FUNCTIONAL ESTIMATION FOR A MULTICOMPONENT AGE REPLACEMENT MODEL

Pierre L'Ecuyer, Benoit Martin, and Felisa J. Vázquez-Abad
Département d'informatique et de recherche opérationnelle
Université de Montréal, C.P. 6128, Succ. Centre-Ville
Montréal, H3C 3J7, CANADA
lecuyer@iro.umontreal.ca vazquez@iro.umontreal.ca

SYNOPTIC ABSTRACT

We consider a multicomponent system with identical components subject to failures, controlled by a replacement rule based on age thresholds. We show how to estimate the expected cost of operating the system as a function of one of the thresholds, in functional form, by simulation. The efficiency of the approach is analyzed for two types of performance measures of the system: (i) The total expected cost over a finite time horizon and (ii) the average cost over an infinite time horizon. For the average cost case, we show under mild conditions that the system is geometrically ergodic, uniformly over the parameter set of interest. Our numerical illustrations suggest that the proposed approach is viable from the practical viewpoint.

Key Words and Phrases: functional estimation, split-and-merge tree, regenerative simulation, maintenance models, age replacement policies.

1. INTRODUCTION

We consider a system with M identical components subject to failures. Each time a component fails it must be replaced immediately by a new one. Components can also be replaced preventively, at any time. The component lifetimes are independent and identically distributed (i.i.d.) and the replacements are assumed instantaneous. Each intervention entails a fixed cost c_1, plus a variable cost c_2 for each component replaced. An additional cost c_3 is incurred if the intervention is forced by a failure.

The objective is to minimize the expected total operating cost of the system, either over a finite time horizon, or as a time-average over an infinite horizon. For this model, we define the *state* of the system as a vector that contains the ages of all the components, plus the current time if the horizon is finite. A (general) control policy is a rule which, for each state, tells us which decision to make: to stop the system for a preventive intervention or to let it go, and which component(s) to replace when the system is stopped. This gives rise to an extremely rich class of policies. This model was studied by L'Ecuyer (1983) and Haurie and L'Ecuyer (1986), where the optimal policy was computed via dynamic programming combined with finite element approximation. The criterion considered there was the total expected discounted cost over an infinite time horizon. That computational approach is practical only for a small number of components (say, $M \leq 6$ or so), and the optimal policy turns out to be extremely complicated, in general. For example, for $M \geq 3$, the optimal policy is not necessarily monotone with respect to the component ages.

As a more practical alternative, L'Ecuyer (1983) (see also L'Ecuyer 1990) proposed to restrict the control to a class of (very simple) suboptimal policies determined by two thresholds $0 < \theta < \bar{\theta}$, as follows: Whenever a component fails or reaches the age threshold $\bar{\theta}$, an intervention is performed and all the components whose age exceeds θ are replaced preventively, in addition to the failed component (if any). Some numerical experiments by L'Ecuyer (1983) with three-component systems showed (for the examples he examined) that the best policy within that class of two-threshold rules was almost as good as the optimal one. Then comes the issue of how to optimize the two thresholds. For example, for a fixed value of $\bar{\theta}$, frequent replacements (a small θ) tend to reduce failure costs but to increase the fixed costs of the interventions as well as the replacement costs. Infrequent replacements (a large θ) do the opposite. For more about replacement models, the reader may consult, e.g., Gertsbakh (1977), Pierskalla and Voelker (1976), and Cho and Parlar (1991).

In this paper we focus on estimating the expected cost (objective function) as a function of θ, in functional form over some interval, for a fixed value of $\bar{\theta}$. For this, we adapt the *split-and-merge* approach, proposed and studied by L'Ecuyer and Vázquez-Abad (1997), and designed to estimate a mathematical expectation as a function of a threshold parameter over a given interval. The main idea of the method is as follows. Up to the time of the first intervention, a single simulation is valid for the entire interval of values of θ, because what happens is independent of θ. The first time when a replacement decision is not the same for all values of θ in the interval of interest, this interval is split at the point(s) where the decision changes as a function of θ, and the simulation is continued separately in each of the subintervals, with a different sample trajectory for each subinterval. This continues recursively: each subinterval is split whenever a replacement decision differs for different values of θ in this subinterval, and so on. This gives rise to a growing *simulation tree*, where each node of the tree corresponds to an ongoing simulation. Merging can also occur, when adjacent nodes in the tree reach a common regeneration point (e.g., all components are replaced by new ones). A functional estimator is obtained when the time horizon is reached in all nodes of the tree. In the infinite-horizon time-average case, we use a regenerative approach. By averaging independent replications of this procedure, a functional estimator with arbitrary precision can be obtained. This can be used to estimate the optimum, or for other purposes.

For stochastic systems where all the decisions are binary and depend on the system's evolution only via a threshold parameter θ, L'Ecuyer and Vázquez-Abad (1997) have shown that the simulation tree in the split-and-merge method grows only linearly as a function of the number of decisions taken during the simulation. These results do not apply directly to the age replacement problem considered here, because the decisions made at the interventions times are not binary and depend on the component lifetimes. Here, at any intervention, any tree node can be split into any number of children nodes from 1 to M. Conceivably, in the worst-case scenario, the tree could grow exponentially fast.

In Section 2, we study the growth behavior of the tree in *expectation*, as a function of the number of interventions n and of the number of components M, both empirically and theoretically. We prove a bound on the expected number of leaves which is *linear* in n and *quadratic* in M. Our empirical experiments suggest that the growth is typically *linear* in both n and M. Since the expected number of transitions is linear in the time horizon T, the total expected computational effort to obtain the functional estimator over some interval is of order $\mathcal{O}(Mn^2) = \mathcal{O}(MT^2)$. In Section 4, we consider the steady-state average cost over an infinite horizon and construct a functional estimator based on a regenerative analysis. Under mild assumptions, we show that the system is geometrically ergodic, uniformly in θ. It follows that the times between the successive regeneration epochs have bounded moments of all orders, uniformly in θ. We discuss implementation and performance issues associated with the regenerative approach. The empirical results are from Martin (1996), where other numerical examples and further discussion can be found. A short preliminary draft of this paper is given by Vázquez-Abad, L'Ecuyer, and Martin (1996).

Another approach to functional estimation by simulation is based on the use of a likelihood ratio to estimate the function in an entire interval from a single simulation. This applies when θ is a parameter of the probability distribution of some random variables that drive the system and when a dominating measure, independent of θ, can be found. For details, see, e.g., L'Ecuyer (1993), Rubinstein and Shapiro (1993). This does not apply to our problem (at least not directly), because θ is not a parameter of a probability distribution.

A very simple functional estimation method which does apply in our case is to perform parallel simulations over a fixed grid of values of θ. If the grid is fine enough, the discretization error can be made negligible, assuming that the expected cost is a continuous function of θ. This may be simpler to implement than the proposed approach, but the latter has the following advantages: (i) There is no discretization error and (ii) we avoid repeating several times the same simulation in the early part, where the sample paths are exactly the same at many adjacent grid points.

2. THE FINITE-HORIZON SETUP

2.1. The Model. In this section, we focus on the case where the objective function is the total expected cost over a finite time horizon T. The model is a stochastic process defined over some probability space (Ω, Σ, P) independent of the threshold parameters, where a sample point $\omega \in \Omega$ can be viewed as the successive lifetimes of all the components. The evolution of the system depends on $(\theta, \bar{\theta})$ only through the replacement decisions. The component lifetimes are i.i.d. with distribution G, density g, reliability function $\bar{G}$ defined by $\bar{G}(t) = 1 - G(t)$, and failure rate function λ defined by $\lambda(t) = g(t)/\bar{G}(t)$, for $t \geq 0$. We assume that the failure rate is bounded over the interval $[0, \bar{\theta}]$; that is, $\lambda(t) \leq \bar{\lambda}$ for some constant $\bar{\lambda}$.

For a given value of $(\theta, \bar{\theta})$, for $n = 0, 1, 2, \ldots$, let $T_n = T_n(\theta)$ be the time of the nth intervention, and $Z_n = Z_n(\theta) = (Z_{n,1}, \ldots, Z_{n,M})$ be the M-dimensional age vector of the components at time T_n, immediately after the replacement(s). Each vector $Z_n = (Z_{n,1}, \ldots, Z_{n,M})$ has all its components in the interval $[0, \theta)$ and at least one component is zero. We take $T_0 = 0$ and assume that initially, all the components are new, so $Z_0 = (0, \ldots, 0)$. The process $\{(T_n, Z_n), n \geq 0\}$

is a Markov renewal process, as explained in Haurie and L'Ecuyer (1986), and its evolution depends on $(\theta, \bar{\theta})$ via the decisions.

The total cost for the first N interventions is the random variable

$$C_N(\theta) = \sum_{n=1}^{N} \left[c_1 + c_2\rho_n + c_3\delta_n\right],$$

where ρ_n is the number of components replaced at intervention n, and δ_n is the indicator function that the nth intervention was forced by a failure. The parameter θ indicates that $C_N(\theta)$ varies with θ for every fixed sample point ω (and fixed $\bar{\theta}$). In other words, $C_N(\cdot)$ is a random function. For a fixed time horizon T, let $N(T, \theta)$ denotes the (random) number of interventions until time T. The total expected cost is then $\mathcal{C}(T, \theta) = E[C_{N(T,\theta)}(\theta)]$.

2.2. Functional Estimation via split-and-merge. For a given sample point ω, the random function $C_{N(T,\cdot)}(\cdot)$ is an unbiased functional estimator of $\mathcal{C}(T, \cdot)$, over any interval $[a, b)$ of interest, where $0 \le a < b \le \bar{\theta}$. Let $\tilde{C}_1(\cdot), \ldots, \tilde{C}_R(\cdot)$ be R i.i.d. copies of $C_{N(T,\cdot)}(\cdot)$ and define

$$\hat{C}(\theta) = \frac{1}{R}\sum_{i=1}^{R} \tilde{C}_i(\theta), \qquad \theta \in [a, b).$$

We already outlined how to compute the sample function $C_{N(T,\cdot)}(\cdot)$ via a split-and-merge simulation tree. We now give more detailed explanations. This computation uses *common random numbers*, in the sense that the same set of component lifetimes is used for all values of θ. This has no impact on the variance of $\hat{C}$, but induces a positive correlation between the values of $\hat{C}(\theta)$, which reduces the "variation" of the function $\hat{C}$ (note that $\hat{C}$ is piecewise constant). The correlation between $\hat{C}(\theta_1)$ and $\hat{C}(\theta_2)$ depends on both θ_1 and θ_2, and converges to 1 as $|\theta_1 - \theta_2| \to 0$ (the proof is highly technical and is omitted; the arguments are similar to those used in Proposition 3 of L'Ecuyer and Perron 1994).

The procedure for computing $C_{N(T,\cdot)}(\cdot)$ is represented by a simulation tree, where each node of the tree corresponds to a simulation valid for a particular subinterval of $[a, b)$. At each *level* of the tree, the interval $[a, b)$ is partitioned into subintervals and the nodes at that level are in one-to-one correspondence with the subintervals of the partition. For $n \geq 0$, let X_n be the number of nodes at level n, and label these nodes (n, i), $i = 1, \ldots, X_n$. Node (n, i) corresponds to some interval $(a_n(i-1), a_n(i)]$, where $a_n(0) = a$ and $a_n(X_n) = b$.

At time 0, we generate M component lifetimes and determine the time T_1 of the first intervention, which is the minimum between $\bar{\theta}$ and the shortest component lifetime. Up to time T_1^- (i.e., time T_1, immediately before the intervention), the evolution is independent of θ and a single simulation is thus valid for all $\theta \in [a, b)$. This is level 0 of the simulation tree. At this level, we have a single node ($X_0 = 1$), labeled $(0, 1)$. When the first intervention takes place, at least one component is replaced and all the remaining components have lived $\tau_0 = T_1$ units of time. If $\tau_0 \geq b$, we replace all components and the process starts afresh. If $\tau_0 \leq a$, this is a forced intervention and we replace only the failed component. In these two situations, the system's evolution is again independent of $\theta \in [a, b)$ until the next intervention time T_2, so there is still a single node at level 1 of the tree, i.e., $X_1 = 1$. But if $\tau_0 \in (a, b)$, this interval splits in two subintervals $(a, \tau_0]$ and $(\tau_0, b]$, giving rise to nodes $(1, 1)$ and $(1, 2)$ at level 1. For $\theta \in (a, \tau_0]$, all components are replaced at time T_1, while for $\theta \in (\tau_0, b]$, only the failed component is replaced. In this case we have $X_1 = 2$ and the trajectory of the Markov renewal process differs, at these two nodes, from time T_1 onwards. From that time, we start to maintain two "parallel" simulations. For node $(1, 1)$, where all components are replaced, we generate new lifetimes, determine the time $T_1 + \tau_1(1)$ of the next intervention, where $\tau_1(1)$ is the minimum between $\bar{\theta}$ and the smallest lifetime, and proceed as we did above, with $(a, \tau_0]$ instead of $[a, b)$. For node $(1, 2)$, the time $\tau_1(2)$ until the next intervention is the minimum between the lifetime of the new component, the residual lifetimes of the other components (of age τ_0), and $\bar{\theta} - \tau_0$. Each of these subintervals (or nodes) may then split again, at level 2, if the age of one or more component(s), immediately before the intervention, lies in the subinterval.

In general, node (n, i) at level n corresponds to the nth intervention time, for θ in the subinterval $(a_n(i-1),\ a_n(i)]$. Let $T_n(i)$ denote the nth intervention time for this interval, $Z_n(i) = (Z_{n,1}(i), \ldots, Z_{n,M}(i))$ the age vector immediately after the intervention, and $\tau_n(i)$ the time until the next intervention. In our implementation, a new lifetime is generated at level n for each component that is replaced at one (at least) of the level-n nodes. If the same component is replaced for several nodes at level n, the same lifetime is used across all these nodes for this component. Whenever a new node (n, i) is created, we use the new lifetimes generated for that level for the components that are replaced and compute the value of $\tau_n(i)$ for this node. If $\zeta_n(i)$ is the minimum among the new component lifetimes and the residual lifetimes of the non-replaced (old) components, then $\tau_n(i) = \min\{\zeta_n(i),\ \bar\theta - Z_{n,1}(i), \ldots,\ \bar\theta - Z_{n,M}(i)\}$. The age of component m at the next intervention, just before the replacements, will be $\tilde Z_{n,m}(i) = \infty$ (by convention) if the next intervention is triggered by a failure of this component, and $\tilde Z_{n,m}(i) = Z_{n,m}(i) + \tau_n(i)$ otherwise. The number of splits of the subinterval $(a_n(i-1),\ a_n(i)]$ at level $n+1$ will be equal to the number of *different* values of $\tilde Z_{n,m}(i)$ lying in the open interval $(a_n(i-1),\ a_n(i))$, and these values will determine the partition of the subinterval. The number of components replaced at intervention $n+1$ will be the same for all θ within each of these new subintervals determined by the partition, but will differ between the subintervals (so the process trajectory will start to differ, from that point on, between these subintervals). For the components whose value of $\tilde Z_{n,m}(i)$ is outside this interval, the (replacement) decision at the next intervention time is the same for all values of θ in the interval.

When the current time $T_n(i)$ has reached the simulation horizon T for each leaf (node without current successor) of the tree, the function $C_{N(T,\cdot)}(\cdot)$ becomes available. This function is piecewise constant (constant over each subinterval of the partition of $[a, b)$ determined by the leaves of the tree).

Since the number of splittings of node (n, i) can be (in general) anywhere in the range 0 to $M-1$, the size of the simulation tree could, in the worst case, increase exponentially fast as a function of n. Fortunately, it turns out that $E[X_n]$ increases only *linearly* in n (see the next subsection). As a consequence, the expected computational effort to obtain the functional estimator $\hat{C}(\theta)$ over the interval $[a, b)$ is $\mathcal{O}(RT^2)$ for R replications.

2.3. Bounding the Expected Tree Width. We now show that for $b < \bar{\theta}/2$, $E[X_n]$ grows only linearly in n, where X_n (defined in Section 2.2) is the number of nodes (or tree width) at level n. This is done via a recurrence formula using conditional expectations. We believe that the result also holds for $\bar{\theta}/2 \le b < \bar{\theta}$, but we do not have a formal proof.

Proposition 1. *Under the assumptions of our model, for any* $0 < a < b < \bar{\theta}/2$, *and all* $n \ge 0$, *one has*

$$E[X_n] \le 1 + (b - a)(M - 1)M\bar{\lambda}n.$$

Proof. Let $\xi_n(i)$ be the number of splits of the subinterval $(a_n(i-1),\ a_n(i)]$ at level $n+1$. For all $i > X_n$, define $a_n(i) = a_n(X_n)$ and $\xi_n(i) = 0$. Observe that

$$X_{n+1} = X_n + \sum_{i=1}^{X_n} \xi_n(i) = X_n + \sum_{i=1}^{\infty} \xi_n(i).$$

Thus, it suffices to show that

$$\begin{aligned} & E[E[\xi_n(i) \mid X_n, a_n(i-1), a_n(i), Z_n(i)]] \\ \le\ & E[(M-1)M(a_n(i) - a_n(i-1))\bar{\lambda}] \end{aligned} \tag{2}$$

for all (n, i), because this would imply

$$\begin{aligned} E[X_{n+1}] &= E[X_n] + E\left[\sum_{i=1}^{\infty} \xi_n(i)\right] \\ &= E[X_n] + E\left[\sum_{i=1}^{\infty} E[\xi_n(i) \mid X_n, a_n(i-1), a_n(i), Z_n(i)]\right] \\ &\le E[X_n] + E\left[\sum_{i=1}^{\infty} (M-1)M(a_n(i) - a_n(i-1))\bar{\lambda}\right] \\ &= E[X_n] + (b-a)(M-1)M\bar{\lambda}, \end{aligned}$$

which, together with the fact that $X_0 = 1$, implies the result. (The sum and expectation in the above display can be interchanged by Fubini's theorem, because all the summands are non-negative.) To prove (2), denote $Z_n(i) = (z_1, \ldots, z_M)$ and $(a_n(i-1), a_n(i)) = (a_0, a_1)$ to simplify the notation, and assume, without loss of generality, that $z_1 \geq z_m$ for all $m > 1$. Note that the (conditional) lifetime density of the system is the product of the (conditional) system's total failure rate and reliability function. That is, at node (n, i), the residual lifetime ζ of the system has conditional density

$$f(\zeta \mid Z_n(i)) = \left(\prod_{m=1}^{M} \bar{G}(z_m + \zeta \mid z_m)\right) \sum_{m=1}^{M} \lambda(z_m + \zeta),$$

for $\zeta > 0$, where $\bar{G}(z_m + \zeta \mid z_m) = \bar{G}(z_m + \zeta)/\bar{G}(z_m)$ is the conditional reliability function of component m. The expectation of $\xi_n(i)$ is bounded above by the expected number of components m for which $z_m + \zeta$ belongs to the interval (a_0, a_1). This gives the following inequalities, explained afterwards:

$$\begin{aligned}
& E[\xi_n(i) \mid a_0, a_1, z_1, \ldots, z_M) \\
\leq\ & \int_0^{\bar{\theta} - z_1} \prod_{m=1}^{M} \bar{G}(z_m + \zeta \mid z_m) \sum_{j=1}^{M} \lambda(z_j + \zeta) \sum_{m=1,\, m \neq j}^{M} I[a_0 < z_m + \zeta < a_1]\, d\zeta \\
& + \prod_{m=1}^{M} \bar{G}(z_m + \bar{\theta} - z_1 \mid z_m) \sum_{m=2}^{M} I[a_0 < z_m + \bar{\theta} - z_1 < a_1] \qquad (3) \\
\leq\ & (M-1) M \bar{\lambda} (a_1 - a_0) + \sum_{m=2}^{M} I[a_0 < z_m + \bar{\theta} - z_1 < a_1], \qquad (4)
\end{aligned}$$

where I denotes the indicator function. The first term (the integral) in (3) is with respect to the density of a failure before the next planned intervention, in ζ units of time, and the other term in (3) corresponds to the event of a preventive intervention because component 1 has reached age $\bar{\theta}$. The sums of indicators in (3) represent the number of components whose age at the next intervention time lies in (a_0, a_1). In the second inequality, we use the facts that the conditional density of ζ never exceeds $M\bar{\lambda}$ and that each indicator function in (3) is nonzero only in an interval of values of ζ of length at most $a_1 - a_0$.

If $b \leq \bar{\theta}/2$, then $a_1 \leq \bar{\theta}/2$ and $z_1 \leq b < \bar{\theta}/2$, so the sum in (4) vanishes and the result follows. □

3. SIMULATION RESULTS FOR THE FINITE-HORIZON CASE

We now illustrate with a numerical example the method and the growth rate of the simulation tree as a function of n and M. The objective here is to estimate the function $\mathcal{C}(T, \cdot)$ over the entire interval (a, b). We want to know how viable and efficient is our proposed functional estimation method for doing that. For all our experiments, the model is the one defined in Sections 1 and 2.1. The components lifetime distribution is $G(t) = 1 - e^{-t^2/100}$, i.e., a Weibull distribution with linearly increasing failure rate $\lambda(t) = 0.02t$, $t > 0$ (as in Haurie and L'Ecuyer 1986). We take $\bar{\theta} = 13$.

Figure 1 shows the results of the functional estimation for an example with the parameter values: $M = 3$, $c_1 = c_2 = 1, c_3 = 2$, $T = 200$, and $R = 500$. The solid line represents the value of $\hat{C}(\theta)$, over the interval $[a, b) = [0, \bar{\theta}) = [0, 13)$. The dotted lines are located 1.64 empirical standard deviations above and below the empirical mean and indicate a *pointwise* 90% confidence interval for our estimator, valid at any given *single* value of θ. (For a true confidence interval for the function, we would need an approximation of the distribution of the random variable $W = \sup_{\theta \in [a,b]} |\hat{C}(\theta) - C(\theta)|$.) As indicated from the figure, the optimal value of θ is around 4.0, although assessing the error of this estimate of the optimizer, or computing a confidence interval for the true optimizer, is beyond the scope of this paper.

The next figure illustrates how the tree grows as a function of M and n. It gives the evolution of $X_{n+1} - X_n$ as a function of n, for three different numbers of components: $M = 5$ (lower), $M = 10$ (middle), and $M = 20$ (upper). Each of these plots was obtained by running $R = 200$ independent runs (simulation trees) over the entire interval $(0, \bar{\theta})$. Note that after some nodes of the tree reach time T (not shown in Figure 2; i.e., for $n > 160$ in this case), the growth

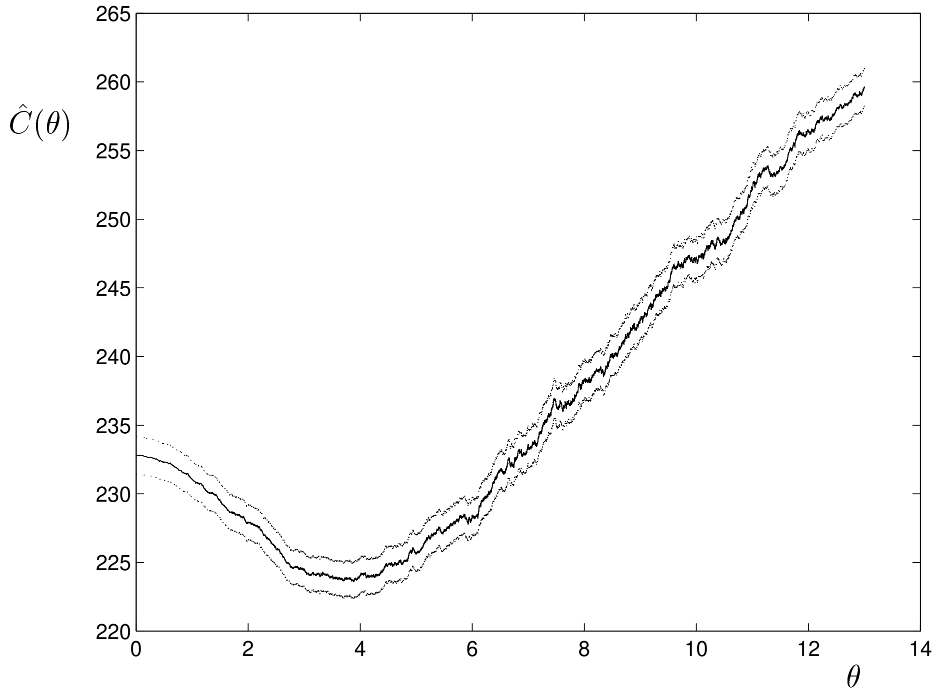

Figure 1: Estimation of the total expected cost $\hat{C}(\theta)$ as a function of θ.

rate of X_n starts to decrease and eventually becomes negative, until $X_n = 0$ when the simulation is over for all the nodes.

We observe from Figure 2 that $E[X_n] \approx Mn/2$, so the growth of the tree appears linear in both n and M. We actually made this experiment with other values of M and the results agreed very well with the empirical formula $E[X_n] \approx Mn/2$. The upper bound of Proposition 1 is quadratic in M, but our empirical results suggest a linear growth rate in M instead, at least for this example. We also observed that the number of leaves in the tree is $\mathcal{O}(MT)$, while the total number of nodes and the total CPU time are both $\mathcal{O}(MT^2)$. Almost all of the CPU time is used to run the simulations themselves; the additional overhead for managing the lists of intervals and computing the functional estimator is practically negligible.

If we take a shorter interval $[a, b)$ instead of the entire interval $[0, \bar{\theta})$, we observe the same properties of the growth rate. However, the number of levels in the tree at the end of the simulation increases with θ; that is, it is larger if $[a, b)$ is close to $\bar{\theta}$ than if it is close to 0. This is because a larger θ means fewer replacements at the interventions, leading ultimately to more interventions. Moreover, the growth rate and the total number of nodes in the tree tend to be smaller (by a constant factor) when $[a, b)$ is close to 0, increase when $[a, b)$ moves to the right (larger θ) up to some point, then decrease again when $[a, b)$ moves further to the right. As a concrete illustration, for our example with $M = 10$ we made 200 replications of the simulation over a horizon $T = 200$. For the intervals $[a, b) = [0, 2)$, $[7, 9)$, and $[11, 13)$, the average total number of nodes in the tree was 1212, 11958, and 10244, respectively. The explanation is that for very small values of θ, the interventions tend to replace many components at a time, so there are usually several components of the same age, and thus less splitting. For an interval $[a, b)$ far to the right (large θ), preventive replacements are rare and the ages become more spread, and more components fail before reaching age a, so the ages tend to fall less frequently into $[a, b)$, which again implies less splitting. The fastest growth rate of the tree lies somewhere in the middle.

4. INFINITE-HORIZON STEADY-STATE AVERAGE CASE

4.1. The Regenerative Model and Stability Issues. The average (expected) cost per unit of time over an infinite horizon is defined as

$$\mathcal{C}(\theta) = \lim_{T\to\infty} \frac{\mathcal{C}(T, \theta)}{T}. \tag{5}$$

The Markov process $\{Z_n, n \geq 0\}$ is regenerative, with regeneration points $\{N_j = N_j(\theta)\}$ defined as $N_0 = 0$ and $N_j = \inf\{n > N_{j-1} : Z_n = (0, \dots, 0)\}$ for $j > 1$. That is, regeneration occurs when all the components are replaced. The process $\{Z_n, n \geq N_j\}$ thus has the same distribution as $\{Z_n, n \geq 0\}$ and is independent of $\{Z_n, n < N_j\}$ conditionally on N_j. As a consequence, if we denote $\Delta_j(\theta) = T_{N_j(\theta)}(\theta) - T_{N_{j-1}(\theta)}(\theta)$ the length of the jth regenerative cycle and $W_j(\theta) = C_{N_j(\theta)}(\theta) - C_{N_{j-1}(\theta)}(\theta)$ the total cost incurred during that cycle,

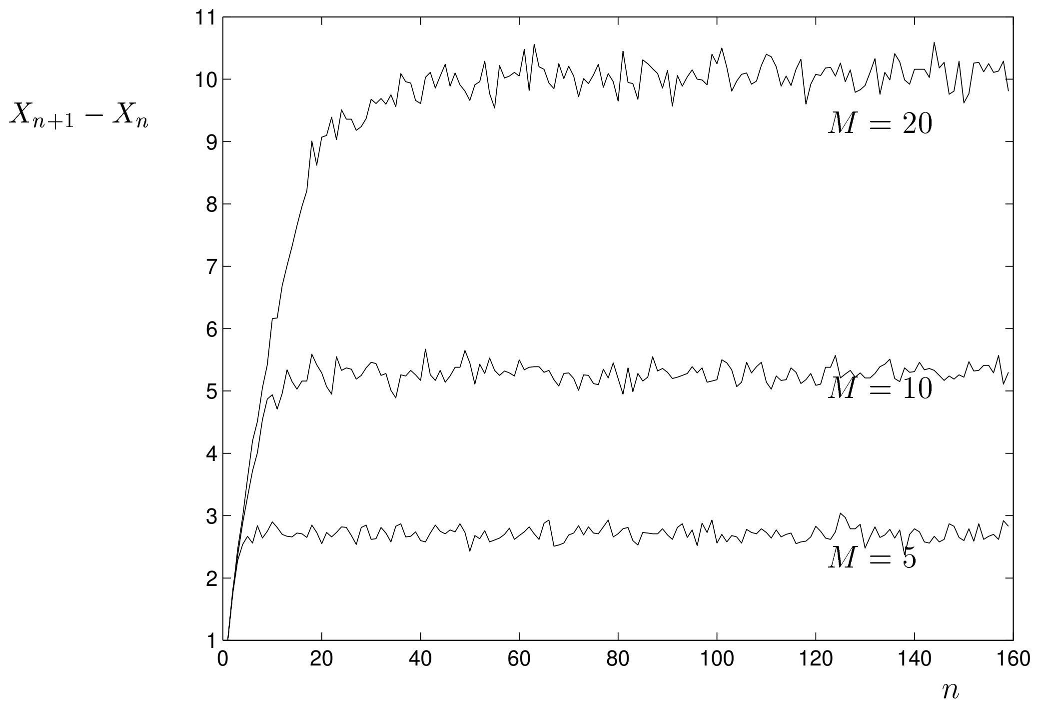

Figure 2: Growth rate of the simulation tree: $X_{n+1} - X_n$ as a function of n.

the vectors $(\Delta_j(\theta), W_j(\theta))$, for $j \geq 1$, are i.i.d. with the same distribution as $(T_{N_1(\theta)}(\theta), C_{N_1(\theta)}(\theta))$.

Under the following (reasonable) lower-bound assumption on the failure rate, we show that $\Delta_j(\theta)$ and $W_j(\theta)$ have bounded moments of all orders. The average cost (5) can then be estimated at any fixed θ by a standard regenerative simulation methodology. We adapt this approach to functional estimation.

Assumption 6. *For any ϵ such that $0 < \epsilon < \bar{\theta}$, there is a constant $\lambda_0 > 0$ such that* $\inf_{\epsilon \leq t \leq \bar{\theta}} \lambda(t) \geq \lambda_0$.

Proposition 7. *Under Assumption 6, for any $\delta > 0$, $N_1(\theta)$ has a bounded moment generating function in some neighborhood of zero, uniformly in $\theta \in (0, \bar{\theta} - \delta]$. Consequently, $N_1(\theta)$, $T_{N_1(\theta)}(\theta)$, and $W_1(\theta)$ have bounded moments of all orders, uniformly in $\theta \in (0, \bar{\theta} - \delta]$.*

Proof. To prove the first part, we will show that $N_1(\theta)$ is stochastically dominated, uniformly in θ, by a geometric random variable independent of θ. The moment generating function of $N_1(\theta)$ is then bounded by that of the geometric random variable, which is finite near zero, and all the moments are thus bounded.

Consider the set of age vectors

$$A_1 = \{(z_1, \ldots, z_M) : z_m \leq \bar{\theta}/2 \text{ for all } m\}.$$

Since the failure rate is bounded, there is a constant $p_1 > 0$ such that if $Z_n \in A_1$, then with probability at least p_1, there will be no failure for the next $\bar{\theta}/2$ units of time, in which case all components will be replaced at the next intervention. If $Z_n \notin A_1$, then for all $\theta \in (0, \bar{\theta} - \delta]$, the age of the oldest component never exceeds $\bar{\theta} - \delta$, so Z_n always belongs to

$$A_2 = \{(z_1, \ldots, z_M) : z_m \leq \bar{\theta} - \delta \text{ for all } m\}.$$

In view of Assumption 6, there is a constant $p_2 > 0$ such that for all $Z_n \in A_2$, with probability at least p_2, all the components whose age exceeds $\bar{\theta}/2 - \delta$ will fail within the next δ units of time and no component will fail more than once during that period. If this happens, then after at most M steps, the system will enter the set A_1. Therefore, from any state in A_2, the system regenerates within the next $M+1$ steps with probability at least $p = p_2 p_1$. The random variable $N_1(\theta)$ is thus stochastically smaller or equal to a geometric random variable with parameter p, for all $\theta \in (0, \bar{\theta} - \delta]$.

To complete the proof, it suffices to observe that the duration and cost of each step are bounded by $\bar{\theta}$ and $\bar{c} = c_1 + c_2 M + c_3$, respectively, so $T_{N_1(\theta)}(\theta) \leq \bar{\theta} N_1(\theta)$ and $W_1(\theta) = C_{N_1(\theta)}(\theta) \leq \bar{c} N_1(\theta)$, and thus the uniform bound on any moment of $T_{N_1(\theta)}(\theta)$ or $W_1(\theta)$ follows from the uniform bound on the corresponding moment of $N_1(\theta)$. □

From this proposition, it follows that the model is geometrically ergodic, uniformly in θ (see Meyn and Tweedie 1993 for more on geometric ergodicity and its consequences). This means in particular that the distribution of Z_n converges to the steady-state distribution, as $n \to \infty$, at a geometric rate. The fact that the first two moments are bounded also implies that

$$\mathcal{C}(\theta) = \frac{E[W_1(\theta)]}{E[\Delta_1(\theta)]} \tag{8}$$

and that the central-limit theorem can be invoked to compute an asymptotically valid confidence interval for $\mathcal{C}(\theta)$.

More specifically, for any $T > 0$, let $R(T, \theta) = \sup\{j : T_{N_j(\theta)}(\theta) \leq T\}$, the number of cycles completed by time T. Define

$$\hat{\mathcal{C}}_+(T, \theta) = \frac{\sum_{j=1}^{R(T,\theta)+1} W_j(\theta)}{\sum_{j=1}^{R(T,\theta)+1} \Delta_j(\theta)}$$

and

$$S^2(T, \theta) = \frac{\sum_{j=1}^{R(T,\theta)+1} (W_j(\theta) - \hat{\mathcal{C}}_+(T, \theta)\Delta_j(\theta))^2}{\left(\sum_{j=1}^{R(T,\theta)+1} \Delta_j(\theta)\right)^2}.$$

For $\theta < \bar{\theta}$, $\hat{\mathcal{C}}_+(T, \theta)$ is a strongly consistent estimator of $\mathcal{C}(\theta)$, and one has

$$\frac{\hat{\mathcal{C}}_+(T, \theta) - \mathcal{C}(\theta)}{S(T, \theta)} \Rightarrow N(0, 1) \tag{9}$$

as $T \to \infty$ ($\Rightarrow N(0, 1)$ means convergence in distribution to the standard normal). This follows from the central-limit theorem for regenerative processes (see, e.g., Crane and Iglehart 1975, Law and Kelton 1991, Meyn and Tweedie 1993, Wolff 1989). Note that the central-limit theorem holds as well if we replace $R(T, \theta) + 1$ by $R(T, \theta)$ or by a fixed constant R (with $R \to \infty$ instead of $T \to \infty$) in (9). Taking $R(T, \theta) + 1$ means that the ongoing cycle at time T must be completed until the next regeneration point. We adopted the latter strategy in our numerical experiments because under this setup, the bias of the regenerative estimator $\hat{\mathcal{C}}_+(T, \theta)$ is $\mathcal{O}(1/T^2)$, as opposed to $\mathcal{O}(1/T)$ if we discard the ongoing cycle and $\mathcal{O}(1/R)$ with a fixed R (Meketon and Heidelberger 1982).

4.2. Functional Estimation: Implementation Issues. Implementing the functional estimation is a bit more complicated in the regenerative setup than for the finite-horizon case, because the number of regenerative cycles and the actual ending time of the simulation vary with θ (recall that we always complete the ongoing regenerative cycle in each branch of the tree). Note that the frequency of regeneration generally *decreases* with θ, because a larger θ means that fewer components are replaced at the interventions. Thus, the extra simulation time (after T) tends to increase with θ. For θ very close to $\bar{\theta}$, regeneration becomes so infrequent that the method is no longer viable.

The growth of the simulation tree can be reduced by *merging* neighboring nodes when their trajectories reach a common state. Reducing the number of nodes in the tree could significantly improve the performance. In our implementation, two (or more) nodes corresponding to adjacent intervals of values of θ are merged into a single node when they both reach a regeneration point at a given step n; i.e., nodes $(n, i-1)$ and (n, i) are merged when $Z_n(i-1) = Z_n(i) = (0, \ldots, 0)$. If the nodes to be merged have a different value on the simulation clock, i.e., $T_n(i-1) \neq T_n(i)$, the clock of the new child node (after the merging) takes the *smallest* value. Because of this, the empirical function obtained with the merging approach will be based on slightly more simulation, at many values of θ, than the non-merging one, and is therefore expected to change its value at more places (i.e., to have more jumps). Note that merging occurs more frequently for small values of θ than for large ones, because the system regenerates more often for small θ. Before merging nodes, the boundaries of the corresponding intervals, and the associated values of $\Delta_j(\theta)$ and $W_j(\theta)$ obtained so far, must be saved to compute the functional estimator at the end. We use *heap* data structures (c.f., Knuth 1973) for that purpose.

We made experiments with the same example as in Section 3, to estimate the steady-state average cost over the interval $[0, \bar{\theta}/2) = [0, 6.5)$. We repeated the entire simulation 100 times, first with the merging technique, then with-

Table 1: Statistics on the performance of the merging method

T	500	1000	2000	4000
With merging				
Total CPU time (sec.)	7.0	15.6	42.5	136.8
Simulation time	5.1	10.6	27.0	77.1
Total no. of nodes	670	1337	2693	5409
Max. tree width	13	15	16	17
No. of jumps	12611	25116	50276	100913
Without merging				
Total CPU time (sec.)	21.3	79.5	305.8	1272.4
Simulation time	20.7	78.1	303.0	1266.4
Total no. of nodes	5884	23041	89739	357960
Max. tree width	90	180	355	714
No. of jumps	10190	20360	40114	80424

out. Table 1 summarizes some of the results regarding the performance of the merging technique, with $T = 500$, 1000, 2000, and 4000. For each of the two methods, we report the total CPU time, the CPU time devoted to the simulation only (the difference is overhead time, including the time to manage the merges, maintain and sort the nodes, compute the function, etc.), the total number of nodes in the tree, the maximum width of the simulation tree at any level, and the total number of jumps in the sample function. All these values are averages over the 100 replications.

Merging brings overhead and additional implementation work, but it clearly wins in terms of performance for the present example. The savings in terms of CPU time increases with T. It is approximately a factor of 10 for $T = 4000$.

The simulation tree remains thin with the merging, and grows wide without. The sample function is piecewise constant, and the number of jumps indicates at how many values of θ this function changes. As expected, this number is approximately proportional to T, i.e., to the total number of simulated regenerative cycles (there is the same expected number of new breakpoints for each regenerative cycle). It is also larger when merging is applied, for the reasons explained earlier.

5. CONCLUSION

We have studied the computational efficiency of a split-and-merge approach to functional estimation for a preventive replacement problem in an M-component system. This is a difficult problem: No analytical solution is available for $M > 2$ and classical numerical methods quickly become powerless as M increases. The split-and-merge approach allows the estimation of functions of a threshold control over an entire interval of interest. It was unclear *a priori* if the size of the simulation tree (and of the computational effort) would increase exponentially fast, or at a more reasonable rate, with the time horizon T. The worst case is exponential, but we proved that the expected growth rate of the tree is bounded linearly in T and quadratically in M. Our formal proof is only for half of the control space, but we also presented experimental evidence that the bound is valid over the entire control space and that it is not tight as a function of the number of components. The growth of the tree seems to be linear in both T and M.

We then showed that the model under study is regenerative, and proved that the length of a regenerative cycle and the cost incurred over a cycle have bounded moments of all orders, uniformly over the control space. Finally, we showed by an example how the performance of the algorithm can be improved by merging adjacent nodes of the simulation tree when they simultaneously reach the regeneration state. We have offered a viable functional estimation approach for this problem. A next step for further investigation could be to study the convergence of an optimization algorithm based on this approach.

ACKNOWLEDGEMENTS

This work has been supported by NSERC-Canada grants # ODGP0110050, # SMF0169893, and # WFA0139015, and by FCAR-Québec grant # 93ER1654. The paper was written while the first author was on sabbatical in the Department of Industrial Engineering at North Carolina State University.

REFERENCES

Cho, D. and M. Parlar. (1991). A survey of maintenance models for multi-component systems. *European Journal of Operations Research*, 51:1–23.

Crane, M. A. and D. L. Iglehart. (1975). Simulating stable stochastic systems III: Regenerative processes and discrete-event simulations. *Operations Research*, 23:33–45.

Gertsbakh, I. B. (1977). *Models of Preventive Maintenance.* Amsterdam, The Netherlands: North Holland.

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

Knuth, D. E. (1973). *The Art of Computer Programming, Vol 3: Sorting and Searching.* Reading, Mass.: Addison-Wesley.

Law, A. M. and W. D. Kelton. (1991). *Simulation Modeling and Analysis.* Second ed. New York: McGraw-Hill.

L'Ecuyer, P. (1983). Processus de décision markoviens à étapes discrètes: Application à des problèmes de remplacement d'équipement. PhD thesis, Dépt. d'IRO, Université de Montréal.

L'Ecuyer, P. (1990). A unified view of the IPA, SF, and LR gradient estimation techniques. *Management Science*, 36(11):1364–1383.

L'Ecuyer, P. (1993). Two approaches for estimating the gradient in a functional form. In *Proceedings of the 1993 Winter Simulation Conference*, ed. G. W. Evans, M. Mollaghasemi, E. C. Russell, and W. E. Biles, 338–346, Piscataway, NJ. IEEE Press.

L'Ecuyer, P. and G. Perron. (1994). On the convergence rates of IPA and FDC derivative estimators. *Operations Research*, 42(4):643–656.

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

Martin, B. (1996). Estimation fonctionnelle dans le cadre de politiques de remplacement pour un système à plusieurs composantes. Master's thesis, Département d'IRO, Université de Montréal.

Meketon, M. S. and P. Heidelberger. (1982). A renewal theoretic approach to bias reduction in regenerative simulations. *Management Science*, 26:173–181.

Meyn, S. P. and R. L. Tweedie. (1993). *Markov Chains and Stochastic Stability*. Springer-Verlag.

Pierskalla, W. P. and A. Voelker. (1976). A survey of maintenance models: The control and surveillance of deteriorating systems. *Naval Research Logistic Quarterly*, 23(3):351–388.

Rubinstein, R. Y. and A. Shapiro. (1993). *Discrete Event Systems: Sensitivity Analysis and Stochastic Optimization by the Score Function Method*. New York: Wiley.

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

Wolff, R. W. (1989). *Stochastic Modeling and the Theory of Queues.* New York: Prentice-Hall.