

A Dynamic Programming Policy Improvement Approach to the Development of Maintenance Policies for 2-Phase Systems With Aging

Andrew J. MacPherson and Kevin D. Glazebrook

Abstract—The deterioration observed in many industrial systems may be modeled in two phases. In the first phase, a period during which the system operates fault free ends with entry into a worn state. In the second phase, the system spends time in the worn state prior to failure. Should the system be found to be in the worn state upon inspection, failure can be pre-empted by preventive maintenance. Transitions into the worn state occur more frequently as the system ages, as does the cost of maintaining the system. The goal of analysis is the design of cost-effective policies for the inspection, maintenance, and renewal of such systems. The paper extends previous work by offering a choice between a (cheap) repair and a (more expensive) renewal of the system, should it be found to be in the worn state upon inspection. The decision-maker may also renew the system at any time without inspection. We propose simple, cost effective heuristic policies, whose design avoids the computational complexities of a full dynamic programming (DP) solution. The closeness to optimality of these policies is investigated, as is their sensitivity to parameter misspecification. A numerical investigation identifies cases for which the inclusion of the repair option is most advantageous.

Index Terms—Cost minimization, heuristic policy, policy improvement, preventive maintenance, reactive maintenance, sensitivity analysis, 2-phase systems.

ACRONYMS

PM	Preventive Maintenance
RM	Reactive Maintenance
DP	dynamic programming
p.d.f	probability density function
NAG	Numerical Algorithms Group

NOTATION

τ	Age of system.
T	Time elapsed, from age 0, until the system enters the worn state.

$T - \tau T > \tau$	Time elapsed, from age τ , until a system enters the worn state.
W	Time spent in the worn state prior to system failure.
$h(\cdot)$	Hazard rate for transitions of the system into the worn state.
$f_T(\cdot)$	Probability density function of T .
$f_{T-\tau}(\cdot)$	Probability density function of $T - \tau T > \tau$.
$f_W(\cdot)$	Probability density function of W .
m	Fixed cost of an inspection of the system.
v	Fixed cost of a repair of the system.
r	Fixed cost of a renewal of the system.
R	Fixed cost of system renewal following failure.
$p(FF \tau, x)$	Probability the system is fully functioning at age $x + \tau$ given the system was fully functioning at age τ .
$p(W \tau, x)$	Probability the system is worn at age $x + \tau$ given the system was fully functioning at age τ .
$p(F \tau, x)$	Probability the system fails before achieving age $x + \tau$ given the system was fully functioning at age τ .
$t(\tau, x)$	S -expected time until whichever is earlier of system failure or the system achieving age $x + \tau$, having started with a fully functioning system aged τ .
$V_h^\pi(\tau)$	S -expected cost incurred over time horizon h by maintenance policy π , having started with a fully functioning system of age τ .
$V_h^\pi(\tau x, a)$	S -expected cost incurred over time horizon h where a system of age τ has its first scheduled PM at age $x + \tau$ at which point action a is taken. Thereafter policy π is operated.
γ^π	Cost rate incurred under maintenance policy π .
$w^\pi(\cdot)$	S -bias function for costs of system operated under maintenance policy π .

Manuscript received January 24, 2009; revised March 17, 2010 and November 02, 2010; accepted November 30, 2010. Date of publication April 25, 2011; date of current version June 02, 2011. This work is supported by the Engineering and Physical Sciences Research Council (EPSRC) through the provision of a doctoral studentship to the first author. Associate Editor: L. Walls.

The authors are with the Department of Management Science, Lancaster University, Lancaster, LA1 4YX, U.K. (e-mail: A.MacPherson1@Lancaster.ac.uk; K.Glazebrook@Lancaster.ac.uk).

Digital Object Identifier 10.1109/TR.2011.2135750

- $(\underline{\tau}, k, N)$ A maintenance policy based on a cycle of interventions at ages given by the N -vector $\underline{\tau}$. At ages τ_i $i = 1, \dots, k - 1$ the system is inspected and repaired if worn. At ages τ_i , $i = k, \dots, N - 1$ the system is inspected and renewed if worn. The system is renewed at age τ_N .
- $C(\underline{\tau}, k, N)$ Total cost incurred in a single cycle of the system under the $(\underline{\tau}, k, N)$ maintenance policy.
- $T(\underline{\tau}, k, N)$ Time incurred in a single cycle of the system under the $(\underline{\tau}, k, N)$ maintenance policy.

I. INTRODUCTION

IN [1], Maillart and Pollock discuss the development of plans for the optimal maintenance of a 2-phase system of a kind to be found in many industrial contexts. In such a system, a transition from a **fully functioning state** in which its operation is fault free into a **worn state** is a precursor to failure. The system will fail following a random sojourn in the worn state unless a maintenance intervention prevents that outcome. Should the system fail, it is replaced by a new one at considerable cost. A preventive maintenance (PM) policy is a schedule of inspections of the system at each of which a determination is made of whether it is in the worn state or not. Should the system be found worn, it is renewed. Both inspections and renewals of the system incur costs, though those are small compared to the cost of system failure. In [1], Maillart and Pollock use the methods of stochastic dynamic programming (DP) to design monitoring schedules to minimize an aggregate rate of costs incurred through system failure, and in the course of inspections and renewals.

Following discussions with Gaver and Jacobs and colleagues in relation to their study of failure patterns of helicopter engines (see [2]–[4]), the model discussed in the paper develops the model of Maillart and Pollock to expand the set of maintenance actions available to the decision maker at inspection. We shall consider (as did Maillart and Pollock) 2-phase systems in which the rate of transition into the worn state increases as the system ages. This implies that the burden of system inspection and repair will also increase with age. At some point it will be cost effective to renew such a system because of its age, even when it is fully functioning. Further, we shall suppose that, should a system be found worn at inspection, two PM actions are available: the system may either be **repaired**, namely restored to the fully functioning state but with no reduction of system age, or it may be **renewed** (with a reduction in age to 0) at somewhat greater cost. The repair option is likely to be favored in the early stages of a system's life, but may not be cost effective later on. To summarize, while for Maillart and Pollock the only decision is when to inspect (and then to automatically renew a system, if worn), we are also required to specify at each inspection whether to repair or renew a worn system, and additionally when we should renew the system without inspection. As we shall see, the introduction of a repair option may lead to substantial cost savings.

While it is possible to apply the methods of stochastic DP to our problem, we shall propose two very simple policies (which we shall call Heuristics 1 and 2) whose underlying rationale is clear, whose cost performance is strong, which are easy to implement, and which are trivial to compute. For example, in Heuristic 2, inspection intervals are always chosen such that the probability of failure in each is the same. To obtain Heuristic 2, we perform a (parametric) optimization over the class of such schedules. Heuristic 1 is similar, but equalizes the probability of transition into the worn state within each inspection interval. In 1,800 problems studied, Heuristic 2 outperforms Heuristic 1 in all but four, and was never more than 0.440% suboptimal. We also develop four further heuristics (Heuristics 3-6) by applying a single DP policy improvement step to either Heuristic 1 or Heuristic 2. These policies are also easy to compute and implement. They can also serve as checks on the quality of Heuristics 1 and 2. If they offer little improvement in performance, then it is reasonable to infer that the simple heuristics are close to optimal. If they offer significant improvement, they may very well be preferred. Of the four additional heuristics, the strongest performing policy, Heuristic 4, was never more than 0.056% suboptimal in the 1,800 problems studied.

To the authors' knowledge, both the model discussed in the paper and the DP-based approaches to the construction of heuristic maintenance policies are new. Other studies which use DP to construct maintenance policies which mandate maintenance decisions based on the outcomes of scheduled inspections include those described in [5]–[8]. The models considered in [9]–[17] also share some of the features of our 2-phase system with aging, but those contributions do not share the current paper's focus on the development of near-optimal maintenance policies. The reader is referred to [18], and the literature cited by Maillart and Pollock in [1] for related discussions of the construction of monitoring policies for 2-phase systems, while in [19], Wang gives a comprehensive survey of maintenance policies for deteriorating systems.

The paper is constructed as follows. Our model is described in Section II, while Section III describes cost rate calculations for policies of simple structure. Heuristics 1 and 2 are introduced in Section IV. A stochastic DP formulation of the problem is introduced in Section V where we discuss the design of Heuristics 3-6. Section VI contains a report of a numerical investigation into the effectiveness of the PM policies proposed.

II. THE MODEL

We consider a system which at time 0 is brand new (has age 0), and is **fully functioning**, which here means that it is operating fault free and as intended. As the system is used, it ages. Its age is here taken to be the time elapsed since it was last renewed. A system which is fully functioning of age τ has a probability $h(\tau)\delta + o(\delta)$ of entering the **worn** state in the next δ time units. The function h is a **hazard rate** for the system. If we use T for the time from 0 until the system enters the worn state (i.e., for the first time), then T has the **probability density function (p.d.f.)**

$$f_T(t) = h(t) \exp \left\{ - \int_0^t h(s) ds \right\}, \quad t \geq 0.$$

The notation $T - \tau | T > \tau$ denotes the time elapsed until a system which is currently fully functioning of age τ enters the worn state. The corresponding p.d.f. is given by

$$f_{T-\tau}(t) = h(t + \tau) \exp \left\{ - \int_{\tau}^{t+\tau} h(s) ds \right\}, \quad t \geq 0. \quad (1)$$

The worn state is a precursor to failure, and is not outwardly visible. In the absence of any maintenance intervention, once the system has entered the worn state, it will spend a random time W there before failure. We use $f_W(w)$ for the p.d.f. of W . System failure is catastrophic, immediately observable, and very expensive. We plan maintenance interventions to minimize the rate at which costs are incurred, **both** from system failure, and from the interventions themselves.

The solution methods we shall develop are quite general, and do not require particular distributional choices for T and W . In our numerical work, reported in Section VI, our choices for $h(\tau)$, and $f_W(w)$ are

$$h(\tau) = \alpha \beta^{-1} \tau^{(\alpha-1)}, \quad \tau \geq 0, \quad (2)$$

and

$$f_W(w) = \mu e^{-\mu w}, \quad w \geq 0, \quad (3)$$

respectively. Note that the choice $\alpha > 1$ in (2) guarantees an increasing hazard, as would typically be appropriate. The choice of an exponential distribution for the duration of the worn phase is consistent with Maillart and Pollock [1]. Natural alternatives are an increasing failure rate (IFR) model or a conditional exponential model in which the rate μ is an increasing function of the age of the system when the transition into the worn state occurs. Please note that, in our numerical studies, we shall suppose that the lengths of sojourns in the worn state are small relative to the time of the first transition into it. See also the comment at the end of Section VI-C.

Possible maintenance interventions (and their consequences) are as follows.

- **Preventive Maintenance (PM):** The system can be **inspected** at any time for a fixed cost m . Each inspection reveals (without error) the condition of the system to be either fully functioning or worn. If the system is found to be worn, it may be **repaired** at a further cost v , returning it to the fully functioning state but **without loss of age**. Alternatively, the system may be **renewed** at a further cost r_1 , returning the system to its **brand new** state with age $\tau = 0$. Should the system be fully functioning at inspection, it is allowed to continue operating. In this event, only the cost of monitoring m is incurred.

Alternatively, the system may be renewed at any time without inspection for a fixed cost r_2 . We shall assume throughout that $r_1 = r_2 = r$, but we emphasize that our methodologies do not require this.

- **Reactive Maintenance (RM):** Should the system fail, it is renewed at fixed cost R .

All maintenance interventions are deemed to take negligible time. The cost parameters typically satisfy $R \gg r > v \geq m$. The goal of analysis is the development of a policy for PM (i.e. a rule for making PM decisions) to minimize the aggregate long

run average cost rate from system failures and maintenance interventions. Such a policy will need to balance the costs of excessive inspections against the increased risk of system failure should inspections be too infrequent. We would expect that an older system would require more frequent inspection. An example policy in [1] has such a structure. Further, there may come a point when a system is too old to be maintained economically, and needs to be renewed without inspection.

In the next section, we describe cost rate calculations for policies of simple structure.

III. COST RATE COMPUTATIONS

This section is a necessary digression which explores how long run average cost rates may be computed for a simple class of PM policies. We shall use the triple $(\underline{\tau}, k, N)$ to denote a policy in the class where $\underline{\tau}$ is an N -vector of system ages with components $\tau_1 < \tau_2 < \dots < \tau_N$. At ages $\tau_1 < \tau_2 < \dots < \tau_{k-1}$, a PM intervention of type 1 is scheduled. Such an intervention involves an inspection of the system together with a repair if the system is found to be in the worn state. PM interventions of type 2 are scheduled at ages $\tau_k < \tau_{k+1} < \dots < \tau_{N-1}$. Such an intervention involves an inspection together with a renewal if the system is found to be worn. The system is renewed without inspection at age τ_N . Of course, all of these scheduled PM interventions may be pre-empted by system failure. Note that we shall use $\tau_0 \equiv 0$ in what follows. Please also note that all of the heuristic policies proposed in the paper are of this form.

Before proceeding to the details of the cost rate calculation, we shall introduce some key quantities which are required. Consider a fully functioning system of age τ under a decision that the next maintenance action is due when it achieves age $x + \tau$ (PM) or failure (RM), whichever comes first. We have the following probabilities.

$$\begin{aligned} p(FF|\tau, x) &\triangleq Pr\{\text{system is fully functioning at } x + \tau \\ &\quad | \text{fully functioning at age } \tau\} \\ &= \int_x^\infty f_{T-\tau}(t) dt; \end{aligned} \quad (4)$$

$$\begin{aligned} p(W|\tau, x) &\triangleq Pr\{\text{system is worn at } x + \tau \\ &\quad | \text{fully functioning at age } \tau\} \\ &= \int_0^x \int_{x-t}^\infty f_{T-\tau}(t) f_W(w) dw dt; \end{aligned} \quad (5)$$

$$\begin{aligned} p(F|\tau, x) &\triangleq Pr\{\text{system fails by } x + \tau \\ &\quad | \text{fully functioning at age } \tau\} \\ &= 1 - p(FF|\tau, x) - p(W|\tau, x). \end{aligned} \quad (6)$$

Further, we write $t(\tau, x)$ for the s -expected time until **either** system failure **or** the system achieves age $x + \tau$, whichever comes first. We have that

$$\begin{aligned} t(\tau, x) &= \int_0^x \int_0^{x-t} (t+w) f_{T-\tau}(t) f_W(w) dt dw \\ &\quad + x \{p(FF|\tau, x) + p(W|\tau, x)\}. \end{aligned} \quad (7)$$

Between two successive visits to the brand new (age 0) state, the system performs a **cycle** under policy $(\underline{\tau}, k, N)$. We write $C(\underline{\tau}, k, N)$ for the total cost incurred during a single cycle, and

$T(\underline{\tau}, k, N)$ for the cycle time. As above, we shall also use FF , W , and F to denote the events that a cycle ends with the system in the fully functioning, worn, and failed states respectively, prior to renewal. The long run average cost rate for the system under policy $(\underline{\tau}, k, N)$, from any initial system state, is given by

$$\gamma^{(\underline{\tau}, k, N)} = \frac{E[C(\underline{\tau}, k, N)]}{E[T(\underline{\tau}, k, N)]}. \quad (8)$$

Below find the calculations needed to compute the cost rate under any such policy for which $1 < k < N$. Other cases are handled similarly. We consider three possibilities in turn.

A. Event F - the System is Renewed Upon Failure

We write

$$E[C(\underline{\tau}, k, N)|F, \tau_{i-1} \leq T(\underline{\tau}, k, N) < \tau_i], \quad 1 \leq i \leq N,$$

for the s -conditional s -expected cycle cost given that the system completes a cycle with failure during the interval $[\tau_{i-1}, \tau_i]$. It is straightforward, utilizing the definitions in (4) to (6), that

$$E[C(\underline{\tau}, k, N)|F, \tau_{i-1} \leq T(\underline{\tau}, k, N) < \tau_i] = \begin{cases} R + (i-1)m + v \sum_{j=1}^{i-1} \frac{p(W|\tau_{j-1}, \tau_j - \tau_{j-1})}{1 - p(F|\tau_{j-1}, \tau_j - \tau_{j-1})}, & 1 \leq i \leq k, \\ R + (i-1)m + v \sum_{j=1}^{k-1} \frac{p(W|\tau_{j-1}, \tau_j - \tau_{j-1})}{1 - p(F|\tau_{j-1}, \tau_j - \tau_{j-1})}, & k < i \leq N, \end{cases} \quad (9)$$

with the corresponding s -conditional s -expected cycle time being given by

$$E[T(\underline{\tau}, k, N)|F, \tau_{i-1} \leq T(\underline{\tau}, k, N) < \tau_i] = \tau_{i-1} + \frac{1}{p(F|\tau_{i-1}, \tau_i - \tau_{i-1})} \times \int_0^{\tau_i - \tau_{i-1}} \int_0^{\tau_i - \tau_{i-1} - t} (t + w) f_{T-\tau}(t) f_W(w) dw dt, \quad 1 \leq i \leq N. \quad (10)$$

Finally, the probabilities of the conditioning events in (9) and (10) are given by (11), shown at the bottom of the page.

B. The System Is Renewed Following Inspection

By definition of the policy $(\underline{\tau}, k, N)$, the system is renewed following inspection when it is found to be in the worn state at one of the ages $\tau_k, \dots, \tau_{N-1}$. We have

$$E[C(\underline{\tau}, k, N)|W, T(\underline{\tau}, k, N) = \tau_i] = r + im + v \sum_{j=1}^{k-1} \frac{p(W|\tau_{j-1}, \tau_j - \tau_{j-1})}{1 - p(F|\tau_{j-1}, \tau_j - \tau_{j-1})}, \quad k \leq i < N, \quad (12)$$

while the probability of the conditioning event in (12) is given by

$$Pr\{W, T(\underline{\tau}, k, N) = \tau_i\} = p(W|\tau_{i-1}, \tau_i - \tau_{i-1}) \prod_{j=1}^{k-1} \{1 - p(F|\tau_{j-1}, \tau_j - \tau_{j-1})\} \times \prod_{j=k}^{i-1} p(FF|\tau_{j-1}, \tau_j - \tau_{j-1}), \quad k \leq i < N. \quad (13)$$

C. The System Is Renewed at Age τ_N

We now consider the possibility that a cycle ends without failure at τ_N . We have

$$E[C(\underline{\tau}, k, N)|FF \cup W, T(\underline{\tau}, k, N) = \tau_N] = r + (N-1)m + v \sum_{j=1}^{k-1} \frac{p(W|\tau_{j-1}, \tau_j - \tau_{j-1})}{1 - p(F|\tau_{j-1}, \tau_j - \tau_{j-1})}, \quad (14)$$

and the probability of the conditioning event in (14) is given by

$$Pr\{FF \cup W, T(\underline{\tau}, k, N) = \tau_N\} = \{1 - p(F|\tau_{N-1}, \tau_N - \tau_{N-1})\} \times \prod_{j=1}^{k-1} \{1 - p(F|\tau_{j-1}, \tau_j - \tau_{j-1})\} \times \prod_{j=k}^{N-1} p(FF|\tau_{j-1}, \tau_j - \tau_{j-1}). \quad (15)$$

$$Pr\{F, \tau_{i-1} \leq T(\underline{\tau}, k, N) < \tau_i\} = \begin{cases} p(F|\tau_{i-1}, \tau_i - \tau_{i-1}) \prod_{j=1}^{i-1} \{1 - p(F|\tau_{j-1}, \tau_j - \tau_{j-1})\}, & 1 \leq i \leq k, \\ p(F|\tau_{i-1}, \tau_i - \tau_{i-1}) \prod_{j=1}^{k-1} \{1 - p(F|\tau_{j-1}, \tau_j - \tau_{j-1})\} \\ \times \prod_{j=k}^{i-1} p(FF|\tau_{j-1}, \tau_j - \tau_{j-1}), & k < i \leq N. \end{cases} \quad (11)$$

All modes of system renewal at the conclusion of a cycle have been considered in Sections III-A–III-C above. The s -unconditional s -expectations $E[C(\underline{\tau}, k, N)]$, and $E[T(\underline{\tau}, k, N)]$ may now be computed from the quantities in (9)–(15) as follows.

$$\begin{aligned}
E[C(\underline{\tau}, k, N)] &= \sum_{i=1}^N E[C(\underline{\tau}, k, N) | F, \tau_{i-1} \leq T(\underline{\tau}, k, N) < \tau_i] \\
&\quad \times Pr\{F, \tau_{i-1} \leq T(\underline{\tau}, k, N) < \tau_i\} \\
&\quad + \sum_{i=k}^{N-1} E[C(\underline{\tau}, k, N) | W, T(\underline{\tau}, k, N) = \tau_i] \\
&\quad \times Pr\{W, T(\underline{\tau}, k, N) = \tau_i\} \\
&\quad + E[C(\underline{\tau}, k, N) | FF \cup W, T(\underline{\tau}, k, N) = \tau_N] \\
&\quad \times Pr\{FF \cup W, T(\underline{\tau}, k, N) = \tau_N\}; \quad (16)
\end{aligned}$$

and

$$\begin{aligned}
E[T(\underline{\tau}, k, N)] &= \sum_{i=1}^N E[T(\underline{\tau}, k, N) | F, \tau_{i-1} \leq T(\underline{\tau}, k, N) < \tau_i] \\
&\quad \times Pr\{F, \tau_{i-1} \leq T(\underline{\tau}, k, N) < \tau_i\} \\
&\quad + \sum_{i=k}^{N-1} \tau_i Pr\{W, T(\underline{\tau}, k, N) = \tau_i\} + \tau_N \\
&\quad \times Pr\{FF \cup W, T(\underline{\tau}, k, N) = \tau_N\}. \quad (17)
\end{aligned}$$

The cost rate $\gamma(\underline{\tau}, k, N)$ may now be inferred from (8), (16), and (17).

IV. TWO SIMPLE HEURISTICS FOR PM

We now introduce our two simple heuristics for PM. Both Heuristics 1 and 2 are of the form $(\underline{\tau}, N, N)$. Hence the system is always repaired if found to be in the worn state, and is only renewed either upon system failure or at age τ_N . The two heuristics differ in how the vector $\underline{\tau}$ is determined for a given N . In the case of Heuristic 1, choose $p \in (0, 1)$, and select $\underline{\tau}$ such that

$$\int_{\tau_{i-1}}^{\tau_i} h(t) dt = -\ln(1-p), \quad 1 \leq i \leq N, \quad (18)$$

where $\tau_0 = 0$. Hence $\underline{\tau}$ is chosen such that the probability of a transition from the fully functioning state into the worn state during each subinterval $[\tau_{i-1}, \tau_i)$ is held fixed at p . For the Weibull hazard in (2), the components in $\underline{\tau}$ are determined by the recursion

$$\tau_i = (\tau_{i-1}^\alpha - \beta \ln(1-p))^{1/\alpha}, \quad 1 \leq i \leq N. \quad (19)$$

Clearly, $\underline{\tau}$ is determined uniquely by N and p , which we write as $\underline{\tau}(N, p)$. For any $N \in \mathbb{Z}^+$, $p \in (0, 1)$, the long run cost rate for policy $\{\underline{\tau}(N, p), N, N\}$ may be obtained by the methods of Section III. Heuristic 1 is the policy $\{\underline{\tau}(N^*, p^*), N^*, N^*\}$ where (N^*, p^*) is a cost rate minimizing choice.

Heuristic 2 differs from Heuristic 1 in how the vector $\underline{\tau}$ is determined for given N , and p . Equation (18) is now replaced by

$$\begin{aligned}
&\int_0^{\tau_i - \tau_{i-1}} \int_0^{\tau_i - \tau_{i-1} - t} f_{T - \tau_{i-1}}(t) f_W(w) dw dt \\
&= p(F | \tau_{i-1}, \tau_i - \tau_{i-1}) \\
&= p, \quad 1 \leq i \leq N. \quad (20)
\end{aligned}$$

Hence $\underline{\tau}$ is chosen such that the probability of system failure during each subinterval $[\tau_{i-1}, \tau_i)$ is held fixed at p . As before, Heuristic 2 is obtained by making a cost rate minimizing choice for the pair (N, p) . In both cases, the cost rate minimization is easily accomplished. The optimization routine e04abf, produced by the Numerical Algorithms Group (NAG), is used to minimize over p for a given N . This is followed by a search over N .

Example: As a running example, we consider a set up with a Weibull model for the rate transitions into the worn state, as in (2); and an exponential model for length of stay in the worn state, as in (3). The parameters for these distributions are $\alpha = 1.671$, $\beta = 100$, and $\mu = 0.630$. The cost parameters are given by $m = 1$, $v = 2$, $r = 23.762$, and $R = 87.593$. This is one of the 1800 randomly generated problems which will be discussed in more detail in Section VI. For this case, we designed both Heuristics 1 and 2. All computations throughout the paper were performed on the High Performance Cluster at Lancaster University. Each execution node comprises 124 Sun Fire X4100 servers, two dual-core 2.4 GHz Opteron CPUs, and a standard 8G of memory. Heuristic 1 has an associated long run average cost rate of 3.448, which is 0.980% above the optimum. The computation of Heuristic 1 and its cost rate for this example took 1.6 seconds. Heuristic 2 has a long run average cost rate of 3.424, which is 0.284% above optimal. The computations associated with Heuristic 2 took 6.3 seconds.

A summary of all the results relating to the running example, including more detailed information on the structure of all the heuristic policies discussed in this section and the next, may be found at the end of Section V.

V. SDP FORMULATION AND FOUR FURTHER HEURISTICS

In this section, we shall describe how to develop four further heuristics (Heuristics 3–6) by the application of a single DP policy improvement step to **either** Heuristic 1 (thus yielding Heuristics 3 and 5) **or** Heuristic 2 (yielding Heuristics 4 and 6). This is in the spirit of Krishnan [20] who developed a strongly performing policy for the dynamic routing of customers to queues by applying a single DP policy improvement step to a simple static proposal. Our motivation in doing this is as follows.

- (a) Heuristics 3–6 are easy and quick to compute.
- (b) In all cases studied to date, Heuristics 3–6 have the form $(\underline{\tau}, k, N)$, namely are in the policy class discussed in Section III. They are thus easy to implement, and their cost rates are easy to compute.
- (c) Heuristics 3–6 provide a ready check on the likely quality of Heuristics 1 and 2. If Heuristics 3 and 5 (respectively

Heuristics 4 and 6) offer little cost improvement over Heuristic 1 (respectively Heuristic 2), then the latter is likely close to optimal, and the preferred choice. If, however, they do offer a substantial cost improvement over the simplest heuristics, then they may be preferred.

A. DP Policy Improvement

In what follows, we shall use π to denote some policy for PM maintenance. We write $V_h^\pi(\tau)$ for the s -expected cost incurred over a (large) time horizon h by operating policy π from an initial state in which the system is fully functioning, and of age τ . We expand this notation to $V_h^\pi(\tau|x, 1)$ to indicate the s -expected cost achieved over horizon h when a system of age τ has its first scheduled PM at age $x + \tau$ where $x < h$. It is inspected and repaired (action 1) should it be found to be in the worn state. This is a PM intervention of type 1. Following the first maintenance action (whether PM or RM), maintenance policy π is operated for the remainder of the horizon. We have that

$$\begin{aligned} V_h^\pi(\tau|x, 1) = & \{m + V_{h-x}^\pi(x + \tau)\}p(FF|\tau, x) \\ & + \{m + v + V_{h-x}^\pi(x + \tau)\}p(W|\tau, x) + Rp(F|\tau, x) \\ & + \int_0^x \int_0^{x-t} V_{h-t-w}^\pi(0) f_{T-\tau}(t) f_W(w) dw dt. \end{aligned} \quad (21)$$

We use $V_h^\pi(\tau|x, 2)$ for the equivalent quantity, but where the first scheduled PM action is inspection followed by a renewal of the system (action 2) if it is found to be in the worn state. This is a PM intervention of type 2. We have that

$$\begin{aligned} V_h^\pi(\tau|x, 2) = & \{m + V_{h-x}^\pi(x + \tau)\}p(FF|\tau, x) \\ & + \{m + r + V_{h-x}^\pi(0)\}p(W|\tau, x) + Rp(F|\tau, x) \\ & + \int_0^x \int_0^{x-t} V_{h-t-w}^\pi(0) f_{T-\tau}(t) f_W(w) dw dt. \end{aligned} \quad (22)$$

We further use $V_h^\pi(\tau|x, 3)$ for the equivalent quantity, but where the first scheduled PM action is renewal without inspection (action 3). We have that

$$\begin{aligned} V_h^\pi(\tau|x, 3) = & \{r + V_{h-x}^\pi(0)\} \{p(FF|\tau, x) + p(W|\tau, x)\} \\ & + Rp(F|\tau, x) + \int_0^x \int_0^{x-t} V_{h-t-w}^\pi(0) f_{T-\tau}(t) f_W(w) dw dt. \end{aligned} \quad (23)$$

Our general approach to heuristic development is through the application of a single policy improvement step to some given maintenance policy π . The core idea here is that, in any system state, we make a cost minimizing choice **given** that all future decisions are made according to π . This idea means that, for a fully functioning system of age τ , we seek a pair (x, a) to achieve the minimum value of $V_h^\pi(\tau|x, a)$ in the limit as $h \rightarrow \infty$. To find this pair, we note that the system **regenerates** (reverts to the brand new state) upon renewal or failure, meaning that $V_h^\pi(\tau)$ has the s -asymptotic form (as $h \rightarrow \infty$)

$$V_h^\pi(\tau) \sim \gamma^\pi h + w^\pi(\tau). \quad (24)$$

In (24), γ^π is the long run average cost rate for the system operating under policy π , and $w^\pi(\tau)$ is the s -bias, defined by

$$w^\pi(\tau) \triangleq \lim_{h \rightarrow \infty} \{V_h^\pi(\tau) - V_h^\pi(0)\}. \quad (25)$$

The bias measures the transient effect of starting the system fully functioning in age τ rather than age 0. Using (24) and (25) within (21)–(23), and minimizing, we infer that the policy obtained by applying a single policy improvement step to π will take the action (x, a) in system state τ to achieve the minimum

$$\begin{aligned} \min_x (m \{p(FF|\tau, x) + p(W|\tau, x)\} + Rp(F|\tau, x) - \gamma^\pi t(\tau, x) \\ + \min [vp(W|\tau, x) + w^\pi(x + \tau) \\ \times \{p(FF|\tau, x) + p(W|\tau, x)\}; \\ rp(W|\tau, x) + w^\pi(x + \tau)p(FF|\tau, x); \\ (r - m) \{p(FF|\tau, x) + p(W|\tau, x)\}]). \end{aligned} \quad (26)$$

Please recall that the quantity $t(\tau, x)$ in (26) is given in (7).

The new policy derived from π is determined by a sequence of pairs $\{(x_n^\pi, a_n^\pi), 1 \leq n \leq N\}$ in which $a_N^\pi = 3$. This sequence is obtained from the following algorithm:

Step 1. Set $\tau = 0$, and obtain (x_1^π, a_1^π) as a minimizer of (26). If $a_1^\pi = 3$, then stop, and set $N = 1$. Otherwise, proceed to step 2.

Step k. For $k \geq 2$, set $\tau = \sum_{n=1}^{k-1} x_n^\pi$, and obtain (x_k^π, a_k^π) as a minimizer of (26). If $a_k^\pi = 3$, then stop, and set $N = k$. Otherwise, proceed to step $k + 1$.

This policy chooses action (x_1^π, a_1^π) at age 0. Should $a_1^\pi = 3$, then the system is renewed at age x_1^π if it has not failed before then. Either way, under such a policy, decisions are only required for brand new systems, and we are done. If $a_1^\pi \in \{1, 2\}$, then further decisions may be called for should the system reach age x_1^π without failing ($a_1^\pi = 1$), or without being in the worn state ($a_1^\pi = 2$). The policy chooses action (x_2^π, a_2^π) at age x_1^π . Should $a_2^\pi = 3$, we are done. We continue in this way until a decision to renew the system without inspection is mandated ($a_N^\pi = 3$). While the above account is correct, it could be argued that a **(full) policy** should prescribe an action for a system of **any age**. We have exploited the fact that **stationary policies** will only be called upon for decisions in relation to systems whose age is in a restricted set (e.g. of the form $\tau = \sum_{n=1}^k x_n^\pi$ in the above account).

B. Heuristics 3 and 4

We shall now use the method described above to develop our remaining heuristic policies (3-6) by applying a single policy improvement step to either Heuristic 1 (and in doing so obtain Heuristics 3 and 5), or Heuristic 2 (to obtain Heuristics 4 and 6). To use the above algorithm based on (26), we require the appropriate long run average cost rates γ^π , and biases $w^\pi(\cdot)$ for this purpose. The cost rates are readily available from Section III, and we shall now describe how appropriate estimates of the biases may be developed to produce Heuristics 3 and 4. We shall describe an alternative, simpler, approach which yields Heuristics 5 and 6 in the next subsection.

We now focus on Heuristic 3, to be developed from Heuristic 1. Recall that (N^*, p^*) is the cost rate minimizing choice in the

design of Heuristic 1. To specify biases for use in (26), and the algorithm following, we shall first develop a **full policy** based on Heuristic 1, namely a choice of PM action for a fully functioning system of **any age**.

Consider a system which is fully functioning, and of age t . We shall propose that **starting from age t** the system will undergo a schedule of PM interventions represented by

$$\{\underline{\tau}(t, N_t^*, p^*), N_t^*, N_t^*\} \longrightarrow \{\underline{\tau}(N^*, p^*), N^*, N^*\}. \quad (27)$$

The expression in (27) indicates that, from age t , the system will be subject to an **initial** schedule of PM interventions determined by an N_t^* -vector $\underline{\tau}(t, N_t^*, p^*)$ whose N_t^* components are determined according to (18), with p^* replacing p on the right hand side, and $\tau_0 = t$. The parameter N_t^* is determined below. The left hand expression in (27) indicates that, in the first phase of maintenance of the system, from initial age t , the system is always repaired if found to be in the worn state at inspection, and is only renewed either upon system failure or at an age equal to the final (N_t^* -th) component in $\underline{\tau}(t, N_t^*, p^*)$. The right hand expression in (27) indicates that, once the system has reverted to the brand new state (at either renewal or failure) for the first time, further PM interventions are scheduled according to Heuristic 1. Before proceeding further, we note that $N_0^* = N^*$, and hence that the policy (27) reduces to Heuristic 1 when $t = 0$.

To take the discussion further, we need the following quantities.

- $E[C\{t, \underline{\tau}(t, N_t^*, p^*), N_t^*, N_t^*\}]$ is the total s -expected cost incurred by a system of initial age t until its first entry into the brand new state. During this period, the system is subject to PM interventions determined by $\{\underline{\tau}(t, N_t^*, p^*), N_t^*, N_t^*\}$ as in (27).
- $E[T\{t, \underline{\tau}(t, N_t^*, p^*), N_t^*, N_t^*\}]$ is the equivalent total s -expected time.

Because of the regenerative nature of the policy in (27) after the first return to the brand new state, the s -expected total cost incurred over a large horizon h is easily shown to be s -asymptotically equivalent to (i.e. within an $o(h)$ quantity of)

$$h\gamma^{\{\underline{\tau}(N^*, p^*), N^*, N^*\}} + E[C\{t, \underline{\tau}(t, N_t^*, p^*), N_t^*, N_t^*\}] - E[T\{t, \underline{\tau}(t, N_t^*, p^*), N_t^*, N_t^*\}] \gamma^{\{\underline{\tau}(N^*, p^*), N^*, N^*\}}. \quad (28)$$

All of the quantities in (28) may be computed readily by the methods of Section III. It is now natural to determine a value of the parameter N_t^* by minimizing the quantity in (28). Call this minimizing value \bar{N}_t^* .

We can now make use of (24)–(26), together with (28), to design Heuristic 3 according to the algorithm following (26). In so doing, our choices of the average cost rate γ^π , and bias $w^\pi(\cdot)$ are given by

$$\gamma^\pi = \gamma^{\{\underline{\tau}(N^*, p^*), N^*, N^*\}}, \quad (29)$$

and (30), shown at the bottom of the page. Because γ^π and $w^\pi(t)$ may be readily computed for any required t , the design of Heuristic 3 is easily accomplished computationally.

The development of Heuristic 4 is as above, excepting only that the starting point is Heuristic 2 rather than Heuristic 1. Hence (20) takes the place of (18) throughout.

Example: We consider again the running example presented at the end of Section IV. For this example, Heuristic 3 has an associated long run average cost rate of 3.415 (0.041% above optimal), and the heuristic and cost rate took 21.6 seconds to compute. Heuristic 4 has a cost rate of 3.414 (0.006% above optimal), and the heuristic and cost rate took 1 minute 27.7 seconds to compute. Given that Heuristic 2 outperforms Heuristic 1 for this example, it is not surprising that Heuristic 4 should outperform Heuristic 3 in cost rate terms.

C. Heuristics 5 and 6

The design of Heuristic 5 from Heuristic 1 is also by means of (26), and the algorithm following, but takes a different approach to the development of bias function $w^\pi(\cdot)$. As usual, we use the notation $\{\underline{\tau}(N^*, p^*), N^*, N^*\}$ for Heuristic 1. Consider the system at some age t where $t = \tau_n(N^*, p^*)$ for some n , $1 \leq n \leq N^* - 1$. Hence t is a system age at which a PM action is scheduled under Heuristic 1. Use $\{\underline{\tau}(\tau_n, N^* - n, p^*), N^* - n, N^* - n\}$ for the **residual** PM schedule under Heuristic 1 from system age $t = \tau_n(N^*, p^*)$ onwards. Adapting (27), we suppose that, from age $\tau_n(N^*, p^*)$, the system will undergo a schedule of PM interventions represented by

$$\{\underline{\tau}(\tau_n, N^* - n, p^*), N^* - n, N^* - n\} \longrightarrow \{\underline{\tau}(N^*, p^*), N^*, N^*\}. \quad (31)$$

Hence, from age $\tau_n(N^*, p^*)$, an initial partial implementation of Heuristic 1 is followed until the system is renewed for the first time. From that point, Heuristic 1 is applied in full. Additionally, we suppose that any fully functioning system of age $t \geq \tau_{N^*}(N^*, p^*)$ is renewed immediately before application of Heuristic 1. We depict this action as

$$\text{Renew system} \longrightarrow \{\underline{\tau}(N^*, p^*), N^*, N^*\}. \quad (32)$$

We note that the s -expected total cost of the system evolving from age $t \geq \tau_{N^*}(N^*, p^*)$ under the PM schedule in (32) over a large time horizon h is easily shown to be s -asymptotically equivalent to

$$h\gamma^{\{\underline{\tau}(N^*, p^*), N^*, N^*\}} + r. \quad (33)$$

The above scheme is specified only from initial system ages appearing in the N^* -vector $\underline{\tau}(N^*, p^*)$, and for $t \geq \tau_{N^*}(N^*, p^*)$. We now derive values of the s -bias function evaluated at these ages. Adapting the notation of the previous sub-sec-

$$w^\pi(t) = \begin{cases} E[C\{t, \underline{\tau}(t, \bar{N}_t^*, p^*), \bar{N}_t^*, \bar{N}_t^*\}] - E[T\{t, \underline{\tau}(t, \bar{N}_t^*, p^*), \bar{N}_t^*, \bar{N}_t^*\}] \gamma^{\{\underline{\tau}(N^*, p^*), N^*, N^*\}}, & t > 0 \\ 0, & t = 0. \end{cases} \quad (30)$$

tion, we write $E[C\{\tau_n, \underline{\tau}(t, N^* - n, p^*), N^* - n, N^* - n\}]$ for the total s -expected cost incurred by a system of initial age $\tau_n(N^*, p^*)$, $1 \leq n \leq N^* - 1$, until its first entry into the brand new state. The system is assumed to be operating under the schedule for PM interventions given in (31). We write $E[T\{\underline{\tau}(t, N^* - n, p^*), N^* - n, N^* - n\}]$ for the equivalent s -expected time. Following the development around (28)–(30), (33) above, we infer the s -bias values. All of the quantities in (34), shown at the bottom of the page, may be computed readily by the methods of Section III. Values of the s -bias $w^\pi(t)$ for t -values not included in (34) are now obtained by interpolation, achieved through fitting a cubic spline to the values $w^\pi\{\tau_n(N^*, p^*)\}$, $0 \leq n \leq N^*$. This fitting was executed using the routines e01baf and e02baf from the NAG Mark 21 Library.

Heuristic 5 is now designed by use of (26), and the algorithm immediately following (26); with s -bias values $w^\pi(\cdot)$ developed as above, and the cost rate γ^π given by

$$\gamma^\pi = \gamma\{\underline{\tau}(N^*, p^*), N^*, N^*\}. \quad (35)$$

The development of Heuristic 6 is as above, excepting only that the starting point is Heuristic 2 rather than Heuristic 1.

Example: For our running example, the long run average cost rate associated Heuristic 5 is 3.419 (0.141% above optimal); and the heuristic and its cost rate took 7.7 seconds to compute. Heuristic 6 has a long run average cost rate of 3.416 (0.059% above optimal); and the heuristic and its cost rate took 52.9 seconds to compute. As before, given that Heuristic 2 outperforms Heuristic 1 for this example, it is unsurprising that Heuristic 6 should outperform Heuristic 5.

Given that the development of Heuristics 3 and 4 involves the minimization of the expression in (28), it is unsurprising that these policies outperform Heuristics 5 and 6, and take longer to compute. An advantage of the approach to the development of Heuristics 5 and 6 based on the cubic spline interpolation following (34) is that it easily lends itself to repeated application for so long as the policies thereby created remain within the class introduced in Section III. While we would expect that such an iterative procedure would lead to successive improvements, we know of no theory which guarantees convergence to the optimum.

D. Optimal PM Policies

In the above discussion of our running example, the quality of the PM heuristics is expressed as a percentage cost rate excess above the optimum. We determine the optimum cost rate by using a procedure based on the value iteration algorithm given in

Section 3.5 of [21]. However, we ease the computational burden by exploiting the fact we have access to strongly performing policies in Heuristics 1 and 2, and approximations to their value functions via (24).

We proceed as follows for any given problem. We first compute Heuristic 2, and determine τ_N , the age at which the system is renewed without inspection under this policy. Note that $\tau_N = \tau_{27} = 25.85$ in the running example above. Our value iteration will then act on successive δ -discretizations of the space of all possible states $[0, (3\tau_N/2)]$ with initial quantum δ chosen to be $0.003\tau_N$. For our initial value setting, we take a suitably scaled version of the bias function in (30). This approach is equivalent to initializing the algorithm with a value function approximation of the form given in (24) where policy π is taken to be Heuristic 2. We obtain the optimum cost rates for successive δ -discretizations (halving δ at each step) until convergence is secured (stop when the difference in cost rate from successive δ -discretizations is below the current cost rate estimate multiplied by 5×10^{-5}). Please note that for each successive discretization the algorithm is initialized in the same manner, taking the estimate of the value function available from its predecessor, and suitably adjusting for the new discretization. Please also note that the high quality of our heuristic policies has meant that, to get meaningful cost rate comparisons with the optimum, we have needed to estimate optimal cost rates with high precision. This requirement in turn has meant long run times for the DP algorithm.

Example: Using the above approach, the optimum cost rate for the running example was found to be 3.414. The computation of the optimal policy and its associated cost rate took 4 hours, 8 minutes, 1.8 seconds.

In Table I, find full details of Heuristics 1-6 as designed for the running example, together with the optimal policy. In the first row, for example, we observe that the minimizing choice of (N, p) for Heuristic 1 is given by $N^* = 26$, $p^* = 0.080$. The corresponding policy in $(\underline{\tau}, k, N^*)$ form makes the choices $k = N^* = 26$, $\tau_1 = 3.56$, and $\tau_{26} = 25.01$. Hence, the first scheduled inspection under Heuristic 1 is at $\tau_1 = 3.56$ while the system is renewed without inspection at $\tau_{26} = 25.01$. In each row of the table, we record the time of the first PM intervention (τ_1), the first PM intervention of type 2 (τ_k), and the occasion of the system's renewal without inspection (τ_{N^*}) for the corresponding heuristic.

The table mentions a new policy, Heuristic 2MP, which will be described in greater detail in Section VI-E. This policy takes its cue from the earlier Maillart and Pollock model, and is of the form $(\underline{\tau}, 1, N)$, namely all inspections will be followed by a

$$w^\pi(t) = \begin{cases} E[C\{\tau_n, \underline{\tau}(t, N^* - n, p^*), N^* - n, N^* - n\}] \\ \quad - E[T\{\underline{\tau}(t, N^* - n, p^*), N^* - n, N^* - n\}] \times \gamma\{\underline{\tau}(N^*, p^*), N^*, N^*\}, & t = \tau_n(N^*, p^*), \\ & 1 \leq n \leq N^* - 1, \\ 0, & t = 0, \\ r, & t \geq \tau_{N^*}(N^*, p^*). \end{cases} \quad (34)$$

TABLE I
SUMMARY OF RESULTS FOR RUNNING EXAMPLE

Heuristic	N^*	k	p^*	τ_1	τ_k	τ_{N^*}	Cost Rate	Suboptimality %	Runtime
1	26	26	0.080	3.56	25.01	25.01	3.448	0.980	1.6 sec
2	27	27	0.019	2.54	25.85	25.85	3.424	0.284	6.3 sec
2MP	35	1	0.022	2.68	2.68	32.96	3.646	6.789	9.0 sec
3	28	19	n/a	2.91	20.17	26.96	3.415	0.041	21.6 sec
4	30	20	n/a	2.77	20.75	28.05	3.414	0.006	1 min 27.7 sec
5	25	19	n/a	2.76	20.09	24.81	3.419	0.141	7.7 sec
6	27	20	n/a	2.74	20.73	25.90	3.416	0.059	52.9 sec
Optimal	30	21	n/a	2.77	21.46	27.86	3.414	0.000	4 hrs 8 min 1.8 sec

renewal of the system, if it is found to be in the worn state. The policy is otherwise structured along the lines of Heuristic 2.

As we shall see in the upcoming section, the cost rate performances of our policies for the running example give a good indication of their relative performance for the large collection of 1,800 problems discussed there.

VI. NUMERICAL STUDY

In an extensive numerical study, the six heuristics proposed in Sections IV and V were developed for 1,800 problem instances generated at random, and described below. In all cases studied, all the heuristic policies were of the form $(\underline{\tau}, k, N)$, and so cost rates were computed for all heuristic policies and for each problem by the methods of Section III. This approach was typically achieved in a few seconds of computing time in each case. The worst case was an instance of Heuristic 4 which took 4 minutes, 9.1 seconds to compute. The cost rates achieved by the heuristics are below compared to the optimum cost rate achievable. The latter was obtained by a DP value iteration algorithm, as described in Section V-D. In contrast to the case with the heuristics, the optimum policy and associated cost rate could on occasion only be obtained to the required accuracy after several hours of computing time. The worst case took just over 24 hours. The running example we have used in Sections IV and V is one of our 1,800 sampled problems, and was chosen because it has features which are fairly typical.

All problems studied were as specified in (2) and (3) above. Hence, transitions of the system into the worn state follow a Weibull hazard with parameters α , and β ; and sojourns in the worn state (prior to failure) are exponentially distributed with mean μ^{-1} . In the numerical study which follows, we are able to assess the quality of performance of our six heuristics when applied to 1,800 randomly generated problems. Of initial interest is the question of whether the quality of heuristic performance varies with values of the key parameters α (shape parameter of the Weibull hazard), and R (the cost of renewing the system after failure).

A. Heuristic Performance as α Varies

In Table II(a)–(c), find summaries of the percentage suboptimality achieved by Heuristics 1–6 for 900 randomly sampled problems, with 300 problems chosen for each of three distinct

ranges of the key parameter α . High values of α indicate that the system ages more rapidly as reflected in an increased rate of transitions into the worn state. In problem set 1, system α are sampled independently from the continuous uniform distribution $U[1.65, 1.75]$. Here the hazard rates in (2) are all increasing, and concave in the system age. Results for this case are given in Table II(a). In problem set 2 (Table II(b)), we have $\alpha \sim U[1.95, 2.05]$, and the resulting Weibull hazard rates are increasing and close to linear. For problem set 3 (Table II(c)), we have $\alpha \sim U[2.25, 2.35]$, and the hazard rates are all increasing and convex. In each problem set, a total of 300 problems were generated, and the average cost rate γ^H for each heuristic H was computed and compared to the optimum cost rate. Hence, for each problem set and each of the six heuristics, 300 percentage suboptimality

$$100(\gamma^H - \gamma^{opt})/\gamma^{opt},$$

were computed and summarized via the order statistics MIN (minimum value), LQ (lower quartile), MED (median), UQ (upper quartile), and MAX (maximum value). These summaries of suboptimality are presented in Table II(a)–(c) below.

In this part of the numerical study, other problem parameters are set/generated as follows.

$$m = 1 \quad (36)$$

$$v = 2 \quad (37)$$

$$r \sim U[20, 30] \quad (38)$$

$$R \sim U[45, 90] \quad (39)$$

$$\beta = 100 \quad (40)$$

$$\left\{ \mu \Gamma \left(1 + \frac{1}{\alpha} \right) \beta^{\frac{1}{\alpha}} \right\}^{-1} | \alpha, \beta \sim U[0.09, 0.16] \quad (41)$$

The sampling in (41) ensures that, once the hazard rate h is determined, the parameter μ is set in such a way that the mean time spent in the worn condition (μ^{-1}), expressed as a percentage of the s -expected time until the system's first entry into the worn state, is uniform between 9% and 16%.

B. Heuristic Performance as R Varies

We now explore the quality of performance of Heuristics 1–6 as we vary R , the cost of system renewal following failure

TABLE II
PERCENTAGE SUBOPTIMALITIES OF SIX HEURISTICS FOR (a) HAZARD h INCREASING AND CONCAVE, (b) HAZARD h INCREASING AND CLOSE TO LINEAR, AND (c) HAZARD h INCREASING AND CONVEX. SEE TEXT FOR DETAILS

	Heuristic					
	1	2	3	4	5	6
MIN	0.574	0.213	0.019	0.002	0.084	0.024
LQ	0.981	0.254	0.060	0.013	0.123	0.058
MED	1.112	0.276	0.081	0.018	0.138	0.067
UQ	1.264	0.298	0.111	0.025	0.160	0.076
MAX	1.686	0.363	0.230	0.049	0.223	0.103

(a)

	Heuristic					
	1	2	3	4	5	6
MIN	0.481	0.239	0.009	0.002	0.051	0.033
LQ	0.826	0.289	0.044	0.007	0.112	0.063
MED	1.009	0.317	0.067	0.012	0.136	0.071
UQ	1.138	0.340	0.098	0.021	0.163	0.078
MAX	1.546	0.385	0.267	0.046	0.245	0.115

(b)

	Heuristic					
	1	2	3	4	5	6
MIN	0.382	0.240	0.004	0.002	0.037	0.042
LQ	0.648	0.315	0.030	0.006	0.084	0.063
MED	0.823	0.340	0.050	0.009	0.111	0.073
UQ	0.983	0.370	0.074	0.018	0.143	0.084
MAX	1.312	0.440	0.186	0.049	0.245	0.117

(c)

(equivalently, the ratio of that cost to the inspection cost m). As in Section VI-A, we study three problem sets containing 900 problems in all. Problem set 4 (reported in Table III(a)) consists of 300 randomly generated problems with $R = 45$. Similarly, problem set 5 (Table III(b)) has 300 randomly generated problems with $R = 67.5$, while problem set 6 (Table III(c)) has $R = 90$. The key shape parameters α for the Weibull hazard are obtained by sampling s -independently from the continuous distribution $U[1.65, 2.35]$. Subsequent to this sampling, the parameters m, v, r, β , and μ are set as in (36)–(38), (40) and (41). The results are presented in the same manner as in Section VI-A.

C. Comments on the Numerical Results

The very simple Heuristics 1 and 2 perform remarkably well with a worst case performance of 2.007% suboptimality (for Heuristic 1) in all 1,800 problems studied. That said, Heuristic 2 comfortably outperforms Heuristic 1 throughout the study, and was never more than 0.440% suboptimal. The notion that constructing PM schedules by equalizing the failure probability in each subinterval is plainly a sound one. Heuristic 1 marginally outperformed Heuristic 2 in just four of the 1,800 problems

TABLE III
PERCENTAGE SUBOPTIMALITIES OF SIX HEURISTICS FOR (a) $R = 45$, (b) $R = 67.5$, AND (c) $R = 90$. SEE TEXT FOR DETAILS

	Heuristic					
	1	2	3	4	5	6
MIN	0.369	0.207	0.007	0.002	0.059	0.025
LQ	0.738	0.257	0.056	0.011	0.111	0.049
MED	0.914	0.275	0.096	0.020	0.136	0.060
UQ	1.165	0.302	0.140	0.027	0.164	0.074
MAX	2.007	0.424	0.323	0.056	0.308	0.131

(a)

	Heuristic					
	1	2	3	4	5	6
MIN	0.467	0.233	0.004	0.003	0.038	0.044
LQ	0.783	0.292	0.040	0.007	0.102	0.064
MED	0.974	0.319	0.064	0.013	0.129	0.073
UQ	1.110	0.349	0.088	0.021	0.153	0.082
MAX	1.561	0.414	0.177	0.040	0.230	0.134

(b)

	Heuristic					
	1	2	3	4	5	6
MIN	0.477	0.237	0.003	0.002	0.047	0.048
LQ	0.837	0.303	0.035	0.006	0.089	0.064
MED	0.971	0.325	0.052	0.009	0.119	0.070
UQ	1.119	0.344	0.072	0.016	0.145	0.079
MAX	1.408	0.401	0.124	0.032	0.209	0.109

(c)

TABLE IV
PERCENTAGE SUBOPTIMALITIES WHEN HEURISTICS ARE DESIGNED FOR PARAMETER ESTIMATES. SEE TEXT FOR DETAILS

	Heuristic	
	1	2
MIN	0.398	0.216
LQ	1.190	0.519
MED	1.836	1.115
UQ	2.922	2.141
MAX	7.417	6.796

studied. All of these exceptional cases were characterized by large α (system ages quickly), and large μ (brief sojourn in the worn condition). Note that, for very large μ , transitions into the worn state lead very quickly to failure, and so Heuristics 1 and 2 become very close to each other. As we would expect from an approach based on DP policy improvement, Heuristics 3 and 5 (respectively Heuristics 4 and 6) improve upon Heuristic 1 (respectively Heuristic 2). Plainly, the policy improvement approach which utilizes directly calculated s -bias values which is described in Section V-B, and which yields Heuristics 3 and 4, is usually more powerful than the simpler approach described in

TABLE V
PERCENTAGE SUBOPTIMALITIES FOR HEURISTIC 2MP. SEE TEXT FOR DETAILS

	Heuristic 2	Overall	α -range			R -value		
			Concave	Near Linear	Convex	45	67.5	90
MIN	0.207	1.473	3.959	2.362	1.473	1.907	1.706	1.579
LQ	0.278	4.034	7.536	4.360	2.805	5.432	4.073	3.311
MED	0.310	5.790	9.834	5.645	3.712	7.713	5.651	4.300
UQ	0.339	8.549	12.140	7.474	4.851	9.898	7.979	6.212
MAX	0.440	23.716	20.343	12.700	8.965	23.716	16.286	13.201

Section V-C, which estimates s -biases by cubic spline interpolation, and which yields Heuristics 5 and 6. The relative performance of the six heuristics is remarkably consistent. Heuristic 4 is the strongest performing policy throughout, with a worst case performance of just 0.056% suboptimality in the 1,800 problems. The performance of the heuristics appeared little affected by varying the key parameters α and R .

Finally, please note that we have already remarked in Section II that our methods are not restricted by any assumption of exponentiality for the distribution of the time spent in the worn state prior to failure. Observe that, in our numerical experiments, these sojourns are assumed to be short in comparison to the time to first failure. When this is the case, we would expect the performances of Heuristics 1 and 2 to remain strong under quite general distributional assumptions for the sojourns concerned. Should sojourns in the worn state be longer, and for example negatively associated with the age of the system, we would expect the performance of Heuristic 2 to remain strong, but would not be surprised if Heuristic 1, whose design takes less account of time spent in the worn state, performed less well.

D. Sensitivity Analysis

In practice, the parameters in our model will be estimated. It is thus of considerable interest to study the performance of Heuristics 1 and 2 when they are designed on the basis of parameter values which are at some distance from the ‘true’ ones. To explore this performance, we revisited our 1,800 randomly generated problems, each with its associated values of α , β , and μ . For each problem, five further problems were generated by sampling new parameter values α' , β' , and μ' from the uniform distributions $U[0.9\alpha, 1.1\alpha]$, $U[0.9\beta, 1.1\beta]$, and $U[0.9\mu, 1.1\mu]$ respectively. The original values of the costs m , v , r , and R were retained. We thus have 9,000 randomly generated problems in 1,800 groups of five, each of which contains problems with a common set of originating system parameters α , β , and μ . In every case, Heuristics 1 and 2 were designed on the basis of the generated values α' , β' , and μ' (to be thought of as parameter estimates); and applied to the system with parameters α , β , and μ (to be thought of as true values). In every case, a percentage cost suboptimality was computed. The 9,000 percentage suboptimalities are summarized by order statistics in Table IV.

The heuristics continue to perform strongly, although unsurprisingly their level of performance is somewhat impaired in comparison with the cases in which all parameter values are assumed to be known. While parametric uncertainty has somewhat

improved the performance of Heuristic 1 **relative to** Heuristic 2, it is clear that the latter continues to be the stronger.

E. The Value of a Repair Option

While it seems appropriate and natural to have a cheap repair option available if possible, we now seek to gain some idea of the cost rate advantage conferred thereby. To achieve this understanding, we proceed as follows. For each of our 1,800 randomly generated problems, we designed a version of Heuristic 2 (which we shall refer to as Heuristic 2MP to signify its relation to the earlier Maillart and Pollock model), constructed as in the discussion around (20), except now each PM intervention before τ_N is assumed to be of type 2. Hence, at these points, the system is renewed if found to be worn. In our earlier terminology, Heuristic 2MP is of the form $(\tau, 1, N)$. As before, an optimization over (N, p) is performed. While we fully acknowledge that Heuristic 2MP will not be optimal for a system disallowing repair, the cost rate comparisons generated should be insightful and adequate for our purposes. For each of our 1,800 problems, the percentage suboptimality of Heuristic 2MP (for the model including repair) was computed, and a summary via order statistics is given in Table V.

At the left hand end of Table V, find a summary of the performance of Heuristic 2 for these problems for ready comparison. In the second column are results for Heuristic 2MP overall (i.e. for all 1,800 problems). In the remaining columns, the results are broken down by α -range (concave, near linear, and convex), and by value of R (45, 67.5, and 90). Here, the breakdowns by α and R are informative. The cost saving achieved by the availability of a repair option are greatest when α and R are small. When α is small, the system deteriorates slowly with age, and a repair option is likely to remain cost effective well into the PM schedule. In contrast, large values of the failure cost R are likely to undermine the cost effectiveness of repair.

VII. CONCLUSION AND FURTHER WORK

We have analyzed a 2-phase system in which transitions into the worn state increase in frequency as the system ages. Should the system be found to be worn upon inspection, then it may either be repaired cheaply, or renewed at a rather greater cost. The system may be renewed without inspection at any point. We have proposed two very simple heuristics (Heuristics 1 and 2) with four more (Heuristics 3-6) derived by methods based on DP policy improvement. All of these policies are easily computed, and close to cost minimizing in all problems studied.

For this model class, we propose the adoption of Heuristic 2 whose inspection intervals are designed to equalize the probability of system failure in each. At each inspection, the system is repaired if found to be worn. A check on the quality of Heuristic 2 is available via the design and evaluation of Heuristic 4. If Heuristic 4 offers little cost rate reduction in comparison to Heuristic 2, then it is reasonable to infer that the latter is close to optimal. If Heuristic 4 does offer significant improvement, it may well be preferred. Heuristic 4 will likely differ from Heuristic 2 in that the inspections which come late in each cycle of the system (i.e. between renewals) will lead to renewal, not repair, of a system found to be worn. If the system ages quickly, and sojourns in the worn state are brief, Heuristic 1 is a serious competitor to Heuristic 2, and may on occasion outperform it in cost terms. Additional insights from our work are (i) that the precision with which model parameters are estimated has significance for the quality of performance of the derived maintenance policies, with some deterioration in performance to be expected when there is parametric uncertainty; and (ii) that inclusion of the repair option in the model yields significant cost rate reductions.

We plan to take the work forward. Lagrangian techniques should allow us to use similar approaches to those of the current paper to tackle problems of maximizing system availability in situations where (i) inspections, repairs, renewals, and failures all lead to system down time; and (ii) there is some fixed budget for maintaining the system. We also plan to consider the deployment of a **sensor** to support system monitoring. Such a sensor **should** give a signal whenever the system enters the worn state, but is typically imperfect, and has associated errors of both kinds (false positives, and false negatives). The question arises of how such a sensor might best be used within an overall policy for maintenance and replacement of the system. How reliable does such a sensor need to be for it to yield significant benefits? Methods based on the approaches described in this paper will assist in answering such questions.

ACKNOWLEDGMENT

The first author would like to thank the Engineering and Physical Sciences Research Council for supporting his doctoral studies with a studentship. Both authors would like to thank Prof. Don Gaver and Prof. Pat Jacobs of the Naval Postgraduate School, and Doctor Chris Kirkbride of Lancaster University (U.K.) for helpful conversations; and two anonymous reviewers whose comments have significantly improved the paper.

REFERENCES

- [1] L. M. Maillart and S. M. Pollock, "Cost-optimal condition-monitoring for predictive maintenance of 2-phase systems," *IEEE Transactions on Reliability*, vol. 51, no. 3, pp. 322–330, September 2002.
- [2] P. A. Jacobs, A nonhomogeneous poisson process model (NHPP) for engine data Dept O.R., Naval Postgraduate School, Monterey, CA, 2006, Working Paper.
- [3] D. P. Gaver and P. A. Jacobs, A nonhomogeneous Poisson process model for CBM Dept. O.R., Naval Postgraduate School, Monterey, CA, 2006, Working Paper.

- [4] S. E. Gauthier, "Decision analysis to support condition-based maintenance plus," Master's thesis, Dept. O.R., Naval Postgraduate School, Monterey, CA, June 2006.
- [5] H. C. Tijms and F. A. van der Duyn Schouten, "A Markov decision algorithm for optimal inspections and revisions in a maintenance system with partial information," *European Journal of Operational Research*, vol. 21, no. 2, pp. 245–253, August 1985.
- [6] D. Chen and K. S. Trivedi, "Optimization for condition-based maintenance with semi-Markov decision process," *Reliability, Engineering and System Safety*, vol. 90, no. 1, pp. 25–29, October 2005.
- [7] L. M. Maillart, "Maintenance policies for systems with condition monitoring and obvious failures," *IIE Transactions*, vol. 38, no. 6, pp. 463–475, June 2006.
- [8] L. M. Maillart and L. Zheltova, "Structured maintenance policies on interior sample paths," *Naval Research Logistics*, vol. 54, no. 6, pp. 645–655, September 2007.
- [9] A. H. Christer, "Prototype modelling of irregular condition monitoring of production plant," *IMA Journal of Mathematics Applied in Business & Industry*, vol. 3, no. 3, pp. 219–232, 1991.
- [10] A. H. Christer, "Developments in delay time analysis for modelling plant maintenance," *Journal of the Operational Research Society*, vol. 50, no. 11, pp. 1120–1137, November 1999.
- [11] R. Baker and W. Wang, "Developing and testing the delay-time model," *Journal of the Operational Research Society*, vol. 44, no. 4, pp. 361–374, April 1993.
- [12] R. Baker and A. H. Christer, "Review of delay-time OR modelling of engineering aspects of maintenance," *European Journal of Operational Research*, vol. 73, no. 3, pp. 407–422, March 1994.
- [13] W. Wang, "A two-stage prognosis model in condition based maintenance," *European Journal of Operational Research*, vol. 182, no. 3, pp. 1177–1187, November 2007.
- [14] T. Aven and I. Castro, "A delay time model with safety constraint," *Reliability, Engineering and System Safety*, vol. 94, no. 2, pp. 261–267, February 2009.
- [15] R. J. Ferreira, A. T. de Almeida, and C. A. Cavalcante, "A multi-criteria decision model to determine inspection intervals of condition monitoring based on delay time analysis," *Reliability, Engineering and System Safety*, vol. 94, no. 5, pp. 905–912, May 2009.
- [16] W. Wang, "An inspection model with two types of inspections and repairs," *Reliability, Engineering and System Safety*, vol. 94, no. 2, pp. 526–533, February 2009.
- [17] W. Wang, "Modeling planned maintenance with non-homogeneous defect arrivals and variable probability of detection," in *Proceedings of the 8th International Conference on Reliability, Maintainability and Safety*, Chengdu, Peoples R. China, 2009, pp. 73–78.
- [18] I. Gertsbakh, *Reliability Theory: With Applications to Preventative Maintenance*. New York: Springer, 2000.
- [19] H. Wang, "A survey of maintenance policies of deteriorating systems," *European Journal of Operational Research*, vol. 139, no. 3, pp. 469–489, June 2002.
- [20] K. Krishnan, "Joining the right queue: A markov decision rule," in *Proceedings of the 28th IEEE Conference on Decision and Control*, Tampa, FL, USA, 1989, pp. 1863–1868.
- [21] H. C. Tijms, *Stochastic Models: An Algorithmic Approach*. Chichester: John Wiley and Sons, 1994.

Andrew J. MacPherson graduated 1st class with honors from the University of Newcastle-Upon-Tyne with a degree in Mathematics and Statistics. He is currently studying for the degree of Doctor of Philosophy (Ph.D.) in Operational Research at Lancaster University (UK) under the supervision of Professor Kevin D. Glazebrook.

Kevin D. Glazebrook is Distinguished Professor of Operational Research (OR) at Lancaster University. He is a former NRC Senior Research Fellow at the Naval Postgraduate School in Monterey, California. He is the author of one book, and over 130 papers in OR, applied probability, and statistics; and has served on the editorial boards of the journals *Operations Research*, *Naval Research Logistics*, *Queueing Systems*, and *Mathematical Methods of Operations Research*. He is currently Director of the LANCOS Initiative - a £13 million project to increase research capacity in OR in the United Kingdom.