Available online at www.sciencedirect.com
SCIENCE DIRECT®
ELSEVIER
Automatica 42 (2006) 1321 – 1330
automatica
www.elsevier.com/locate/automatica

A risk-sensitive approach to total productive maintenance☆

Abhijit Gosavi*

Department of Industrial Engineering, University at Buffalo, 317 Bell Hall, SUNY, Buffalo, NY 14260, USA

Received 16 March 2005; received in revised form 16 January 2006; accepted 4 February 2006
Available online 18 May 2006

Abstract

While risk-sensitive (RS) approaches for designing plans of total productive maintenance are critical in manufacturing systems, there is little in the literature by way of theoretical modeling. Developing such plans often requires the solution of a discrete-time stochastic control-optimization problem. Renewal theory and Markov decision processes (MDPs) are commonly employed tools for solving the underlying problem. The literature on preventive maintenance, for the most part, focuses on minimizing the *expected* net cost, and disregards issues related to minimizing risks. RS maintenance managers employ safety factors to modify the risk-neutral solution in an attempt to heuristically accommodate elements of risk in their decision making. In this paper, our efforts are directed toward developing a formal theory for developing RS preventive-maintenance plans. We employ the Markowitz paradigm in which one seeks to optimize a function of the expected cost and its variance. In particular, we present (i) a result for an RS approach in the setting of renewal processes and (ii) a result for solving an RS MDP. We also provide computational results to demonstrate the efficacy of these results. Finally, the theory developed here is of sufficiently general nature that can be applied to problems in other relevant domains.

Keywords: Variance; Renewal reward theorem; Markov decision processes; Maintenance; Risk sensitive

1. Introduction

Total productive maintenance (TPM) is a management initiative that has been widely embraced in the industry. A positive *strategic* outcome of such implementations is the reduced occurrence of unexpected machine breakdowns that disrupt production and lead to losses which can exceed millions of dollars annually. Additionally, frequent machine breakdowns indirectly can lead to a host of other problems, e.g., difficulties in meeting customer deadlines, which makes the transition from make-to-stock to make-to-order difficult (Suri, 1998) and magnifies the need to keep extra safety stocks, increasing inventory-holding costs (Askin & Goldberg, 2002). An important tool of a TPM program is the stochastic model used to determine the optimal time for preventive maintenance (PM) (Askin & Goldberg, 2002). PM can help reduce the frequency of unexpected repairs when the failure *rate* is of an increasing nature (Das & Sarkar, 1999; Lewis, 1994).

Renewal processes (Kao, 1997; Ross, 1992) and Markov decision processes (MDPs) (Bertsekas, 1995; Puterman, 1994) are frequently used as the underlying stochastic models in a TPM program. A critical drawback of a traditional approach in TPM is to use the *expected* value of the long-run cost as the objective function. Such an approach overlooks the risk associated with the occasional high cost that can occur in system optimized with respect to the expected cost. As a result, risk-sensitive (RS) managers, whom we interacted with in a local automobile industry, modify the predicted optimal (with respect to the expected cost) time for maintenance, τ^*, by using a factor of safety, η, where $\eta > 1$, such that the time for PM is then: τ^*/η. While this certainly results in a more *conservative* time for maintenance, it is a heuristic approach. What managers really need is a more sophisticated approach that would help them (i) quantify their risk sensitivity on a scale from 0 to 1 and (ii) determine the optimal maintenance time using a model that incorporates this factor. This clearly motivates the need for

☆ This paper was not presented at any IFAC meeting. This paper was recommended for publication in revised form by Associate Editor Qing Zhang under the direction of Editor Suresh Sethi.

* Tel.: +1 716 645 2357; fax: +1 716 645 3302.
E-mail address: agosavi@buffalo.edu.

0005-1098/$ - see front matter © 2006 Elsevier Ltd. All rights reserved.
doi:10.1016/j.automatica.2006.02.006

embedding the well-known Markowitz criterion (Markowitz, 1952) within the stochastic model. Another significant demand of managers from the model is the ability to *quantify* risks in terms of dollars (or Euros) and hours—units that they are comfortable with. In particular, senior managers involved in developing long-term plans for an enterprise are familiar with the idea of using variance per unit time as a measure of risk in *strategic* decision making (see Ruefli, Collins, & Lacugna, 1999 for an extensive survey). Since TPM has a significant strategic impact on the organization, the units of risk in these calculations should ideally match those used in strategic management. One of the goals of this research is to develop models that can be conveniently used by managers.

A general cost function (objective function) using the Markowitz criterion is

$$g(\tau) = \mu_C + \theta\sigma^2 \quad \text{with } \theta > 0, \tag{1}$$

where μ_C and σ^2 denote the long-run mean and the long-run variance, respectively, of the net cost per unit time incurred from following a preventive maintenance plan that prescribes τ as the time for PM. An alternative formulation in terms of *rewards*, in which the objective function is *maximized*, is $g_R(\tau) = \mu_R - \theta\sigma^2$, with $\theta > 0$, where μ_R and σ^2 denote the long-run mean and variance of the net *reward* per unit time, respectively. Since $\mu_R = -\mu_C$, both formulations are *equivalent*.

Risk-neutral (RN) statistical models for PM use $\theta = 0$. Typically, θ is selected by experimentation by the manager and is a function of the variability in the system. A very large value for θ is undesirable, since that could produce a solution with a very low variability but also with a very high cost. This is because a very large value for θ amplifies the importance of the variance and diminishes that of the mean. A very low value for θ, on the other hand, is indicative of a manager who is neutral to risks. Clearly, the smaller the value of θ, the closer the model gets to becoming RN.

The time for PM, it must be understood, is the time since the last repair or PM. A common assumption is that the unit or the line is as good as new when it is repaired or preventively maintained. A typically made second assumption is that when the machine is not working, it is assumed not to age. We will stick to these two assumptions here. The main focus of this paper is to develop a theory when $\theta > 0$. The work of Chen and Jin (2003) also employs the Markowitz criterion, but their approach is quite different than ours; this will be clarified via our discussions below.

TPM plans for the production line in its entirety tend to be distinct from those for individual units that operate independently of the line. Most factories are full of such units, e.g., fork-lift trucks, electrical pumps, etc. We will develop separate models for the individual-unit scenario and the production-line scenario. For the case of the individual unit, we will present a renewal-theory model and for the case of the production line, we will present a more involved model based on MDPs. The analysis will involve presentation of some key results that could be applied to a large number of other management-science problems involving control theory. Thereafter, we will present results from computational experiments with both models. The remainder of this paper is organized as follows. Section 2 presents the renewal-theory model, and Section 3 presents the Markov decision model. Section 4 describes empirical work done using these two models, and Section 5 concludes the paper.

2. A renewal-theory model

A commonly used model in most TPM programs employs renewal theory and goes by the name “age-replacement”. In the setting of renewal processes, every failure or a maintenance triggers a so-called *renewal event*. A classical result in renewal theory, called the renewal reward theorem, provides an expression for the expected reward per unit time in a renewal process. An extension of this concept to the variance in the rewards of the renewal process can be found in Chen and Jin (2003). However, we present a different result that leads to a significantly different mechanism for measuring variance. Our result was influenced by the need of managers in a local industry for measuring risk in practical units that could be explained to senior management; the unit of risk (variance) in our result is hence *dollar*^{2}/*hour* or *Euro*^{2}/*hour*, which is the same as that for variance in rewards per unit time. As mentioned above, strategic managers are known to use variance to measure risk (Ruefli et al., 1999).

Consider a counting process, $\{N(t), t \geqslant 0\}$, and let T_n denote the time between the $(n-1)$th and the nth event in this process with $n \geqslant 1$. If $\{T_1, T_2, \ldots\}$ denotes a sequence of non-negative random variables that are independent and identically distributed, then the counting process is called a *renewal* process. When there is a reward associated with each renewal, we have a renewal reward process. Let R_n denote the reward associated with the nth renewal or nth cycle. We will let $R(t) = \sum_{n=1}^{N(t)} R_n$ denote the sum of the individual rewards earned by time t and $R^2(t) = \sum_{n=1}^{N(t)} (R_n)^2$ denote the sum of the square of the individual rewards earned by time t. Further, let

$$E[R] \equiv E[R_n], \quad E[R^2] \equiv E[(R_n)^2] \quad \text{and} \quad E[T] \equiv E[T_n],$$

where E denotes the expectation operator. The *long-run* variance in the rewards in Chen and Jin (2003) (see Lemma 1 of their paper) is $\lim_{t\to\infty} R^2(t)/t - [\lim_{t\to\infty} R(t)/t]^2$, which requires a subtraction of two quantities of which the first has the unit *dollar*^{2} per *hour* and the second *dollar*^{2} per *hour*^{2}. We present the following definition for the long-run variance, which measures *a time average of the total variance in infinitely many renewals*: $\sigma^2 = \lim_{t\to\infty} V(t)/t$, where $V(t) = \sum_{n=1}^{N(t)} (R_n - E[R])^2$. Thus σ^2 represents the sum of the squared deviations of the cycle rewards from their means, along an infinite number of renewals, divided by the total duration of the renewals. Hence, it can be interpreted as the average variance per unit time measured over the long run. What is key is that the above definition produces a consistent unit of *dollar*^{2} per hour for the variance. We now prove the following result to measure the variance in this style.