Heuristic Adaptive Replacement Policies

John Wilson

In the literature on replacement/maintenance policies for systems of stochastically failing machines, it is generally assumed that the failure time parameters are known and the objective is to find a replacement policy that minimizes expected cost per unit of time. In practice, however, the variability of costs can be an important consideration. In addition, most managers would want to adapt to new information obtained while operating the machines. In this paper, a class of policies is analyzed where the manager is allowed to adapt the replacement policy according to the statistical information obtained while operating the machines. We allow for Bayesian updating of failure parameters and demonstrated that it is feasible to explicitly calculate both the expected cost per unit time and the variance associated with this class of replacement policies.

The advent of inexpensive and powerful computer methodology is leading to an explosion of information available in all sorts of production settings. For instance, many new production machines incorporate software that tracks system failures, downtime, etc. This new availability of real time data presents interesting opportunities for the modeling of production systems. This is especially true for the scheduling of replacement and/or maintenance times. Most models assume that the parameters of the failure time distributions are known with certainty. In practice, a manager might decide to bring forward replacement or maintenance, if the history of the process to date indicates a more failure prone system than expected. However, once one allows this kind of flexibility, complex modeling issues arise. In particular, decisions made today have ramifications for the future evolution of the system. For instance, replacing a working system means the introduction of a censored observation into the data set and, consequently, has an impact on statistical inferences that will be made in the future.

The problem of finding a policy that incorporates information obtained while operating a machine has been in the literature for many years. Fox (1967) provided existence results regarding an optimal placement policy for the case of one machine with a Weibull failure time. Mazzuchi and Soyer (1996) and Silver and Fiechter (1995) provided heuristic Bayesian approaches for scheduling replacement in this situation. (However, to date, there is no nontrivial solution —heuristic or otherwise — for this problem where the expected cost can be explicitly calculated.)

In most of the maintenance literature, the decision maker is only interested in calculating the expected cost associated with the given policy. Most decision makers would also be interested in the variability of the cost associated with a given replacement policy. Most elementary statistics textbooks warn of the dangers of using location measures without some measure of dispersion. Consequently, it is surprising that most of the literature on replacement policies ignores the variance associated with the policy being used. Knowledge of the variance associated with the given policy is useful for a number of reasons: for a given policy (e.g., the one that minimizes expected cost per unit time) it provides a measure of the variability in costs that can be expected; in some instances, a decision maker might decide to use a suboptimal policy (in terms of expected cost) in order to obtain lower variability.

Allowing the parameters of the failure distribution to be unknown and incorporating information from past observations leads to very high dimensional problems — the so called “curse of dimensionality”. However, for some systems — such as the one considered in this paper — it is possible to develop tractable expressions for intuitively reasonable class of decision rules. The adaptive rules considered in this paper are essentially of the form “replace early if one observes many failures at the beginning of the process”. Such rules seem a natural and intuitive way to react to incoming data. Embedding such rules in the software that collects real time data seems a natural extension to the standard replacement and inspection models. The purpose of this paper is to provide a complete analysis for a system of parallel machines. It is important to note that, while the algebra might be complex, implementation and evaluation of the decision rules is straightforward.

We consider n machines with independent identically distributed lifetimes operating in parallel. A failed machine incurs downtime costs. A salvage value for functioning machines is assumed. It costs a fixed amount to purchase a batch of new machines. The costs can include fixed costs as well as the purchase price for n machines. A number of policies have been considered in the literature. An m-failure policy calls for replacement at the time of the mth failure (see Assaf and Shanthikumar (1987). The class of m-failure models can be thought of as being embedded in the class of shock models (see, e.g., Boland and Proschan (1983)). A T-policy calls for replacement every T units of time (see, e.g., Okumoto and Elsayed (1983)). An (m,T)-policy calls for replacement at the time of the mth failure or time T, which ever occurs first. It is generally assumed in the literature that the failure time distribution is completely specified. Assaf and Shanthikumar (1987) proved that an m-failure policy minimizes the long run expected cost per unit time when the machines have i.i.d. exponential failure times with known parameter λ. We assume a Bayesian approach where the underlying parameter is not known and the replacement policy incorporates both the prior distribution for and statistical information obtained while operating the machines.