Using Markov Models of Fault Growth Physics and Environmental Stresses to Optimize Control Actions

Brian Bole [*]

Georgia Institute of Technology, Atlanta, GA, 30332

Kai Goebel [†]

NASA Ames Research Center, Moffett Field, CA, 94035

George Vachtsevanos [‡]

Georgia Institute of Technology, Atlanta, GA, 30332

A generalized Markov chain representation of fault dynamics is presented for the case that available modeling of fault growth physics and future environmental stresses can be represented by two independent stochastic process models. A contrived but representatively challenging example will be presented and analyzed, in which uncertainty in the modeling of fault growth physics is represented by a uniformly distributed dice throwing process, and a discrete random walk is used to represent uncertain modeling of future exogenous loading demands to be placed on the system. A finite horizon dynamic programming algorithm is used to solve for an optimal control policy over a finite time window for the case that stochastic models representing physics of failure and future environmental stresses are known, and the states of both stochastic processes are observable by implemented control routines. The fundamental limitations of optimization performed in the presence of uncertain modeling information are examined by comparing the outcomes obtained from simulations of an optimizing control policy with the outcomes that would be achievable if all modeling uncertainties were removed from the system.

I.   Introduction

Predicting the future dynamics of component fault modes is typically complicated by a lack of comprehensive run to failure data-sets that could be used to develop and validate prognostic models, and a limited knowledge of the future stresses to be placed on a system's components. However, even highly uncertain prognostic predictions may still provide invaluable information that can be utilized to improve the operation of safety critical and expensive to maintain systems such as aircraft and spacecraft.

Two types of modeling uncertainties affecting prognostic predictions are considered in this publication; these are classified as uncertainty incorporated into physics of failure models and uncertainty in the estimation of future exogenous stresses on a system. Uncertainty in estimating the dynamics of component degradation as a function of the loads or stresses applied to a specimen is often included as a noise term in lumped parameter and data driven fault growth prediction models.[1,2] A bounded environmental disturbance term is typically used to incorporate uncertain modeling of future exogenous system stresses into the analysis of robust and reconfigurable control design techniques, such as H_∞ control, L_1 control, and gain scheduling.[3,4] This paper is motivated by a desire to analyze an analytical formulation of the prognostics-based decision making problem, in which models of fault growth physics and environmental stresses are represented by two independent stochastic processes. An invented example of a dice throwing game is described to illustrate some fundamental challenges associated with assessing and managing the prognostic uncertainty introduced by incorporating stochastic modeling of physics of failure and future environmental stresses. Section II of the paper introduces a generalized representation of component degradation dynamics based on the modeled dynamics of future environmental loading conditions, the estimated statistics of a process noise term, and the actions taken by onboard

---

[*]Graduate Research Assistant, Intelligent Control Systems Laboratory, AIAA Member.
[†]Director of NASA AMES Prognostics Center of Excellence, AIAA Member.
[‡]Professor Emeritus and Director of the Intelligent Control Systems Laboratory at Georgia Tech, AIAA Member

American Institute of Aeronautics and Astronautics

control routines in response to observations of environmental and component health states. Section III describes a contrived but challenging example of a system in which uncertainty in the modeling of fault growth physics is represented by a uniformly distributed dice throwing process, and uncertain modeling of future environmental loading demands on the system is represented by a discrete random walk. The use of finite horizon dynamic programming to identify control policies that optimize the expectation of a given cost function over a finite time window is discussed in Section IV. Simulation results for selected optimized control policies on the dice throwing example are presented and analyzed in Section V.

II.   Stochastic Modeling of Fault Growth Dynamics

Consider a Markov chain representation of component damage accumulation that is expressed in terms of a metric representing the load applied to a specimen at each control time-step,

$$p_{i,j}(u) = \Pr\left(\gamma(k+1) = s_j | \gamma(k) = s_i, u(k) = u\right), \quad \gamma, s_i, s_j \in S, \quad u \in U(k), \quad k \in \mathbb{N} \tag{1}$$

$$\sum_{s_j \in S} p_{i,j} = 1 \tag{2}$$

where $U(k)$ represents a domain of feasible loads that may be applied at a given time-instant, γ is a random variable representing a component's state of health (SOH), S represents a quantized state space for component SOH, and $p_{i,j}(u)$ represents the probability of transitioning from SOH s_i to SOH s_j given an applied load u. The component loading variable may represent pressure, force, torque, or a wide variety of other stressors that drive component damage. Eq. (2) specifies that the sum of all transition probabilities defined at each system state in the Markov chain representation of the damage accumulation process must always equal one.

Let the component degradation dynamics be defined by the following generic nonlinear mapping:

$$\gamma(k+1) = f\left(\gamma(k), u(k), \xi(k)\right) \tag{3}$$

where $\xi(k)$ is a random variable representing uncertainty in physics of failure models. In this case, knowledge of the mapping function f and the statistics of ξ could be used to evaluate the transition probability matrix defined in Eq. (1).

The component loading profile $u(k)$ is assumed to be dictated partly by the dynamics of the system's operating environment, which may not be entirely predictable, and partly by the control actions taken in response to online observations of environmental states and the component SOH. The relationship between environmental conditions, control actions selected by onboard control routines, and component loading is considered to be generically represented as:

$$u(k) = \rho(k) \cdot w(k) \tag{4}$$

where w is taken to represent the net load that would be exerted by a nominally controlled component at a given time-instant, and ρ is taken to represent a relative deviation from nominal component loading that is induced by the actions of an implemented prognostics-based control routine. The response of a nominal control law to variations in environmental conditions is assumed to be known, thus a stochastic model of w can be directly derived from a stochastic model of environmental conditions. A formal description of potential challenges and benefits associated with this definition of the action space for onboard control routines is given in ref. [5].

Substitution of Eq. (4) into Eq. (3) yields:

$$\gamma(k+1) = f\left(\gamma(k), \rho(k) \cdot w(k), \xi(k)\right) \tag{5}$$

Component SOH dynamics are now defined in terms of the control actions taken at each control time-instant and the joint distribution of uncertain models for environmental conditions and fault growth physics. Given that the two sources of uncertainty can generally be assumed to be independent of each other, Eq. (1) can be expressed in terms of the metric ρ as follows:

$$p_{i,j}(\rho) = \sum_{w \in W} \sum_{\xi \in \Xi} \Pr\left(w(k) = w\right) \cdot \Pr\left(\xi(k) = \xi\right) \quad \text{where } f\left(s_i, \rho \cdot w, \xi\right) = s_j \tag{6}$$

where W and Ξ represent quantized state spaces for w and ξ respectively. The random variable ξ is considered to be represented by a stationary model of process noise and a potentially nonstationary distribution is considered for

American Institute of Aeronautics and Astronautics

representing exogenous inputs to the system. A process model for w is generically represented using the notation of conditional probabilities:

$$\Pr\left(w\left(k\right)=w\right)=\Pr\left(w\left(k\right)=w|w\left(k-1\right)=v\right)\cdot\Pr\left(w(k-1)=v\right),\quad\forall\,v\in W,\quad k\in[a+1,\infty)\qquad(7)$$

where a possibly uncertain observation of w at time-index a is used to initialize the random process.

The two dimensional state transition matrix given in Eq. (1) can now be expressed as a four dimensional matrix that incorporates the process models for ξ and w.

$$p_{(i,j),(l,m)}\left(\rho\right)=\sum_{\xi\in\Xi}\Pr\left(w\left(k\right)=v_l|w\left(k-1\right)=v_m\right)\cdot\Pr\left(\xi\left(k\right)=\xi\right)\quad\text{where }f\left(s_i,\rho\cdot v_m,\xi\right)=s_j,$$

$$s_i,s_j\in S,\,v_l,v_m\in W\qquad(8)$$

where $p_{(i,j),(l,m)}\left(\rho\right)$ represents the probability of transitioning from SOH s_i and exogenous loading demand v_l to SOH s_j and exogenous loading demand v_m. This notation enables available knowledge of system kinematics, fault growth modeling, exogenous demand modeling, measurement uncertainties, and modeling uncertainties to be represented by a finite set of state transition probabilities. The state transition probabilities may be chosen to approximate some analytical formulation for the stochastic process, or they may be identified purely by fitting a system model to a history of observations of the fault growth process, as is the case with hidden Markov model learning techniques.[6,7]

III. A Dice Throwing Analogue to the General Prognostics-Based Control Problem

This section considers an invented dice throwing game that is intended to illustrate some of the fundamental difficulties in assessing and managing the modeling uncertainty introduced by incorporating stochastic models of fault growth physics and environmental loading demands.

Consider a component degradation model of the form:

$$\gamma\left(k+1\right)=\gamma(k)-\lambda\cdot|u\left(k\right)|\cdot\xi\left(k\right)\qquad(9)$$

where the rate of component health lost is defined to be proportional to the magnitude of component load u multiplied by a process noise variable ξ, and λ represents a constant of proportionality.

Substituting Eq. (4) into Eq. (9) yields:

$$\gamma\left(k+1\right)=\gamma(k)+\lambda\cdot\rho\left(k\right)\cdot|w\left(k\right)|\cdot\xi\left(k\right)\qquad(10)$$

where ρ represents a degree of freedom that control routines may use to regulate component loading at the expense of degrading a system's nominal performance, and w represents exogenous loading demands.

Exogenous loading demands will be taken to be represented by the sum of consecutive throws of a three sided die with faces numbered {-1,0,1}, constituting a discrete random walk process. The probability mass function for w at time-index k is:

$$\Pr\left(w\left(k\right)=w|w\left(k-1\right)=v\right)=\begin{cases}\frac{1}{3}&\omega-v\in\{-1,0,1\}\\0&\text{else}\end{cases},\quad k\in[a+1,\infty)\qquad(11)$$

where the variable a represents the time-index at which prognostic predictions are made.

Process noise in the fault growth model will be taken to be represented by independent throws of a six sided die with faces numbered $\{.7,.8,.9,1.1,1.2,1.3\}$. The probability mass function for ξ at time-index k is:

$$\Pr(\xi\left(k\right)=\xi)=\begin{cases}\frac{1}{6}&\xi\in\{.7,.8,.9,1.1,1.2,1.3\}\\0&\text{else}\end{cases},\quad k\in[a,\infty)\qquad(12)$$

Both w and γ are be assumed to be observable at all time-instants occurring before a, thus the fault growth processes may be represented by deterministic equations for time-indexes less than a.
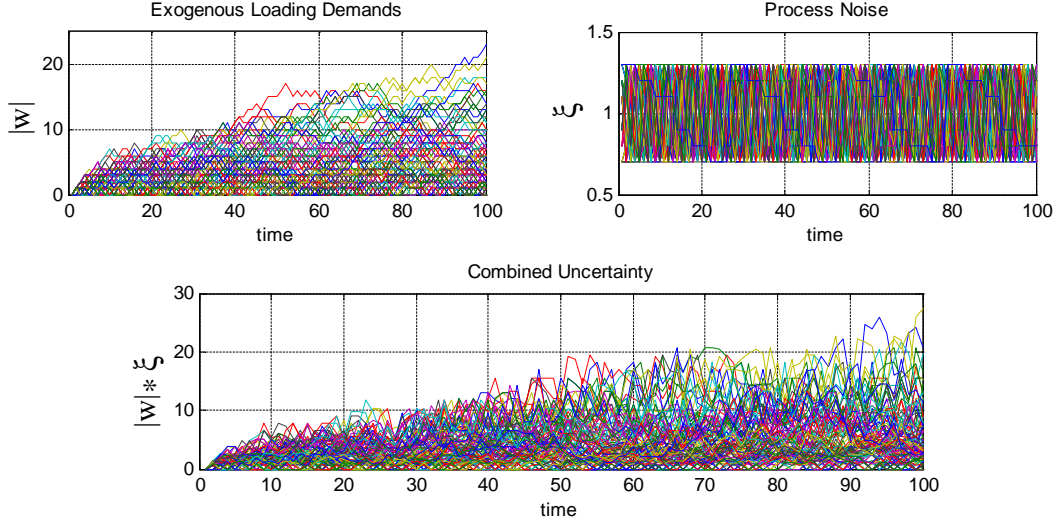
Figure 1: Time-series plots for 100 simulations of w and ξ (top) and the product of w and ξ (bottom).

A domain of feasible values for the control degree of freedom ρ is defined as $\rho \in \{.2, .3, .., 1\}$. The time-series behavior of the dice throwing system was simulated using a pseudo random number generator initialized with a unique 'seed' value in each simulation. Figure 1 shows time-series profiles for $|w|$, ξ, and $|w| \cdot \xi$ that were generated using a pseudo random number generator initialized with 100 unique seeds. Figure 2 shows simulation results for two sample control policies; one control policy commands no modifications to the system's nominal load allocations and the other control policy commands the maximum allowable degradation of the system's nominal load allocations. Note that modeling uncertainties included in these simulations becomes relatively large over the time-interval of interest, as is typical for prognostics-based control applications. Also note that the risk of component failure within the time frame of interest is seen to be relatively large for the case that the system's nominal loading performance is not at all degraded, and the risk of component failure is seen to be essentially nonexistent for the case that the system's loading performance is degraded to the maximum allowable extent.

The box plots shown on the right hand side of Figure 2 provide a convenient means of representing the statistics of the fault growth process. The top and bottom of the boxes plotted in Figure 2 represent the 25^{th} and 75^{th} percentiles of the data range at a given time-index, the notch in each box represents the median value of the data points, the dashed line represents the mean value, and the whiskers in the box plots extend to the most extreme points falling within the range,

$$q_1 - 1.5 \cdot (q_3 - q_1) < d_i < q_3 + 1.5 \cdot (q_3 - q_1) \tag{13}$$

where q_1 and q_3 are the 25^{th} and 75^{th} quantiles of the data respectively, and d_i represents a datapoint.

The next section describes the use of finite horizon dynamic programming for identifying control policies over a fixed time interval that optimize a given risk metric, where the risk metric quantifies a relative aversion to predictions of nominal loading performance degradations and predictions of fault growth over a given prognostic horizon.

IV.   Identifying Optimal Finite Horizon Control Policies With Dynamic Programming

The prognostics-based control problem can generally be viewed as an optimization problem, in which an implemented control routine will select values of ρ at each control time-step in an attempt to minimize the risks posed by the continued application of load to degrading components, while also minimizing deviation from a system's nominal output loading performance. The stochastic models given in Eqs. (10-12) yield predictions of a system's future deviation from its nominal output loading performance and predictions of future component health deterioration in terms of the controlled variable ρ.

If aversion to the potential degradation of a system's nominal output loading performance and aversion to the potential

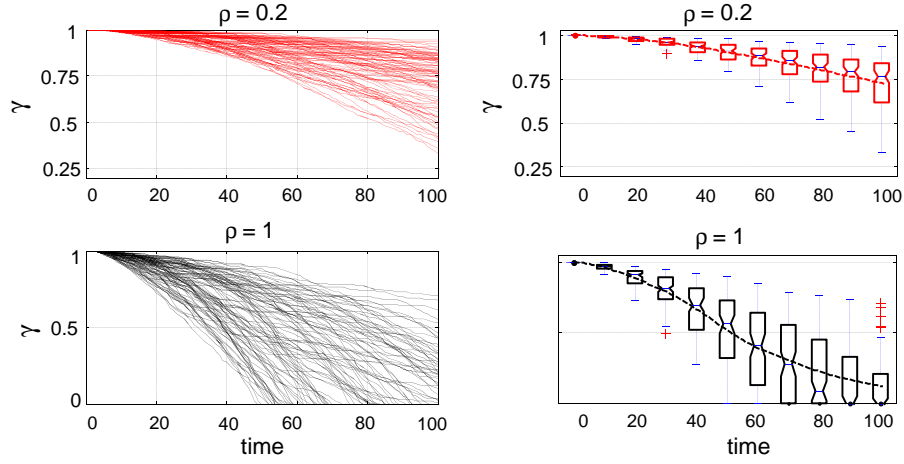American Institute of Aeronautics and Astronautics

Figure 2: Time-series data (left) and box plots (right) for 100 simulations of the dice throwing process with $\rho = 0.2$, representing the maximum allowable degradation of the system's nominal load allocation (top), and with $\rho = 1$, representing no degradation of the system's nominal load allocation (bottom).

degradation of component SOH are expressed in terms of an expectation of accumulated state transition costs, then the search for a control policy that optimizes a stochastic system of the form disclosed in Eq. (8) is expressed as a Markov decision process (MDP). MDPs are commonly used to analyze problems involving decision making in the presence of uncertain or stochastic information. Optimizing control policies for MDPs may be identified using well known MDP optimization techniques such as dynamic programming for finite horizon optimization problems, and linear programming, value iteration, or policy iteration for discounted and average-reward infinite horizon optimization problems.

An MDP formulation of the prognostics-based control problem will optimize a cost function that is defined as:

$$g_N\left(\gamma\left(N\right)\right) + \sum_{k=0}^{N-1} g_k\left(\gamma\left(k\right), \rho\left(k\right), w\left(k\right), \xi\left(k\right)\right) \tag{14}$$

where state transition costs, denoted by $g_k\left(\gamma\left(k\right), \rho\left(k\right), w\left(k\right), \xi\left(k\right)\right)$, assign a cost to the possibility of transitioning from one system state to another at time-index k, and a terminal cost, denoted by $g_N\left(\gamma\left(N\right)\right)$, penalizes the total degradation component SOH over a simulated time window.

If the values to be taken by the random variables w and ξ are known over the time window $k = [0, N-1]$, then Eq. (14) could be optimized directly; however, because future values of w and ξ are unknown in prognostic applications, the optimization problem is formulated in terms of the expectation of summed state transition costs,

$$E\left\{g_N\left(\gamma\left(N\right)\right) + \sum_{k=0}^{N-1} g_k\left(\gamma\left(k\right), \rho\left(k\right), w\left(k\right), \xi\left(k\right)\right)\right\} \tag{15}$$

A sequence of control actions taken over the domain $k = [0, N-1]$ is referred to as a control policy, and is denoted as $\pi = \{\mu_0, ..., \mu_{N-1}\}$, where μ_k maps observations of γ and w obtained at time-index k into a control action,

$$\rho\left(k\right) = \mu_k\left(\gamma\left(k\right), w\left(k\right)\right) \tag{16}$$

The expected cost of enacting a particular control policy π when starting at given initial values of γ and w is:

$$J_\pi\left(\gamma\left(0\right), w\left(0\right)\right) = E\left\{g_N\left(\gamma\left(N\right)\right) + \sum_{k=0}^{N-1} g_k\left(\gamma\left(k\right), \mu_k\left(x\left(k\right), w\left(k\right)\right), w\left(k\right), \xi\left(k\right)\right)\right\} \tag{17}$$

where an optimal control policy is defined as a one that minimizes this cost,

$$J_{\pi^*}\left(\gamma\left(0\right), w\left(0\right)\right) = \min_{\pi \in \Pi} J_\pi\left(\gamma\left(0\right), w\left(0\right)\right) \tag{18}$$

American Institute of Aeronautics and Astronautics

Here π^* represents an optimal control policy. Optimal control policies may be identified over a finite time window using the well known dynamic programming algorithm. This algorithm implements a backwards induction searching strategy to identify an optimal control policy over the time window $k \in \{N-2, N-1\}$, and then for $k \in \{N-3, N-2, N-1\}$ and so on, until the optimal policy is found over the entire window of interest.

$$J_N\left(\gamma\left(N\right), w\left(N\right)\right) = g_N\left(\gamma\left(N\right)\right)$$
$$\mu_k^*\left(\gamma, w\right) = \min_\rho E\left\{g_k\left(\gamma\left(k\right), \rho, w, \xi\left(k\right)\right) + J_{k+1}\left(f\left(\gamma, \rho, w, \xi\left(k\right)\right), w\left(k+1\right)\right)\right\}$$
$$J_k\left(\gamma, w\right) = E\left\{g_k\left(\gamma\left(k\right), \mu_k^*\left(\gamma, w\right), w, \xi\left(k\right)\right) + J_{k+1}\left(f\left(\gamma, \mu_k^*\left(\gamma, w\right), w, \xi\left(k\right)\right), w\left(k+1\right)\right)\right\}$$
$$\gamma \in S, \quad w \in W, \quad \pi^* = \left\{\mu_0^*, ..., \mu_{N-1}^*\right\}, \quad k = \{0, 1, ..., N-1\} \tag{19}$$

V.   Simulation Results

The dynamic programming algorithm given in Eq. (19) was used to identify an optimizing control policy over a given finite time window for the dice throwing game defined in Eqs. (10-12). State transition costs for the dice throwing game were designated to penalize the proportional deviation from the system's nominal output loading commands at each control time-index

$$g_k\left(\gamma\left(k\right), \mu_k\left(x\left(k\right), w\left(k\right)\right), w\left(k\right), \xi\left(k\right)\right) = 1 - \mu_k\left(x\left(k\right), w\left(k\right)\right) \tag{20}$$

The terminating cost was designated to be proportional to the square of component SOH at the end of the time-window.

$$g_N\left(\gamma\left(N\right)\right) = 100\left(1 - \gamma\left(N\right)\right)^2 \tag{21}$$

Figure 3a shows the distribution of component SOH observed at several time-increments over 100 repeated simulations of the control policy that optimizes the expected value of the cost function defined in Eq. (20) and (21). For comparative purposes, the profiles observed for the system's random variables in each simulation were fed into a non-causal implementation of dynamic programming, that solved for the optimal control policy given perfect information for the future values to be taken by the system's random variables. Figure 3b shows the distribution of component SOH profiles observed over repeated simulations of the optimal control policy using perfect future knowledge. The mean and standard deviation of the total costs evaluated over repeated simulations of the two optimized control policies and the two example control policies discussed in Section III are given in Table 1.

The results given in Table 1 show that the performance of the control policy identified by dynamic programming is clearly superior to the two sample control policies discussed in Section III. Comparison of the control policy computed using a stochastic model of the system with the control policy computed using perfect future knowledge, shows a somewhat more conservative behavior early in the mission from the control policy lacking prefect future knowledge, which is to be expected. Also, note that, even though the uncertainty included in the dynamics of the dice throwing game is seen to be substantial, relatively little variation is observed over repeated simulations of fault growth dynamics when the computed optimal control policy is used, indicating good disturbance rejection in the controlled system. The formulation and implementation of prognostics-based control on the dice throwing example is intended to be structurally analogous to the prognostics-based control problem on a wide variety of real-world applications, and future work will demonstrate the adaptation of the notational tools developed in this publication to more complex and practical problems.

	avg $\left(J_\pi\right)$	std(J_π)
Control policy enacting no deviation from nominal load allocations	81.6	26.7
Control policy enacting maximum deviation from nominal load allocations	79.6	10.5
Optimal control policy computed using stochastic modeling	65.7	21.1
Control policies computed using perfect future knowledge of each simulation	56.6	23.6

Table 1: Statistics of the total costs computed over repeated simulations of four simulated control policies

American Institute of Aeronautics and Astronautics

(a) Optimal control with uncertain modeling   (b) Optimal control with future knowledge
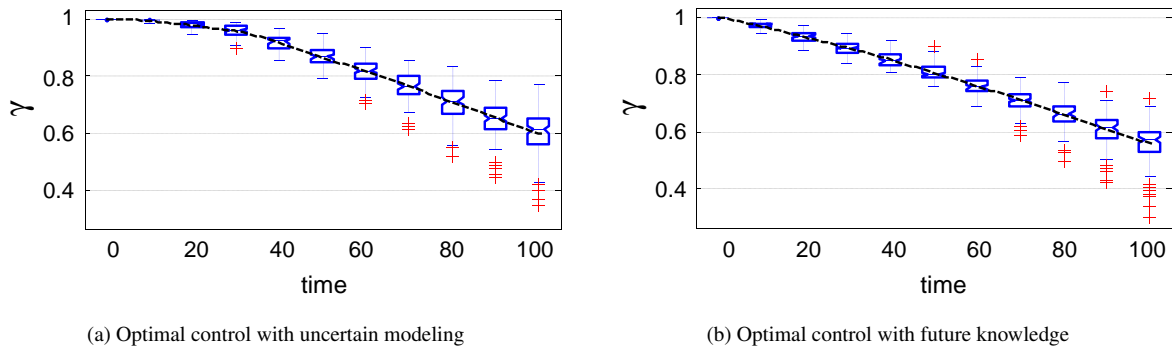
Figure 3: Plots of component SOH over repeated simulations of the optimal control policy computed with stochastic models (a) and plots of component SOH using an optimal control computed for the particular evolution of random variables observed in each simulation of the dice throwing game (b).

VI.   Conclusions

A contrived example of a dice throwing game was considered in order to provide some insight into the general problem of developing prognostics-based control routines that utilize uncertain models of component fault dynamics and future environmental stresses to assess and mitigate risk. A generalized Markov modeling representation of fault dynamics was developed for the case that available modeling of fault growth physics and available modeling of future environmental stresses may be represented by two independent Markov process models. A finite horizon dynamic programming algorithm was given for a generalized Markov decision process representation of the prognostics-based control problem and was used to identify an optimal control policy for the dice throwing game that is considered in this paper. The outcomes obtained from simulations of the optimizing control policy were observed to differ only slightly from the outcomes that would have been achieved if all modeling uncertainties were removed from the example dice throwing game.

Acknowledgment

References

[1] Orchard, M., Kacprzynski, G., Goebel, K., Saha, B., and Vachtsevanos, G., "Advances in Uncertainty Representation and Management for Particle Filtering Applied to Prognostics", *International Conference on Prognostics and Health Management PHM*, 2008

[2] Goebel, K. and Saha, B. "Uncertainty Management for Diagnostics and Prognostics of Batteries using Bayesian Techniques", *IEEE Aerospace Conference*, 2008

[3] Zhou, K., "A New Controller Architecture for High Performance, Robust, and Fault-Tolerant Control", *IEEE Transactions on Automatic Control*, 2001, V. 46, pp 1613-1619

[4] Zhang, Y., and Jiang, J., "Bibliographical Review on Reconfigurable Fault-Tolerant Control Systems", *Annual Reviews in Control*, 2008, V. 32, pp 229-252

[5] Bole, B., Goebel, K., and Vachtsevanos, G., "Stochastic Modeling of Component Fault Growth Over A Derived Domain of Feasible Output Control Effort Modifications" , *International Journal of Production Research*, 2012

[6] Baruah, P. and Chinnam, R. B., "HMMs for Diagnostics and Prognostics in Machining Processes", *International Journal of Production Research*, 2005, V. 43, pp 1275-1293

[7] Sankararaman, S., Ling, Y., Shantz, C., and Mahadevan, S., "Uncertainty Quantification in Fatigue Damage Prognosis", *International Conference of the Prognostics and Health Management*, 1982.