Technical Memorandum Number 807

Selective Maintenance Decision-Making over Extended Planning Horizons

by

Lisa Maillart
Richard Cassady
Chase Rainwater
Kellie Schneider

October 2005

Department of Operations
Weatherhead School of Management
Case Western Reserve University
330 Peter B Lewis Building
Cleveland, Ohio 44106

Selective Maintenance Decision-Making over Extended Planning Horizons

Authors

Lisa M. Maillart, Ph.D.
Weatherhead School of Management
Assistant Professor of Operations
Case Western Reserve University

C. Richard Cassady, Ph.D., P.E.
Associate Professor of Industrial Engineering
Department of Industrial Engineering
University of Arkansas

Chase Rainwater
Department of Industrial and Systems Engineering
University of Florida

Kellie Schneider
Department of Industrial Engineering
University of Arkansas

Abstract

"Selective maintenance" models determine the optimal subset of desirable maintenance actions to perform when maintenance resources are constrained. We analyze a corrective selective maintenance model that identifies which components to replace in the finitely long periods of time between missions performed by a series-parallel system. We formulate this multi-mission problem as a stochastic dynamic program and compare the resulting optimal policy to the optimal single-mission policy by executing a large numerical experiment. Our results indicate that these policies rarely differ, and that when they do, the difference in long-run mission reliability is minimal, which suggests that future work should concentrate on extending results for the single-mission problem.

Keywords

selective maintenance, stochastic dynamic programming, myopic policy, infinite horizon

1. Introduction

Most industrial and military organizations depend upon the efficient utilization of repairable systems (vehicles, machinery, computers, etc.) for the successful operation of their organization. Mathematical models of repairable system performance and the design of optimal maintenance policies for these systems have received an extensive amount of attention in the literature [6, 7, 8, 9, 14, 15, 16]. Unfortunately, the vast majority of this research ignores potential limitations on the resources required to perform maintenance actions. For example, the time between flights may limit maintenance on space vehicles; budgetary constraints on replacement parts may limit maintenance on manufacturing equipment; usage patterns may limit opportunities for maintaining computer networks. The common occurrence of resource-constrained maintenance decision-making, and the corresponding shortcoming in the literature, motivates the development of models for selective maintenance: the process of identifying the subset of actions to perform from a set of desirable maintenance actions.

In the original study on selective maintenance, Rice *et al*. [12] define a system of components that must complete a series of missions. Corrective maintenance (repair of failed components) is performed only during finitely-long breaks between missions. Due to the limited maintenance time, it may not be possible to repair all failed components before the next mission. Rice *et al.* develop a nonlinear, discrete, mathematical program which maximizes the reliability of the next mission in which the decision variables are the number of each component type to repair, and the limitation on maintenance time serves as the primary functional constraint. Due to the complexity of the model, the authors recommend total enumeration as a solution approach. However, because total enumeration is ineffective for large scenarios, they also develop a heuristic solution procedure.

Cassady *et al.* [1, 3] extend the work of Rice *et al.* [12] in several ways. First, they analyze more complex systems. Specifically, they consider a system comprised of independent subsystems connected in series, with the components in each subsystem connected in any fashion. Next, they extend Rice *et al.*'s model to the case in which both time and cost are constrained, which leads to the development of three different selective maintenance models: maximizing system reliability subject to both time and cost constraints; minimizing system repair costs subject to a time constraint and a minimum required reliability level; and minimizing total repair time subject to both cost and reliability constraints. Schneider and Cassady [13] also extend the work of Rice *et al*. [12] by considering multiple systems, each of which needs to perform one of a set of missions.

Cassady *et al.* [4] generalize the work of Rice *et al*. [12] in two additional ways. First, they allow Weibull component lifetime distributions. This relaxation permits systems to experience an increasing failure rate (IFR) and requires monitoring of the age of components.

Second, Cassady *et al.* include three maintenance actions: minimal repair of failed components, replacement of failed components, and preventive maintenance. Pohl *et al.* [10] extend the work of Cassady *et al.* [3] by adding component upgrade as a fourth maintenance option.

Lastly, Chen *et al.* [5] extend the work of Rice *et al.* [12] and Cassady *et al.* [1] by considering systems in which each component and the system itself may be in $K + 1$ possible states, 0, 1, …, K. They formulate an optimization model to minimize the total cost of maintenance activities subject to a minimum required system reliability.

These selective maintenance models considered to date treat decision-making relative to a single, future mission. As previously stated, the computational effort required to solve an instance of the single-mission model – a non-separable, nonlinear knapsack problem with integer-valued decision variables – is nontrivial [1]. Extending the existing selective maintenance literature to account for multi-mission planning, in which a system is required to perform a sequence of missions over time, increases the already formidable computational burden. Therefore, we seek to determine whether surmounting this so-called "curse of dimensionality" associated with large dynamic programs [11] is justified by superior long-run expected performance.

Our contribution is three-fold. First, we extend the work of Cassady *et al.* [2] by formulating the multi-mission finite-horizon and infinite-horizon selective maintenance problems as stochastic dynamic programs. Second, we compare the optimal multi-mission selective maintenance policy, with two missions remaining in the planning horizon and with an infinite number of missions remaining in the planning horizon, to the optimal single-mission policy. Specifically, we execute a large number of numerical experiments to assess how often these multi-mission optimal policies differ from the optimal single-period policy. Finally, we compare the long-run performance of the optimal infinite-horizon selective maintenance policy to that of the two-mission optimal policy and the single-mission (myopic) policy. Specifically, we execute a large number of numerical experiments to assess the improvement in long-run mission reliability that results from using the optimal infinite-horizon policy rather than the optimal two-period policy or the myopic policy.

2. System Definition and Performance

Consider a system comprised of m independent subsystems connected in series, and suppose subsystem i contains n_i independent and identical copies of a constant failure rate (CFR) component connected in parallel. Let $\underline{n}$ be a vector denoting the number of components in each subsystem, i.e., $\underline{n} = \left[n_1, n_2, \ldots, n_m\right]$. Each component in the system is labeled (i,j) where i denotes the subsystem number and j denotes the component number (Figure 1). At any point in time, a component is either functioning or failed, which implies that the subsystems and the system also have binary status.

The system must perform a sequence of identical missions with identical, finite-duration breaks between successive missions. Failures only occur during missions, and maintenance is performed only during breaks between missions. Let r_i denote the mission reliability of a component in subsystem i. The reliability of component (i,j) is defined to be the probability that component (i,j) is functioning at the end of the next mission and is given by

$$R_{ij} = \begin{cases} r_i & \text{if component } (i, j) \text{ is functioning at the start of the next mission} \\ 0 & \text{otherwise.} \end{cases}$$

Likewise, the reliability of subsystem *i* is the probability that subsystem *i* is functioning at the end of the next mission. Since each subsystem is a parallel arrangement of its components, subsystem reliability is given by

$$R_i = \coprod_{j=1}^{n_i} R_{ij} = 1 - \prod_{j=1}^{n_i} \left(1 - R_{ij}\right). \tag{1}$$

Let b_i denote the number of functioning components in subsystem *i* at the start of the next mission (i.e., after maintenance), and let $\underline{b} = \left[b_1, b_2, \ldots, b_m\right]$. Since each subsystem contains identical copies of its components, equation (1) simplifies to

$$R_i = 1 - \left(1 - r_i\right)^{b_i}. \tag{2}$$

Finally, the reliability of the system is the probability that all subsystems are functioning at the end of the next mission. Since the system is a series arrangement of its subsystems, system reliability is given by

$$R = \prod_{i=1}^{m} R_i ,$$

or after substituting equation (2),

$$R = \prod_{i=1}^{m} 1 - \left(1 - r_i\right)^{b_i} . \tag{3}$$

Let S_i' denote the number of failed components in subsystem *i* at the end of the next mission, so that S_i' is a random variable given by

$$S_i' = n_i - b_i + Z_i \tag{4}$$

where Z_i is the number of component failures in subsystem *i* during the next mission. Note that Z_i is a binomial random variable comprised of b_i independent and identical Bernoulli trials each having probability of success $1 - r_i$, i.e.,

$$Z_i \sim \text{binomial}(b_i, 1-r_i),$$

which implies that

$$S_i' \in \{n_i - b_i, n_i - b_i + 1, \dots, n_i\}.$$

Let $\underline{S}' = \left[S_1', S_2', \dots, S_m'\right]$, s_i' denote a specific realization of S_i', and $\underline{s}' = \left[s_1', s_2', \dots, s_m'\right]$.

Then,

$$\begin{aligned}
\Pr\left(S_i' = s_i'\right) &= \Pr\left(n_i - b_i + Z_i = s_i'\right) \\
&= \Pr\left(Z_i = s_i' - n_i + b_i\right) \\
&= \binom{b_i}{s_i' - n_i + b_i} (1-r_i)^{s_i' - n_i + b_i} (r_i)^{n_i - s_i'} \\
&\equiv \text{bin}\left(s_i' - n_i + b_i, b_i, 1-r_i\right),
\end{aligned}$$

and

$$\Pr\left(\underline{S}' = \underline{s}'\right) = \prod_{i=1}^{m} \text{bin}\left(s_i' - n_i + b_i, b_i, 1-r_i\right). \tag{5}$$

3. The Multi-Mission Selective Maintenance Model

Consider the point in time at which a mission ends and the system returns to its base of operation for maintenance. Let $\alpha_{i\ell}$ denote the amount of resource ℓ required to repair a component in subsystem *i*, and let β_ℓ denote the amount of resource ℓ available during a single break. Define the collections of these parameters by

$$\underline{\alpha} = \begin{bmatrix} \alpha_{11}, \alpha_{12}, \dots, \alpha_{1\rho} \\ \alpha_{21}, \alpha_{22}, \dots, \alpha_{2\rho} \\ \vdots \\ \alpha_{m1}, \alpha_{m2}, \dots, \alpha_{m\rho} \end{bmatrix},$$

and $\underline{\beta} = \left[\beta_1, \beta_2, \dots, \beta_\rho\right]$, where ρ denotes the number of limited maintenance resources. Next, let $\underline{s} = \left[s_1, s_2, \dots, s_m\right]$ and $\underline{a} = \left[a_1, a_2, \dots, a_m\right]$ where s_i denotes the current number of failed components in subsystem i and a_i denotes the number of components in subsystem i to be repaired prior to the beginning of the next mission. We refer to $\underline{s}$ as the state of the system and $\underline{a}$ as the maintenance action. If $A(\underline{s})$ denotes the set of feasible actions when the system is in state $\underline{s}$, then $\underline{a} \in A(\underline{s})$ if and only if

$$0 \le a_i \le s_i, \text{ integer, } i = 1, 2, \dots, m,$$

and

$$\underline{a}\,\underline{\alpha} \le \underline{\beta} \ .$$

Recall that b_i denotes the number of functioning components in subsystem i at the beginning of the next mission, in which case

$$b_i = n_i - s_i + a_i \,,$$

which implies that

$$\underline{b} = \underline{n} - \underline{s} + \underline{a} \ . \tag{6}$$

Substituting equation (6) into equation (4) yields

$$S_i' = s_i - a_i + Z_i$$

where

$$Z_i \sim \text{binomial}(n_i - s_i + a_i, 1 - r_i)$$

which implies that

$$S_i' \in \left\{s_i - a_i, s_i - a_i + 1, \dots, n_i\right\}$$

and equation (5) can be re-written as

$$\Pr\left(S_i' = s_i'\right) = \Pr\left(s_i - a_i + Z_i = s_i'\right) = \Pr\left(Z_i = s_i' - s_i + a_i\right)$$
$$= \text{bin}\left(s_i' - s_i + a_i,\, n_i - s_i + a_i,\, 1 - r_i\right) = \binom{n_i - s_i + a_i}{s_i' - s_i + a_i}\left(1 - r_i\right)^{s_i' - s_i + a_i}\left(r_i\right)^{n_i - s_i'},$$

and

$$\Pr\left(\underline{S}' = \underline{s}'\right) = \prod_{i=1}^{m} \text{bin}\left(s_i' - s_i + a_i,\, n_i - s_i + a_i,\, 1 - r_i\right) \equiv p\left(\underline{s}' \middle| \underline{s}, \underline{a}\right).$$

The reliability of the system for the next mission is a function of both $\underline{s}$ and $\underline{a}$. Specifically, using equation (5), system reliability for the next mission can be expressed as

$$R\left(\underline{s}, \underline{a}\right) = \prod_{i=1}^{m} 1 - \left(1 - r_i\right)^{n_i - s_i + a_i} .$$

All previous selective maintenance studies seek to maximize this reliability for the next mission only. We, however, consider the more general case in which t future missions remain in the planning horizon. Therefore, we seek to maximize the total expected number of successful missions over the remainder of the planning horizon.

Let $V\left(t, \underline{s}\right)$ denote the maximum expected number of successful missions with t missions remaining in the planning horizon and a system status of $\underline{s}$. When no missions remain, the expected value of the number of successful missions remaining is equal to zero, i.e.,

$$V\left(0, \underline{s}\right) = 0$$

for all states $\underline{s}$. When a single mission remains, the number of successful missions remaining is a Bernoulli random variable with success probability $R\left(\underline{s}, \underline{a}\right)$. Therefore, $V(1,\underline{s})$ equivalent to the maximum value of system reliability for the next mission, i.e.,

$$V\left(1, \underline{s}\right) = \max_{\underline{a} \in A(\underline{s})} \left\{ R\left(\underline{s}, \underline{a}\right) \right\},$$

which is equivalent to the single-mission selective maintenance problem as defined by Cassady *et al.* [2]. In general, when t missions remain in the planning horizon, the maximum expected number of successful missions remaining is given by the stochastic dynamic programming recursion

$$V\left(t,\underline{s}\right)=\max_{\underline{a}\in A(\underline{s})}\left\{R\left(\underline{s},\underline{a}\right)+\sum_{\underline{s}'}V\left(t-1,\underline{s}'\right)p\left(\underline{s}'\mid\underline{s},\underline{a}\right)\right\}. \tag{7}$$

For a given $\underline{s}$ and t, we denote the optimal maintenance action by $\underline{d}^*\left(t,\underline{s}\right)$, i.e.,

$$\underline{d}^*\left(t,\underline{s}\right)=\arg\max_{\underline{a}\in A(\underline{s})}\left\{R\left(\underline{s},\underline{a}\right)+\sum_{\underline{s}'}V\left(t-1,\underline{s}'\right)p\left(\underline{s}'\mid\underline{s},\underline{a}\right)\right\}.$$

The collection of $\underline{d}^*\left(t,\underline{s}\right)$ for all $\underline{s}$ and t is called the optimal maintenance policy.

Consider the case in which the planning horizon, i.e., the number of missions remaining, is infinite. Since $\lim_{t\to\infty}V\left(t,\underline{s}\right)=\infty$, the stochastic dynamic program represented by equation (7) is not tractable. Therefore, we formulate the infinite-mission case using an average reward criterion (long-run average mission reliability) rather than a total reward criterion (total expected number of successful missions). Over an infinite-horizon, doing so results in the same optimal selective maintenance policy as would the total reward criterion formulation since the intermediate reward, i.e., the number of successes in the next mission, is a Bernoulli random variable with probability of success $R\left(\underline{s},\underline{a}\right)$. Furthermore, since the problem clearly meets the unichain condition, we employ the following result [11],

$$\lim_{t\to\infty}V\left(t,\underline{s}\right)=t\gamma+W\left(\underline{s}\right), \tag{8}$$

where γ is the steady-state system reliability under the optimal infinite-horizon selective maintenance policy, and $W(\underline{s})$ is the bias associated with starting in state $\underline{s}$. Taking the limit of both sides of equation (7) and substituting using (8) yields

$$W(\underline{s}) = \max_{\underline{a} \in A(\underline{s})} \left\{ R(\underline{s}, \underline{a}) + \sum_{\underline{s}'} W\left(\underline{s}'\right) p\left(\underline{s}' \mid \underline{s}, \underline{a}\right) \right\} - \gamma .$$

Analogously to the finite-horizon case, we define the optimal infinite-horizon action when the system is in state $\underline{s}$ to be

$$\underline{d}^*(\underline{s}) = \arg\max_{\underline{a} \in A(\underline{s})} \left\{ R(\underline{s}, \underline{a}) + \sum_{\underline{s}'} W\left(\underline{s}'\right) p\left(\underline{s}' \mid \underline{s}, \underline{a}\right) \right\} - \gamma .$$

Likewise, we refer to the collection of $\underline{d}^*(\underline{s})$ for all $\underline{s}$ as the optimal infinite-horizon selective maintenance policy.

4. Numerical Experimentation

4.1 An Insightful Example

As an example, consider the system portrayed in Figure 1, suppose that r_1 = 0.7962, r_2 = 0.8623 and r_3 = 0.9558, and suppose that the between-mission maintenance is constrained by ρ = 3 limited maintenance resources such that

$$\underline{\alpha} = \begin{bmatrix} 1.93 & 2.85 & 2.84 \\ 3.82 & 3.28 & 1.06 \\ 1.52 & 3.47 & 3.76 \end{bmatrix} \tag{30}$$

and $\underline{\beta} = [11.7, 20.1, 20.2]$. Since $\underline{n} = [5, 3, 2]$, there are

$$(5+1)(3+1)(2+1) = 72 \tag{31}$$

$\underline{s}$ vectors.

Define Φ to be the set of all states that require selective maintenance, i.e.,

$$\Phi \equiv \left\{ \underline{s} \;\middle|\; \exists\, \ell \in \{1, 2, \ldots, \rho\} \ni \sum_{i=1}^{m} \alpha_{i\ell} s_i > \beta_\ell \right\},$$

and let R_{max} denote the system reliability if the system is fully functioning after maintenance, i.e.,

$$R_{max} \equiv \prod_{i=1}^{m} 1-\left(1-r_i\right)^{n_i} .$$

If $\underline{s} \notin \Phi$, then the optimal action is to repair all failed components and $R\left(\underline{s}, \underline{d}^*\right) = R_{max}$.

For this example, repairing all failed components in state $\underline{s} = [3, 2, 1]$ would require 14.95 units of resource one. Clearly, since there are only 11.7 units of resource one available, $[3, 2, 1] \in \Phi$. In fact, half of the system states require selective maintenance in the example problem. These 36 states and their corresponding values of $\underline{d}^*\left(1, \underline{s}\right)$ and $V\left(1, \underline{s}\right)$, obtained via total enumeration, are presented in Table 1 (R_{max} = 0.99587 for this system). Table 1 also contains the $\underline{d}^*\left(2, \underline{s}\right)$ values for the 36 states in Φ. We solved the two-mission problems using backward induction (Puterman [11]).

As highlighted in Table 1, only four system states require a different optimal maintenance action when $t = 2$ as opposed to when $t = 1$. For each of these states, the $\underline{d}^*\left(2, \underline{s}\right)$ action replaces all failed components in the third subsystem, whereas the $\underline{d}^*\left(1, \underline{s}\right)$ action does not; the $\underline{d}^*\left(1, \underline{s}\right)$ action favors replacing an additional component in subsystem one instead. Intuition suggests that this difference stems from the fact that in general, subsystem three's components are fewer in number and more resource-intensive to repair than subsystem one's components. That is, with two missions remaining, it is optimal to sacrifice the reliability of the next mission in order to protect the system from entering a state in which both resource-hungry, type-three components are failed, thus requiring at least one type-three repair. With one mission remaining, however, this notion of "protection" is irrelevant.

For the 36 $\underline{s}$ vectors in Φ, when we use policy iteration (Puterman [11]) to solve for the optimal infinite-horizon selective maintenance action, $\underline{d}^*(\underline{s})$, we find that it is the same as $\underline{d}^*(2,\underline{s})$ for all $\underline{s}$. To assess the benefit of implementing $\underline{d}^*(\underline{s})$ as opposed to the more readily-computable myopic policy, $\underline{d}^*(1,\underline{s})$, we compute $\gamma - \gamma\left(\underline{d}^*(1,\underline{s})\right)$ where $\gamma\left(\underline{d}^*(1,\underline{s})\right)$ is the long-run system reliability under policy $\underline{d}^*(1,\underline{s})$. We compute $\gamma\left(\underline{d}^*(1,\underline{s})\right)$ by executing the policy evaluation step of the policy iteration algorithm (Puterman [11]). For this example, $\gamma - \gamma\left(\underline{d}^*(1,\underline{s})\right) = 2.9\times10^{-10}$. These results suggest that the multi-mission approach results in superior performance over long horizons, but on a very small scale. The difference is small because the four states for which the myopic policy is suboptimal correspond to relatively unlikely states.

4.2 Experimental Design

The numerical example in Section 4.1 provides some insight into the relative performance of multi-mission selective maintenance policies in comparison to single-mission policies. To gain a better understanding of this behavior, we executed a large set of numerical experiments for the case in which m = 3. We randomly generated 1000 example systems using the following procedure. We drew the number of components in a given subsystem from a discrete uniform probability distribution over the integer set {2, 3, 4, 5}, i.e., a DU(2,5) distribution. If the number of components in a given subsystem is two, then we generated the reliability of a single component in that subsystem from a uniform probability distribution over the interval (0.9,1), i.e., a U(0.9,1) distribution. If the number of components in a given subsystem is three, then we generated the reliability of a single component in that subsystem from a U(0.85,0.95) distribution. If the number of components in a given subsystem is four, then we generated the reliability of a single component in that subsystem from a U(0.8,0.9) distribution. If the number of components in a given subsystem is five, then we generated the reliability of a single component in that subsystem from a U(0.75,0.85) distribution.

We generated the number of limited maintenance resources from a DU(1,4) distribution, and the amount of a each resource consumed by the repair of a given type of component from a U(1,4) distribution. Let $\beta_{\ell,min}$ denote the amount of resource ℓ required to repair a single component in each subsystem, i.e.,

$$\beta_{\ell,min} = \sum_{j=1}^{m} \alpha_{j\ell} \quad .$$

Let $\beta_{\ell,max}$ denote the amount of resource ℓ required to repair all the components in the system, i.e.,

$$\beta_{\ell,min} = \sum_{j=1}^{m} \alpha_{j\ell} n_j \quad .$$

Then, β_ℓ is given by

$$\beta_\ell = \beta_{\ell,min} + \Delta\left(\beta_{\ell,max} - \beta_{\ell,min}\right)$$

where Δ was drawn from a U(0.25,0.75) distribution.

4.3 Numerical Results

For each experiment, for all $\underline{s}$, we solved for $\underline{d}^*(1,\underline{s})$, the single-mission (myopic) policy, $\underline{d}^*(2,\underline{s})$, the two-mission policy, and $\underline{d}^*(\underline{s})$, the infinite-horizon policy. Using code developed in VB.net and a personal computer with a 1.7 GHz processor and 512 MB of RAM, a total of 2 seconds were required to solve the 1000 single-mission problems, an additional 40 seconds were required to solve the two-mission problems, and an additional 19 minutes were required to solve the infinite-horizon problems. For the 1,000 experiments, $\underline{d}^*(\underline{s}) = \underline{d}^*(1, \underline{s})$ for all $\underline{s}$ for all but 34 experiments. For these 34 instances in which the myopic selective maintenance policy and the infinite-horizon selective maintenance policy were different, we computed the average fraction of non-trivial states (i.e., those in Φ) for which the policies differed and the percent loss associated with implementing the myopic policy over the infinite-horizon,

$$\delta \equiv \frac{\gamma - \gamma\left(\underline{d}^*(1, \underline{s})\right)}{\gamma}$$

(Table 2). When these two policies differ, it is only for 4.3 percent of the non-trivial states on average, and the maximum percent loss in long-run system reliability associated with implementing the myopic selective maintenance policy is just 0.000009435%.

Despite the fact that these results suggest that the myopic selective maintenance policy is "close" to optimal over the infinite-horizon, we investigated the characteristics of these 34 problem instances as compared to the other 966 instances by computing three metrics: the average number of state vectors per problem, the average percentage of state vectors that require selective maintenance (that is, the average fraction of states in Φ) and the average number of resources (Table 2). Based on the data, it appears that the selective maintenance problems with different myopic and infinite-horizon policies tend to be "larger" in all three senses. That is, these problems tend to have more components, more system states that require selective maintenance and more resources. This trend is intuitive in that "larger" problems yield more complicated maintenance policies, which have greater opportunity to disagree over different horizons.

It is also interesting that for 33 of the 34 experiments of interest, the two-mission optimal selective maintenance policy is the same as the infinite-horizon optimal policy. For the remaining problem instance, the two-period optimal policy is the same as the myopic policy, and the three-period policy is the same as the infinite-horizon optimal policy.

5. Summary and Conclusion

In this paper, we formulate the finite-horizon and infinite-horizon multi-mission selective maintenance problem as stochastic dynamic programs. Using numerical experimentation, we compare the optimal infinite-mission policy, the optimal two-mission policy and the optimal single-mission policy. Based on these experiments, we find that these policies rarely differ, and when they do, the difference in long-run performance is minimal. Therefore, because the infinite-horizon problems are more computationally intensive, our results suggest that implementing the myopic policy or the two-period policy is likely to be a wise tradeoff for the type of system considered here. Furthermore, these results suggest that the area of future study with greatest benefit to the practitioner may be the development of improved solution procedures for the single-mission problem.

Relative to the multi-mission scenario, several opportunities for further study are warranted. First, our results are limited (due to computational burdens) to the case in which m = 3. Therefore, future work should focus on determining if our results extend to larger systems. Second, our results are limited to series-parallel systems of CFR components that perform a sequence of identical missions with identical breaks between missions. Therefore, future work should also focus on evaluating the characteristics of multi-mission selective maintenance solutions for systems with IFR components, nonlinear resource requirements, and varying mission profiles.

References

1. Bretthauer KM, Shetty B. The nonlinear knapsack problem – algorithms and applications. European Journal of Operational Research 2002;138(3):459-472.

2. Cassady CR, Murdock WP, Pohl EA. A deterministic selective maintenance model for complex systems. Recent Advances in Reliability and Quality Engineering (H. Pham, Editor). Singapore: World Scientific, 2001:311-325.

3. Cassady CR, Pohl EA, Murdock WP. Selective maintenance modeling for industrial systems. Journal of Quality in Maintenance Engineering 2001;7(2):104-117.

4. Cassady CR, Murdock WP, Pohl EA. Selective maintenance for support equipment involving multiple maintenance actions. European Journal of Operational Research 2001;129(2):252-258.

5. Chen C., Meng M Q-H, Zuo MJ. Selective maintenance optimization for multi-state systems, In: Proceedings of the 1999 IEEE Canadian Conference on Electrical and Computer Engineering. (CCECE – '99). Edmonton, 1999.

6. Cho DI, Parlar M. A survey of maintenance models for multi-unit systems. European Journal of Operational Research 1991;51(1):1-23.

7. Dekker R. Applications of maintenance optimization models: a review and analysis. Reliability Engineering and System Safety 1996;51:229-270.

8. McCall J.J. Maintenance policies for stochastically failing equipment: a survey. Management Science 1965:11: 493-524.

9. Pierskallya WP, Voelker JA. A survey of maintenance models: the control and surveillance of deterioration systems. Naval Research Logistics Quarterly 1976;23:353-388.

10. Pohl EA, Cassady CR, Kwinn M. A selective maintenance model for serial manufacturing systems involving multiple maintenance actions. Proceedings of the 17th International Conference on Production Research. Blacksburg, 2003

11. Puterman, ML. Markov decision processes: discrete stochastic dynamic programming. New York: John Wiley & Sons, 1994.

12. Rice WF, Cassady CR, Nachlas JA . Optimal maintenance plans under limited maintenance time. Industrial Engineering Research '98 Conference Proceedings. Banff, BC, Canada, 1998.

13. Schneider K, Cassady CR. Fleet-level selective maintenance, submitted for review at Quality and Reliability Engineering International 2004.

14. Sherif YS, Smith ML. Optimal maintenance models for systems subject to failure – a review. Naval Research Logistics Quarterly 1981;28:47-74.

15. Valdez-Florez C, Feldman R. Survey of preventive maintenance models for stochastically deteriorating single-unit systems. Naval Research Logistics 1989;36:419-446.

16. Wang H. A survey of maintenance policies of deteriorating systems, European Journal of Operational Research 2002:139(3):469-489.

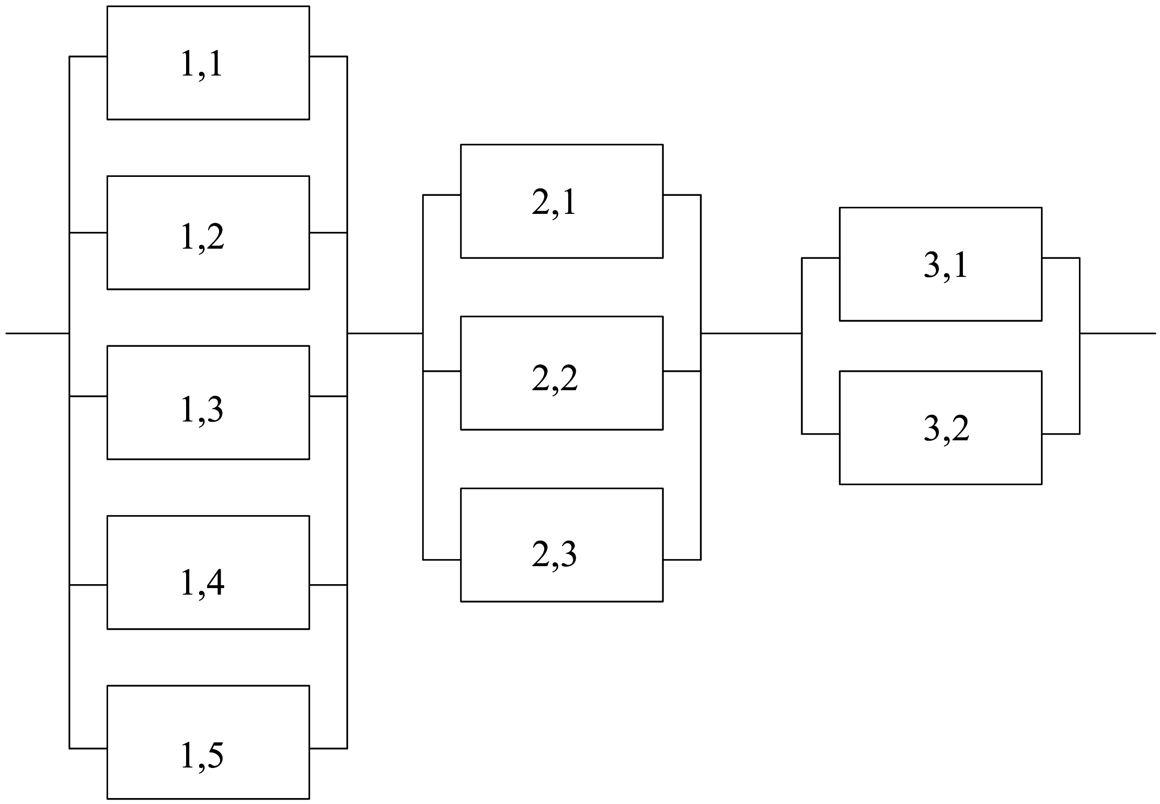

Figure 1. Hypothetical system with $\underline{n}$ = [5, 3, 2].

Table 1. Optimal selective maintenance actions (for t = 1 and t = 2) and the corresponding $V\left(1, \underline{s}\right)$ values for all $\underline{s} \in \Phi$.

$\underline{s}$	$\underline{d}^*(1,\underline{s})$	$V(1,\underline{s})$	$\underline{d}^*(2,\underline{s})$		$\underline{s}$	$\underline{d}^*(1,\underline{s})$	$V(1,\underline{s})$	$\underline{d}^*(2,\underline{s})$
[0, 3, 1]	[0, 2, 1]	0.97955	[0, 2, 1]		[4, 1, 1]	[3, 1, 1]	0.99450	[3, 1, 1]
[0, 3, 2]	[0, 2, 2]	0.97955	[0, 2, 2]		[4, 1, 2]	[2, 1, 2]	0.98779	[2, 1, 2]
[1, 2, 2]	[0, 2, 2]	0.99450	[0, 2, 2]		[4, 2, 0]	[2, 2, 0]	0.98779	[2, 2, 0]
[1, 3, 0]	[0, 3, 0]	0.99450	[0, 3, 0]		[4, 2, 1]	[3, 1, 1]	0.97820	[3, 1, 1]
[1, 3, 1]	[1, 2, 1]	0. 97955	[1, 2, 1]		[4, 2, 2]	[2, 1, 2]	0.97160	[2, 1, 2]
[1, 3, 2]	[0, 2, 2]	0.97820	[0, 2, 2]		[4, 3, 0]	[2, 2, 0]	0.97160	[2, 2, 0]
[2, 2, 1]	[1, 2, 1]	0.99450	[1, 2, 1]		[4, 3, 1]	[2, 2, 0]	0.93947	**[1, 2, 1]**
[2, 2, 2]	[0, 2, 2]	0.98779	[0, 2, 2]		[4, 3, 2]	[1, 2, 1]	0.90813	[1, 2, 1]
[2, 3, 0]	[0, 3, 0]	0.98779	[0, 3, 0]		[5, 0, 2]	[4, 0, 2]	0.99450	[4, 0, 2]
[2, 3, 1]	[1, 2, 1]	0.97820	[1, 2, 1]		[5, 1, 0]	[4, 1, 0]	0.99450	[4, 1, 0]
[2, 3, 2]	[0, 2, 2]	0.97160	[0, 2, 2]		[5, 1, 1]	[3, 1, 1]	0.98779	[3, 1, 1]
[3, 1, 2]	[2, 1, 2]	0.99450	[2, 1, 2]		[5, 1, 2]	[4, 0, 2]	0.97820	[4, 0, 2]
[3, 2, 0]	[2, 2, 0]	0.99450	[2, 2, 0]		[5, 2, 0]	[4, 1, 0]	0.97820	[4, 1, 0]
[3, 2, 1]	[1, 2, 1]	0.98779	[1, 2, 1]		[5, 2, 1]	[3, 1, 1]	0.97160	[3, 1, 1]
[3, 2, 2]	[2, 1, 2]	0.97820	[2, 1, 2]		[5, 2, 2]	[3, 1, 1]	0.93947	**[2, 1, 2]**
[3, 3, 0]	[2, 2, 0]	0.97820	[2, 2, 0]		[5, 3, 0]	[2, 2, 0]	0.93919	[2, 2, 0]
[3, 3, 1]	[1, 2, 1]	0.97160	[1, 2, 1]		[5, 3, 1]	[2, 2, 0]	0.90813	[2, 2, 0]
[3, 3, 2]	[1, 2, 1]	0.93947	**[0, 2, 2]**		[5, 3, 2]	[3, 1, 1]	0.82576	**[2, 1, 2]**

Table 2. Summary of numerical results for $m = 3$

	all 1,000 experiments	966 experiments for which $\underline{d}^*(\underline{s}) = \underline{d}^*(1,\underline{s})\,\forall \underline{s}$	34 experiments for which $\underline{d}^*(\underline{s}) \neq \underline{d}^*(1,\underline{s})\,\forall \underline{s}$
average number of states	90.27	89.51	112.00
average fraction of states in Φ	39.15%	38.54%	56.65%
average number of resources	2.489	2.467	3.118
average fraction of states $\underline{s}$ in Φ for which $\underline{d}^*(\underline{s}) \neq \underline{d}^*(1,\underline{s})\,\forall \underline{s}$	n/a	n/a	4.334%
maximum δ	n/a	0%	0.000009435%