In this section, we perform a few simulation experiments illustrating the use of the phantom estimator(s) for optimization in the case of a 10-component maintenance system. We use the SA iteration algorithm in Equation (2). For SA, the parameter c in Equations (32) and (33) needs to be updated per iteration, and the value of c will be denoted by c_k for the kth iteration. The choice of the gain sequences (a_k in Equation (2) and c_k) is critical to the performance, and we follow the practical suggestions for choosing these parameters as proposed by Spall (1998).

EXAMPLE 3. We first perform an optimization for a single-type 10-out-of-10 maintenance system; i.e., $N = 1, F = 1$,

Figure 4. The SA iteration trace in the optimization of a single-type 10-out-of-10 maintenance system.
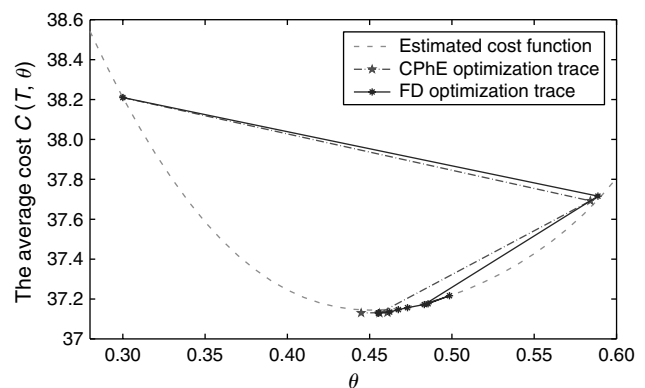
Table 2. The optimization performance comparison between FD and CPhE.

	θ^*	$C(T, \theta^*)$	No. of iterations	Computation time (s)
FD	0.455579	37.130482	10	331.578
CPhE	0.444829	37.130362	5	163.578

Table 3. The optimization performance comparison between SP and CPhE.

	(θ_1^*, θ_2^*)	$C(T, \theta^*)$	No. of iterations	Computation time (s)
SP	(0.547110, 0.674446)	23.818845	17	1459.576
CPhE	(0.536216, 0.687814)	23.816568	9	745.25

and the horizon $T = 20{,}000$. Because $N = 1$, one can just use FD to approximate the gradient information rather than SP. The number of simulation runs in each iteration is chosen as 20. By following the suggestions presented in Spall (1998), we choose $a = 0.02$ and $a_k = a/(k + 1)$. We terminate the algorithm when the difference between two successive iterations is smaller than 0.005 in three successive iterates; i.e., $|C(T, \theta_{k+1}) - C(T, \theta_k)| < 0.005$. We use CPhE with $N_{Ph} = 1{,}000$ to estimate the gradient information. The iteration traces of both methods during the SA optimization of a 10-component system are shown in Figure 4.

The best solutions of both methods are shown in Table 2. FD takes 10 iterations to obtain the optimal result of $C(T, \theta^*) = 37.130482$ with $\theta^* = 0.455579$, whereas CPhE only needs 5 iterations to obtain the optimal result of $C(T, \theta^*) = 37.130362$ with $\theta^* = 0.444829$. Note that the fitted curve of $C(T, \theta)$ as a function of θ is rather flat on $[0.4, 0.5]$; see Figure 4. This explains why FD and CPhE yield similar values for $C(T, \theta^*)$ with different values for θ^*. The algorithm using CPhE converges around 2 times as fast as the one using the FD estimator.

Next, we optimize 2-type maintenance systems. In these cases, we use the SP method to approximate the gradient.

EXAMPLE 4. We test a 2-type 9-out-of-10 maintenance system. We use CPhE with $N_{Ph} = 100$ to estimate the gradient information, and the number of runs per iteration is 20. The value of parameters are chosen as $a = 0.165$, $A = 2$, $\alpha = 0.602$, $c = 0.001$, and $\gamma = 0.101$. The algorithm evolves as shown in Figure 5. The best solutions of both methods are shown in Table 3. The SA algorithm using SP takes 17 iterations converging to the optimal value, whereas the one using CPhE only takes 9 iterations. Furthermore, SP consumes almost twice the CPU time as CPhE to reach the best solution.

7. Conclusion and Further Work

In this paper, we have established unbiased gradient estimators based on measure-valued differentiation for the derivative of the average cost for a multicomponent system with a general lifetime distribution with respect to the parameter of an age-replacement policy. On a theoretical level, the paper presents the first gradient estimator for a multicomponent system with modified F-group failure replacement.

Figure 5. The SA iteration trace in the optimization of a 2-type 9-out-of-10 maintenance system.
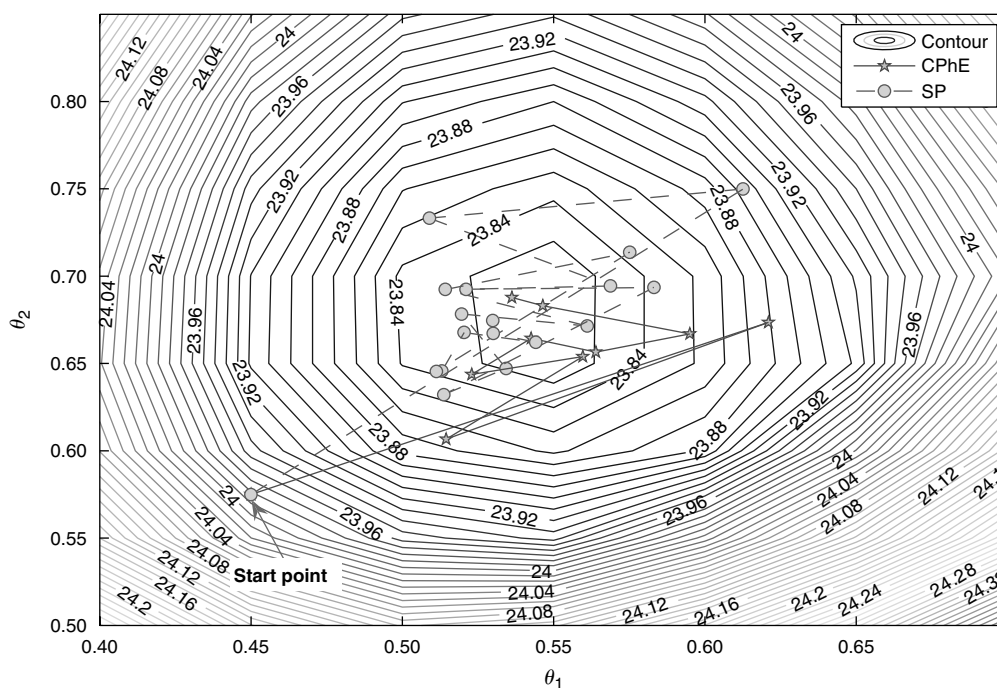
On a practical level, we derive gradient estimators tailored to situations typically encountered in practice, where the best gradient estimation has to be found within a fixed simulation budget.

On a more general note, the study of this paper exemplifies the difficulty of comparison between different methods. We believe that our study highlights the need for a framework for comparing gradient-based and gradient-free methods as initiated by Pasupathy and Henderson (2006).

The study of maintenance systems with highly reliable components has not been addressed in this paper. Simulations of highly reliable systems take a long time to execute because failures happen so rarely that gathering sufficient failure statistics is extremely slow. For such highly reliable systems, rare-event simulation techniques (e.g., importance sampling) need to be applied. By following the analysis presented in §5.1, one can derive a differential Markov kernel for such a system and then obtain the phantom estimator. However, a detailed analysis is a topic of further research.

Phantom estimators obtain gradient information from the difference between phantom processes. The shorter the phantom processes need to be simulated, the more efficient the phantom estimator becomes. How to find an efficient coupling for the phantoms to improve the performance is a topic of further research.

8. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at http://or.journal.informs.org/.

References

Billingsley, P. 1966. *Weak Convergence of Probability Measures.* John Wiley & Sons, New York.

Brémaud, P., F. J. 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 Appl.* **10**(3) 249–270.

Dekker, R., J. R. van der Meer, R. Ph. Plasmeijer, R. E. Wildeman, J. de Bruin. 1998. Maintenance of light-standards—A case-study. *J. Oper. Res. Soc.* **49** 132–143.

Fu, M. C., J. Q. Hu. 1997. *Conditional Monte Carlo: Gradient Estimation and Optimization Applications.* Kluwer Academic Publishers, Boston.

Fu, M. C., J. Q. Hu, L. Shi. 1993. An application of perturbation analysis to a replacement problem in maintenance theory. *25th Winter Simulation Conf.*, ACM Press, Los Angeles, 329–337.

Glasserman, P. 1988. Performance continuity and differentiability in Monte Carlo optimization. *20th Winter Simulation Conf.*, ACM, San Diego, 518–524.

Glasserman, P. 1991a. *Gradient Estimation via Perturbation Analysis.* Kluwer Academic Publishers, Boston.

Glasserman, P. 1991b. Structural conditions for perturbation analysis of queueing systems. *J. Assoc. Comput. Machinery* **38**(4) 1005–1025.

Glasserman, P. 1991c. Structural conditions for perturbation analysis derivative estimation: Finite-time performance indices. *Oper. Res.* **39**(5) 724–738.

Glasserman, P., W. Gong. 1990. Smoothed perturbation analysis for a class of discrete-event systems. *IEEE Trans. Automatic Control* **35**(11) 1218–1230.

Glynn, P. W., W. Whitt. 1992. The asymptotic efficiency of simulation estimators. *Oper. Res.* **40**(3) 505–520.

Gong, W., Y. Ho. 1987. Smoothed (conditional) perturbation analysis of discrete event dynamical systems. *IEEE Trans. Automatic Control* **32**(10) 858–866.

Hammersley, J. M., D. C. Handscomb. 1964. *Monte Carlo Methods.* Methuen & Co., London, and John Wiley & Sons, New York.

Haurie, A., P. L'Ecuyer. 1986. Approximation and bounds in discrete event dynamic programming. *IEEE Trans. Automatic Control* **31**(3) 227–235.

Heidelberger, P., X. Cao, M. A. Zazanis, R. Suri. 1988. Convergence properties of infinitesimal perturbation analysis estimates. *Management Sci.* **34**(11) 1281–1302.

Heidergott, B. 2001. A weak derivative approach to optimization of threshold parameters in a multicomponent maintenance system. *J. Appl. Probab.* **38**(2) 386–406.

Heidergott, B., T. Farenhorst-Yuan. 2007. A note on gradient estimation for maintenance systems. *IEEE Trans. Automatic Control* **52** 2141–2145.

Heidergott, B., F. Vázquez-Abad. 2006. Measure-valued differentiation for random horizon problems. *Markov Processes and Related Fields* **12** 509–536.

Heidergott, B., F. Vázquez-Abad. 2008. Measure-valued differentiation for Markov chains. *J. Optim. Appl.* **136**(2) 187–209.

Heidergott, B., T. Farenhorst-Yuan, F. Vázquez-Abad. 2010a. A perturbation analysis approach to phantom estimators for waiting times in the G/G/1 queue. *Discrete Event Dynamic Systems: Theory Appl.* Forthcoming.

Heidergott, B., A. Hordijk, H. Weisshaupt. 2006. Measure-valued differentiation for stationary Markov chains. *Math. Oper. Res.* **31**(1) 154–172.

Heidergott, B., A. Hordijk, H. Weisshaupt. 2008. Derivatives of Markov kernels and their Jordan decomposition. *J. Appl. Anal.* **14**(1) 13–26.

Heidergott, B., F. Vázquez-Abad, G. Pflug, T. Farenhorst-Yuan. 2010b. Gradient estimation for discrete event systems by measure-valued differentiation. *ACM Trans. Modeling Comput. Simulation.* Forthcoming.

Ho, Y. C., X. Cao. 1991. *Perturbation Analysis of Discrete Event Dynamic Systems.* Kluwer Academic Publishers, Boston.

Kushner, H., D. Clark. 1978. *Stochastic Approximation for Constrained and Unconstrained Systems.* Springer-Verlag, Berlin and New York.

L'Ecuyer, P. 1983. Processus de décision Markoviens á étapes discrètes: Application á des problèmes de remplacement d'équipment. Ph.D. thesis, Université de Montréal, Montréal, Quebec.

L'Ecuyer, P. 1993. Two approaches for estimating the gradient in functional form. *25th Winter Simulation Conf.*, ACM Press, Los Angeles, 338–346.

L'Ecuyer, P., F. Vázquez-Abad. 1997. Functional estimation with respect to a threshold parameter via dynamic split and merge. *J. Discrete Event Dynam. Systems* **7**(1) 69–92.

L'Ecuyer, P., F. Vázquez-Abad, B. Martin. 1999. Functional estimation for a multicomponent age replacement model. *Amer. J. Math. Management Sci.* **19**(1–2) 135–156.

Nicolai, R., R. Dekker. 2006. Optimal maintenance of multi-component systems: A review. *Complex System Maintenance Handbook.* Springer, London, 263–286.

Pasupathy, R., S. Henderson. 2006. A testbed of simulation-optimization problems. *38th Winter Simulation Conf.*, Monterey, CA, 255–263.

Pflug, G. 1996. *Optimization of Stochastic Models.* Kluwer Academic Publishers, Norwell, MA.

Rubinstein, R. 1992. Sensitivity analysis of discrete event systems by the push out method. *Ann. Oper. Res.* **39**(1–4) 229–250.

Rubinstein, R. 1998. *Modern Simulation and Modeling.* Kluwer Academic Publishers, Boston.

Spall, J. C. 1987. A stochastic approximation technique for generating maximum likelihood parameter estimates. *Proc. Amer. Control Conf.*, Minneapolis, 1161–1167.

Spall, J. C. 1992. Multivariate stochastic approximation using a simultaneous perturbation gradient approximation. *IEEE Trans. Automatic Control* **37**(3) 332–341.

Spall, J. C. 1998. Implementation of the simultaneous perturbation algorithm for stochastic optimization. *IEEE Trans. Aerospace Electronic Systems* **34** 817–823.

Strickland, S. G. 1993. Gradient/sensitivity estimation in discrete-event simulation. *25th Winter Simulation Conf.*, ACM Press, Los Angeles, 97–105.

Suri, R. 1987. Infinitesimal perturbation analysis for general discrete event systems. *J. Assoc. Comput. Machinery* **34**(3) 686–717.

Suri, R. 1989. Perturbation analysis: The state of the art and research issues explained via the GI/G/l queue. *Proc. IEEE* **77**(1) 114–137.

Van der Duyn Schouten, F., B. Van Vlijmen, S. Vos de Wael. 1998. Replacement policies for traffic control signals. *IMA J. Management Math.* **9**(4) 325–346.

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

Vázquez-Abad, F. J., P. L'Ecuyer, B. Martin. 1996. On the linear growth of the split-and-merge simulation tree for a multicomponent age replacement model. *WODES'96: Internat. Workshop on Discrete Event Systems*. Edinburgh, Scotland, 57–62.