

Research

A Modeling Approach to Maintenance Decisions Using Statistical Quality Control and Optimization

Julie Simmons Ivy^{1,*},[†] and Harriet Black Nembhard²

¹*School of Business Administration, University of Michigan, Ann Arbor, MI 48109, U.S.A.*

²*Harold and Inge Marcus Department of Industrial and Manufacturing Engineering, The Pennsylvania State University, University Park, PA 16802, U.S.A.*

Maintenance concerns impact systems in every industry and effective maintenance policies are important tools. We present a methodology for maintenance decision making for deteriorating systems under conditions of uncertainty that integrates statistical quality control (SQC) and partially observable Markov decision processes (POMDPs). We use simulation to develop realistic maintenance policies for real-world environments. Specifically, we use SQC techniques to sample and represent real-world systems. These techniques help define the observation distributions and structure for a POMDP. We propose a simulation methodology for integrating SQC and POMDPs in order to develop and evaluate optimal maintenance policies as a function of process characteristics, system operating and maintenance costs. A two-state machine replacement problem is used as an example of how the method can be applied. A simulation program developed using Visual Basic for Excel yields results on the optimal probability threshold and on the accuracy of the decisions as a function of the initial belief about the condition of the machine. This work lays a foundation for future research that will help bring maintenance decision models into practice. Copyright © 2005 John Wiley & Sons, Ltd.

KEY WORDS: maintenance; decision making; partially observable Markov decision processes (POMDPs); statistical quality control (SQC); simulation

1. INTRODUCTION

Maintenance is a hidden—and often ignored—expense for all businesses, impacting every industry from manufacturing to service to health care. We define maintenance as any repair or renewal (replacement) action performed on a system. Since systems used in the production of goods and delivery of services constitute the vast portion of capital in most industries, maintenance of such systems is crucial. If such systems deteriorate or fail, the effects can be widespread. Indeed, system deterioration is often reflected in higher production costs and lower product quality¹.

*Correspondence to: Julie Simmons Ivy, School of Business Administration, University of Michigan, Ann Arbor, MI 48109, U.S.A.

[†]E-mail: jsimmons@umich.edu

To address the issue of system maintenance, models and mechanisms that can accurately describe the dynamic nature of quality costs as well as the interaction between such costs, resulting activities, and system improvement, are needed. An accurate model can help decision makers plan their quality activities and estimate the expected costs and benefits of such activities².

Our research focuses on developing methods for determining and evaluating maintenance policies for deteriorating systems under conditions of limited and costly information and with noisy data observations. We propose a model of cost-effective maintenance policies that minimizes the total expected cost of maintaining a machine, process, or service level by integrating statistical quality control (SQC) and partially observable Markov decision processes (POMDPs) through simulation in order to develop realistic maintenance policies for real-world environments. The ultimate goal of this research is to lay the groundwork for bringing the theory of maintenance decision making into practice on an array of important problems in the industrial and medical arenas.

A key element in the maintenance decision is that the action depends on the *state* of the system. The system state is the condition of the system, e.g. in the simplest case whether the system is *good* or *bad*. The dynamic behavior of the system is given by the *state transition*. The state transition is a change in the system condition, e.g. going from good to bad as a result of wear. Most state transitions also entail some uncertainty. For example, the wear may have different rates that may be unknown to the engineer. In situations where the transition depends only on the current state information, a Markov decision process (MDP) is a natural model of the system. A MDP is an optimization model for discrete-stage, stochastic sequential decision making. (Refer to Chen and Feldman³, Chen and Trivedi⁴, and Hontelez *et al.*⁵.) Irvani and Duenyas⁶ use a MDP to emphasize the necessity of addressing maintenance and production control decisions in an integrated manner.

A POMDP is a generalization of a MDP for gathering *incomplete* information regarding the system state, i.e. the decision maker does not know the current state of the system. For example, the state of a machine may include the condition of various internal and unobservable components of the machine. In addition, sensors used to measure the state may give noise-corrupted readings. Finally, exact state observations may be costly to obtain.

In our POMDP modeling, we concentrate on the case where maintenance is imperfect. Imperfect maintenance refers to deterministic maintenance actions whose effect *depends* on the current state of the system. Therefore, since states are probabilistically known, their affect on the state of the system is also probabilistic (i.e. the maintenance performed may not be the exact maintenance needed). Therefore, making a maintenance decision is a management decision.

All management decisions are really management with partial information, and one always faces the issue of whether to use the current information or delay action and learn more (or learn more at cost). Despite numerous studies that define the structure for POMDPs, very little of this research can be applied to a real industrial setting. Indeed, Sinuany-Stern *et al.*⁷ state that practical applications of solution schemes for real-life POMDPs are limited. Most of the existing structural results require many simplifying assumptions, e.g. constant state transition rates. Because of the complexity, stochastic relations, and other uncertainties, not all real-world problems can be represented adequately using a POMDP. Attempts to use POMDPs to optimize such systems require so many simplifying assumptions that the solutions are likely to be inferior or inadequate for implementation.

One such assumption is that the POMDP assumes that if the state is known, then the outcome of the sample is known. This assumption, however, raises many issues when POMDPs are used in practice. For example, suppose a system can be in a good or bad state. A typical POMDP approach is to define one Bernoulli distribution for products produced during the good state and another Bernoulli distribution for products during the bad state. The actual sample from either Bernoulli distribution is then determined by that distribution to be good or bad. However, a supposedly good sample may be out of control with respect to a statistical monitoring chart (likewise, a bad sample may be in control). Furthermore, the monitoring chart to which it is compared has a probability of false positive (the type I error of deeming that a sample is in control when it is not) and false negative (the type II error of deeming that a sample is out of control when it is not). And finally, the sample has some noise associated with its measurement or collection. All of these points impact the decision determined by the POMDP. In our methodology, we integrate an SQC model for the samples that addresses these issues.

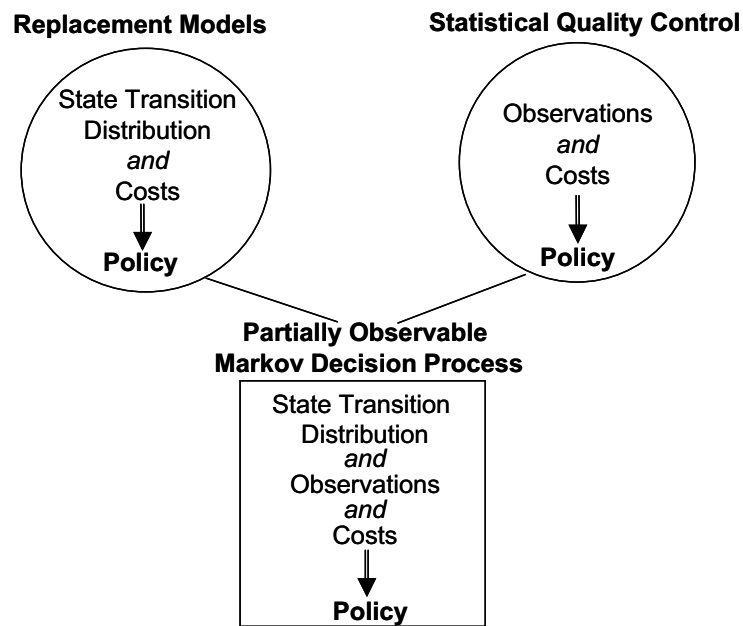


Figure 1. Maintenance literature

Simulation is a technique that will allow us to imitate the operation of a real-world system as it evolves over time. We use simulation to generate the observations according to a SQC model that is representative of a real system (i.e. with noise and monitoring effects) and to evaluate the maintenance decisions made according to the POMDP model. Specifically, we are interested in developing and evaluating optimal maintenance plans as a function of process characteristics, system operating and maintenance costs.

The remainder of the paper is organized as follows. Section 2 discusses the relationship between SQC and POMDPs in the maintenance literature. Section 3 discusses our methodology for integrating SQC and POMDPs. Section 4 presents a case example where the methodology is applied. We conclude with an overview of the future research opportunities in this area in Section 5.

2. THE RELATIONSHIP BETWEEN SQC AND THE POMDP

As shown in Figure 1, the maintenance literature can be divided into three major categories: those for replacement (reliability) models, those for SQC models, and those for POMDP models. We focus on the latter two—the SQC and POMDP models—in our approach to making maintenance decisions.

Replacement models use state transition distributions—with the assumption that ‘time to failure’ or a related measurement gives complete information about the condition of the system—and costs to develop an optimal maintenance policy. The maintenance policy is a set of rules, or a procedure, prescribing action based on observed information. (For example, refer to Wang and Pham¹, and Cremona and Lukic⁸.)

By contrast, much of the research using SQC in maintenance decision making relies on the assumption that the observations give essentially complete information about the condition of the system and the costs to develop the optimal maintenance policy. For example, it may be assumed that the sample or observation comes from a probability distribution that determines the state or that the state completely determines the distribution. There is usually no attempt to make any statistical inference about an observation by comparing it to a statistical monitoring chart. So while an extreme observation from the distribution may be flagged as out of control,

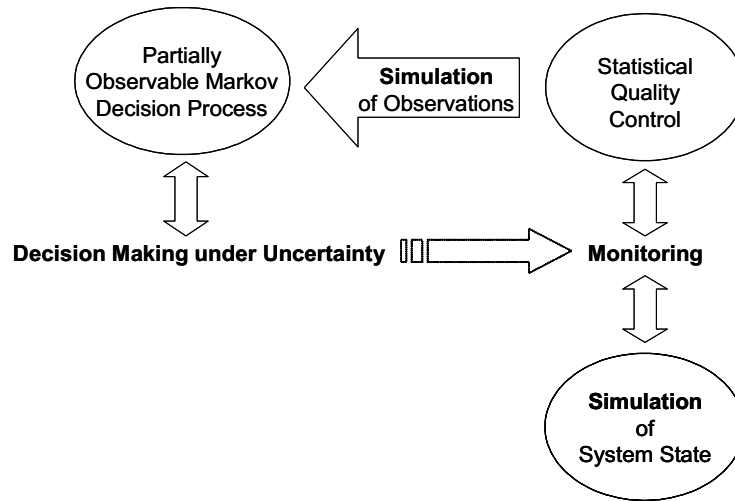


Figure 2. Integration of POMDPs and SQC for decision making under uncertainty using simulation

this approach still maintains that the process is in the in-control state. Furthermore, the fact that any observation will be corrupted with noise is often overlooked. (For example, see Margalio⁹ and von Collani¹⁰.)

POMDPs combine the state transition distribution, process observations, and costs, to determine the optimal maintenance policy. McCall¹¹, Pierskalla and Voeller¹², Monahan¹³, Valdez-Flores and Feldman¹⁴ and White III and White¹⁵ present overviews of theoretical and computational results, applications, and several generalizations of POMDPs and the standard MDP problem formulation. Most stochastic maintenance models make at least one of two assumptions: (1) perfect information is available about the system state through monitoring or inspection (refer to Luss¹⁶ and Rosenfield¹⁷), or (2) maintenance on a system with imperfect information will reveal the true state of the system (e.g. the system *renews* following a maintenance action, refer to Hopp and Wu¹⁸, White III^{19–22}, and Sinuany-Stern *et al.*⁷).

We go beyond these earlier models to focus on the case where maintenance is imperfect. Imperfect maintenance models using POMDPs for simple systems with only two states—new and failed—have been studied extensively. Within this body of literature, Shiryayev²³ proved the Fixed Probability Threshold Rule for the two-state model with geometric failure times, i.e. that a ‘probability threshold rule’ minimizes average cost per unit time. This fixed probability rule states that: *if π , the posterior probability of the system operating in the ‘good’ condition, is less than or equal to some threshold value, then a checking action should be taken.* For similar proofs, see Taylor²⁴ and Teneketzis and Varaiya²⁵. There are many structural results in the literature that suggest, given certain assumptions about the cost structure, that the policy that minimizes the total expected cost has a convenient structure which simplifies the computational complexity of a POMDP model. (See Lovejoy²⁶, Simmons²⁷, and Ivy and Pollock²⁸.)

Because both SQC and POMDPs use the observations of a process, they naturally lend themselves to a modeling integration as we discuss in the next section.

3. METHODOLOGY

As discussed in the introduction, our goal is to present a model of cost-effective maintenance policies that minimizes the total expected cost of maintaining a machine, process, or service level. Our methodology is shown in Figure 2. It integrates POMDPs with SQC through simulation to enable complex decision making under conditions of uncertainty.

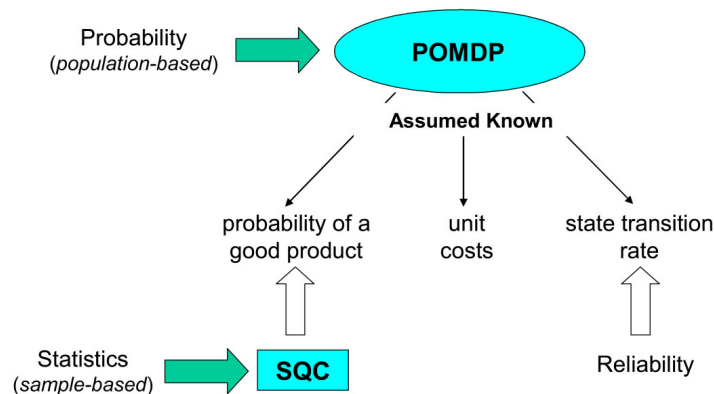


Figure 3. Graphical representation of interrelationship between POMDPs and SQC

We use POMDPs to model deteriorating processes under conditions of uncertainty and to develop the structure of optimal maintenance policies. The POMDP structure allows us to develop a method for selecting maintenance and monitoring actions that minimize the total expected system operation cost over a finite number of observations. This cost includes the cost per unit time for operating the system in each state as well as the cost to change the system state or gather additional system information. We use information provided by the signals machines or processes generate (such as sound, light, pressure, temperature, complaints, etc.), knowledge of the state transition distributions and maintenance and operating costs, to characterize the general structure of the maintenance policy that minimizes total expected cost.

SQC is used to relax the assumption that the distributions of the observations for each system state are known (as done in previous analyses). In applied conditions for studying deteriorating processes under uncertainty, a great deal is often known about the ‘good’ state of a system, so the distribution of observations may be characterized under these conditions. However, little is known about the state of a deteriorating or failed system, making the distributions of observations for these states unknown. We apply real-time monitoring to address this issue.

Simulation is used in four ways. First, it allows us to represent the system state. Second, it allows us to incorporate noise in the observation mechanism and monitoring effects. Third, it allows us to perform sensitivity and robustness analyses for the state transition rates and cost parameters. Fourth, it allows us to evaluate the accuracy of the POMDP maintenance policy.

Although POMDP models have previously been used in maintenance decision making, they require that assumptions about the parameter values, state transition rates, state dependent observation distributions and costs, are all assumed known (as shown in Figure 3). In reality these parameter values are rarely known and must be estimated empirically. Furthermore, since POMDP models are rooted in a probabilistic approach, they deal with the population. In contrast, SQC models are rooted in a statistics approach and deal with samples or observations. Recognizing these core issues will allow us to better extend POMDPs to a real industrial environment. Our methodology allows us to use a SQC model that is representative of a real system to estimate the state dependent observation distribution parameters. We consider an application of the method in the next section.

4. AN APPLICATION OF THE SQC–POMDP METHOD TO THE TWO-STATE MACHINE REPLACEMENT PROBLEM

4.1. Model formulation

Consider a system that can be in either a *good* (0) or *bad* (1) state, so that $S_t \in \{0, 1\}$. The system starts in state 0 and enters state 1 after operating time I , where I is a discrete random variable with a probability mass function

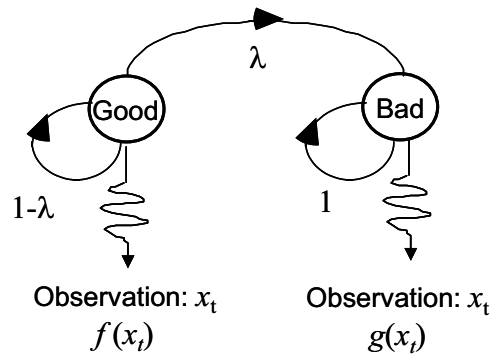


Figure 4. State transition diagram for the two-state model with hazard rate λ and state conditional observation distributions, $f(x_t)$ and $g(x_t)$

$m_I(i) = \lambda(1 - \lambda)^{i-1}$, $i = 1, 2, \dots$. Once the system enters state 1, it remains there until exogenous action is taken. The state transition diagram is shown in Figure 4.

The *state occupancy probability* is the probability that the system is operating in a particular state given the information provided by monitoring. We define

$$\pi_t \equiv \text{Probability}\{S_t = 0\}$$

and

$$(1 - \pi_t) \equiv \text{Probability}\{S_t = 1\}$$

(For notational convenience, state occupancy probabilities written without the subscript t refer to the current state probabilities.)

The SQC plan yields independent observations, X_t , at time t , $t = 1, 2, \dots$, with probability density function (PDF) $f(X_t)$ if $S_t = 0$, or $g(X_t)$ if $S_t = 1$. This implies that the distribution of X_t depends on the system state. We define the vector of observations through time t as $\underline{X}_t \equiv [X_1, X_2, \dots, X_t]$ with the realization $\underline{x}_t \equiv [x_1, x_2, \dots, x_t]$. For simplicity, we assume the state conditional observation distributions for the system if it is in the good state or bad state follow Bernoulli distributions with parameters p_1 and p_2 , respectively. The Bernoulli distribution can reasonably well represent a process in which a single observation is taken at each decision epoch. If the system is in the good state, the probability that the observation is defective is p_1 and if the system is in the bad state the probability that the observation is defective is p_2 .

In our model, decisions are made at discrete points of time referred to as decision epochs. Let T denote the set of decision epochs, $T \equiv \{1, 2, \dots, M\}$, where M is the *horizon* of the problem. We assume that decisions are made at the beginning of each decision epoch. At any decision time t there are two possible actions: the 'do nothing' action, a_0 , allows the system to continue operating; the 'replace' action, a_1 , stops the system from operating and the system is replaced and re-starts in state 0. The goal is to detect when system goes 'bad' or fails, and to do so as soon as possible. Usually, an economic or other advantage is gained when a policy raises an alarm that detects the change soon after it happens. However, it is also desirable to avoid costly false alarms that may occur while the system is still 'good'. The decision tree shown in Figure 5 maps the progression of the decision maker's actions and belief states for the two-state maintenance decision. Initially, the decision maker has a belief, π , about the condition of the machine. If the machine is replaced, the machine is considered 'new' and $\pi = 1$. If the machine is not replaced, it will continue to deteriorate. Following the observation, x , of a sample from the machine, the decision maker's belief state is updated by applying Bayes' rule to $\pi'(x)$.

We associate a fixed cost, d_1 , with performing action, a_1 , and a fixed cost, $d_0 < d_1$ with action, a_0 . Without loss of generality, we can assume $d_0 = 0$. The cost, d_1 , serves as a penalty for false alarms. The cost per unit time for operating the system in state 1 is c_1 , and the cost per unit time for operating in state 0 is c_0 , where $c_0 < c_1$. Without loss of generality, we can set $c_0 = 0$. The cost, c_1 , is the penalty per unit time for failure to

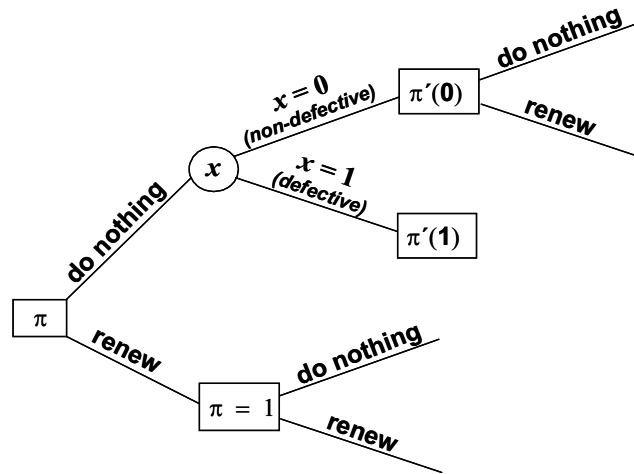


Figure 5. Decision tree for maintenance decisions for two-state model, where π is the probability system is good and $\pi'(x)$ is the updated probability that the system is good given observation x

detect. We define

$$\pi'_{a_l}(x) \equiv \Pr(S_{t+1} = 0 \mid \text{perform action } a_l \text{ at time } t \text{ and observe } X = x \text{ at time } t + 1), \quad l = 0, 1$$

and

$$C_k^*(\pi) \equiv \text{minimum total expected cost with } k \text{ observations available, and probability } \pi \text{ that the system is in state 0, the good state}$$

When $t < M$ and action $a_l \in (a_0, a_1)$ is taken, define $L(a_l|\pi)$, $l = 0, 1$, to be the immediate cost accrued at time t . $L(a_0|\pi) = c_1(1 - \pi)$ and $L(a_1|\pi) = d_1$.

The minimum total expected cost with k observations remaining is

$$C_k^*(\pi) = \min_{l \in (0,1)} \{L(a_l|\pi) + E_X[C_{k-1}^*(\pi'_{a_l}(X))]\} \tag{1}$$

where

$$E_X[C_{k-1}^*(\pi'_{a_l}(X))] \equiv \text{expectation of the minimum total cost with respect to the next observation } X$$

and

$$\pi'_{a_l}(X) \equiv \text{probability the state is 0 given observation } x \text{ and action } a_l \text{ is selected, } \quad l = 0, 1$$

Using Bayes' rule gives

$$\pi'_{a_0}(x) \equiv \frac{(1 - \lambda)\pi f(x)}{\pi f(x) + (1 - \pi)g(x)}$$

and since a_1 is the *renew* action

$$\pi'_{a_1}(x) \equiv 1$$

For notational convenience, by writing $\pi'(x) = \pi'_{a_0}(x)$, we can rewrite Equation (1) as

$$C_k^*(\pi) = \min \left\{ \begin{array}{l} (1 - \pi)c_1 + \int_x C_{k-1}^*(\pi'(x))h_X(x) dx \\ d_1 + C_{k-1}^*(1) \end{array} \right\}$$

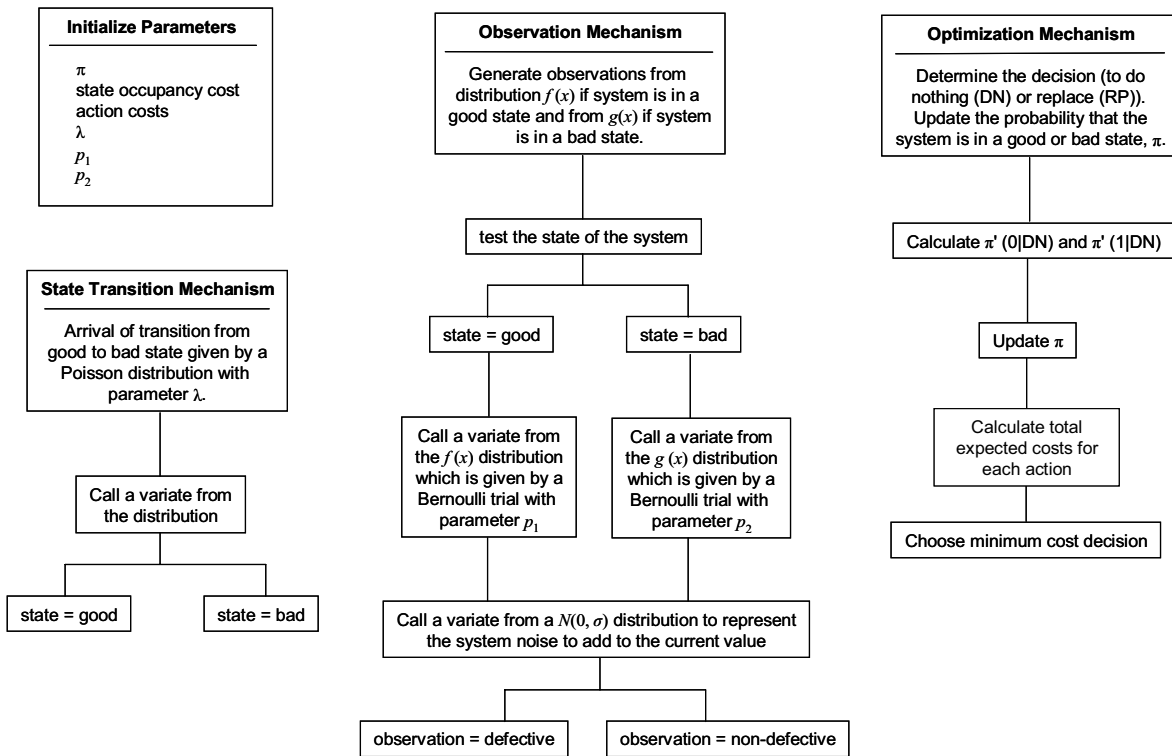


Figure 6. Flow chart of simulation methodology for two-state variables

where

$$h_X(x) \equiv \pi f(x) + (1 - \pi)g(x)$$

and $C_M^*(\pi)$ is the optimal return for an M -period finite horizon. When $M < \infty$ we assume that

$$C_o^*(\pi) = 0, \quad \forall \pi$$

i.e. the salvage value of the system (the value of the system when no observations are remaining and the process terminates) is zero, regardless of the state in which the system is operating.

4.2. Simulation methodology

Figure 6 details the simulation methodology for the two-state model. The State Transition Mechanism defines the underlying actual system condition that is only visible via incomplete and noisy observations. As described in the model formulation, state transition rates are assumed to be constant following an exponential distribution. We therefore model the arrival of the transition from the good to the bad state with the Poisson distribution. The Observation Mechanism generates observations for the good and bad states from state conditional Bernoulli distributions with a variate from a normal distribution to represent system noise. These observations are fed into the Optimization Mechanism which implements the dynamic program for the POMDP, selecting the cost minimizing maintenance action as a function of the belief about the system condition (π).

4.3. Results

For the two-state model, we developed a simulation code in Visual Basic for Excel according to the simulation flow chart in Figure 6. The parameter values shown in Table I are used for the POMDP state occupancy

Table I. Parameter values for two-state example

	State occupancy cost	Action cost	Observation distribution	White noise distribution	Failure rate
i	c_i , per unit time	d_i , fixed	Bernoulli	$N(\mu, \sigma)$	Exponential
0	0	0	$p_1 = 0.2$	$\mu = 0$	$\lambda = 0.05$
1	2	1	$p_2 = 0.8$	$\sigma = 0.2$	

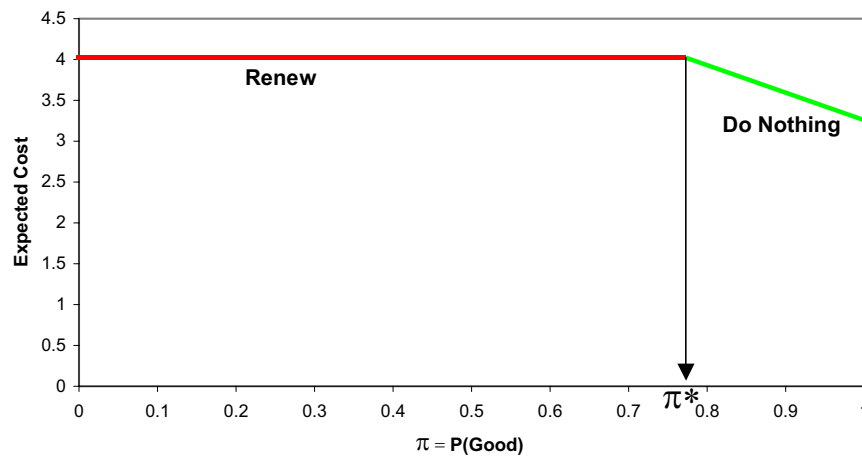


Figure 7. Illustration of optimal probability threshold for the two-state example




Figure 8. POMDP accuracy as a function of initial belief state, π

and action costs, and the state, observation and noise distributions for the simulation. We use the simulation program to identify the optimal probability threshold for the example defined by the parameter values in Table I. As shown in Figure 7, the optimal probability threshold is $\pi^* \approx 0.77$ when there are 15 observations remaining (i.e. $k = 15$), i.e. if $\pi < \pi^*$ the machine should be replaced.

In addition, the simulation model is used to evaluate the accuracy of the POMDP decision rules. Figure 8 details the impact of the decision maker's initial state of knowledge, π , on the accuracy of the decision selected by the POMDP. Accuracy is defined as the probability that the decision maker makes the correct decision, i.e. the probability that the decision maker replaces when the system has failed and does nothing when the system is good. It should be noted that the variability in the estimate of the probability correct is approximately 0.16.

Table II. Impact of observation distribution parameters on optimal probability threshold for two-state example (assuming all other parameter values remain fixed as shown in Table I)

	Observation distribution for good state Bernoulli p_1	Observation distribution for bad state Bernoulli p_2	Optimal probability threshold π^*
Original	0.2	0.8	0.77
a	0.3	0.6	0.83
b	0.4	0.4	0.86
c	0.01	0.9	0.71

While there is not a distinct pattern in the accuracy as a function of π , the accuracy of the POMDP model appears to improve as the π approaches 1.

It is interesting to note that the optimal probability threshold is relatively insensitive to changes in the state conditional observation distribution parameter values. Table II shows that when the state conditional observation distribution parameters are changed slightly, from $p_1 = 0.2$ and $p_2 = 0.8$ to $p_1 = 0.3$ and $p_2 = 0.6$ (making the observation distributions for the bad and good states less distinct), the optimal threshold increases slightly to $\pi^* \approx 0.83$. Similarly, if the observation distribution parameters are changed such that they are independent of the system state, e.g. $p_1 = 0.4$ and $p_2 = 0.4$, the optimal threshold increases slightly to $\pi^* \approx 0.86$. In contrast, when the observation distributions are distinct for each system state, e.g. $p_1 = 0.01$ and $p_2 = 0.9$, the optimal threshold decreases slightly to $\pi^* \approx 0.70$. In general, the optimal threshold is robust under changes to the state conditional observation distribution parameters—the threshold increases slightly when the observation distribution parameters are more similar and decreases slightly when these parameters are more distinct and provide more information about the underlying system state.

5. SUMMARY AND FUTURE RESEARCH

In this paper, we present a methodology for monitoring and maintenance decision making for deteriorating systems under conditions of uncertainty that integrates SQC and the POMDP through simulation in order to develop user-friendly heuristics. Specifically, we use advanced SQC techniques for real-world system monitoring. These techniques are used to define the observation distributions and structure for a POMDP. The integration of SQC and POMDPs through simulation allows for a realistic representation of a true production or service environment. This methodology enhances the POMDP model that assumes known observation distributions and process characteristics that are noise free. We incorporate the inherent process noise through simulation of ‘noisy observations’. Further, we use the simulation model to perform sensitivity analysis and investigate the robustness of the POMDP optimal solutions, i.e. decision policy (sequence), optimal threshold, and decision accuracy with respect to cost and process parameter variation.

We believe this work lays a foundation for future research that will help bring maintenance decision models into practice. In future investigations, we hope to use the methodology to develop a numerical solution to an efficient frontier of optimal policies for the POMDP as a function of process characteristics, system operating and maintenance costs. We will explore the competing objectives of the producer and consumer through the simulation model by developing efficient frontiers. Initially, we will explore the efficient frontier for the producer’s objective of minimizing total expected cost versus the consumer’s objective of maximizing accuracy, i.e. as defined by the probability that the decision maker makes the ‘correct’ decision. Other extensions include investigating the impact of the cost minimizing decisions on production, downtime, service, etc. We will also investigate utilizing these efficient frontiers for developing cost premiums as a function of quality. Furthermore, although our focus will be on finite horizon optimization, we will analyze the impact of the decision horizon length on the optimal decision strategy and evaluate the necessity of incorporating a discount factor.

Beyond the theoretical and empirical aspects of this future work, we hope that this work will lead to a broader understanding of maintenance by extending the application of these ideas beyond machine maintenance to product, service, and health maintenance. We expect to use our methodology to bring the maintenance theory into practice in an array of important problems in the industrial and medical arenas.

REFERENCES

1. Wang H, Pham H. Some maintenance models and availability with imperfect maintenance in production systems. *Annals of Operations Research* 1999; **91**:305–318.
2. Gautreau N, Yacout S, Hall R. Simulation of partially observed Markov decision process and dynamic quality improvement. *Computers in Industrial Engineering* 1997; **32**:691–700.
3. Chen M, Feldman R. Optimal replacement policies with minimal repair and age-dependent costs. *European Journal of Operational Research* 1997; **98**:75–84.
4. Chen D, Trivedi K. Closed-form analytical results for condition-based maintenance. *Reliability Engineering and System Safety* 2002; **76**:43–51.
5. Hontelez J, Burger H, Wijnmalen J. Optimum condition-based maintenance policies for deteriorating systems with partial information. *Reliability Engineering and System Safety* 1996; **51**:267–274.
6. Iravani S, Duenyas I. Integrated maintenance and production control of a deteriorating production system. *IIE Transactions* 2002; **34**:423–435.
7. Sinuany-Stern Z, David I, Biran S. An efficient heuristic for a partially observable Markov decision process of machine replacement. *Computers in Industrial Engineering* 1997; **24**:117–126.
8. Cremona C, Lukic M. Probability-based assessment and maintenance of welded joints damaged by fatigue. *Nuclear Engineering and Design* 1998; **182**:253–266.
9. Margalot N. Selecting a quality control attribute sample: An information-economics method. *Annals of Operations Research* 1999; **91**:83–104.
10. von Collani E. Control of production processes subject to random shocks. *Annals of Operations Research* 1999; **91**:289–304.
11. McCall J. Maintenance policies for stochastically failing equipment: A survey. *Management Science* 1965; **11**:493–524.
12. Pierskalla W, Voeller J. A survey of maintenance models: Control and surveillance of deteriorating systems. *Naval Research Logistics Quarterly* 1976; **23**:353–388.
13. Monahan G. A survey of partially observable Markov decision processes: Theory, models, and algorithms. *Management Science* 1982; **28**:1–16.
14. Valdez-Flores C, Feldman R. A survey of preventative maintenance models for stochastically deteriorating single-unit systems. *Naval Research Logistics* 1989; **36**:419–446.
15. White C III, White D. Markov decision processes. *European Journal of Operational Research* 1989; **39**:1–16.
16. Luss H. Inspection policies for a system which is inoperable during inspection periods. *IIE Transactions* 1977; **9**:189–194.
17. Rosenfield D. Markovian deterioration with uncertain information. *Operations Research* 1976; **24**:141–155.
18. Hopp W, Wu S. Multiaction maintenance under Markovian deterioration and incomplete state information. *Naval Research Logistics* 1988; **35**:447–462.
19. White C III. A Markov quality control process subject to partial observation. *Management Science* 1977; **23**:843–851.
20. White C III. Optimal inspection and repair of a production process subject to deterioration. *Journal of Operational Research Society* 1978; **29**:235–243.
21. White C III. Bounds on optimal cost for replacement problem with partial observations. *Naval Research Logistics Quarterly* 1979; **6**:415–422.
22. White C III. Monotone control laws for noisy, countable state Markov chains. *European Journal of Operational Research* 1980; **5**:124–132.
23. Shiryaev A. An optimum methods in quickest detection problems. *Theory of Applied Probability* 1968; **8**:22–46.
24. Taylor H. Markovian sequential replacement processes. *The Annals of Mathematical Statistics* 1965; **36**(6):1677–1694.
25. Teneketzis D, Varaiya P. The decentralized quickest detection problem. *IEEE Transactions on Automatic Control* 1984; **AC-29**:641–644.
26. Lovejoy W. Some monotonicity results for partially observed Markov decision processes. *Operations Research* 1987; **35**:736–743.
27. Simmons J. Maintenance and replacement policies for a multi-state deteriorating process with probabilistic monitoring. *PhD Thesis*, Industrial and Operations Engineering Department, University of Michigan, Ann Arbor, 1998.
28. Ivy J, Pollock S. Maintenance of a deteriorating process with probabilistic monitoring and silent failures. *American Mathematical Society Contemporary Mathematics Series* 2001; **275**:69–93.

Authors' biographies

Julie Simmons Ivy is an Assistant Professor of Business Administration at the University of Michigan. Her research interests include mathematical modeling of stochastic dynamic systems with emphasis on statistics and decision analysis. Her PhD degree is in Industrial and Operations Engineering from The University of Michigan. She is a member of INFORMS and the Society for Medical Decision Making.

Harriet Black Nembhard is an Associate Professor of the Harold and Inge Marcus Department of Industrial and Manufacturing Engineering at the Pennsylvania State University. She has previously held manufacturing positions with Dow Chemical, General Mills, and Pepsi-Cola. Her PhD degree is in Industrial and Operations Engineering from The University of Michigan. She is a member of ASQ, IIE, and INFORMS.