Computers & Industrial Engineering 77 (2014) 103–114

Contents lists available at ScienceDirect

Computers & Industrial Engineering

journal homepage: www.elsevier.com/locate/caie

A condition-based maintenance strategy for heterogeneous populations ☆

Mimi Zhang [a,*], Zhisheng Ye [b], Min Xie [a]

[a] *Department of Systems Engineering and Engineering Management, City University of Hong Kong, Hong Kong*
[b] *Department of Industrial & Systems Engineering, National University of Singapore, Singapore*

ARTICLE INFO

Article history:
Received 20 February 2013
Received in revised form 15 June 2014
Accepted 2 September 2014
Available online 22 September 2014

Keywords:
Burn-in test
Gamma process
Mixture distribution
Preventive replacement
Wiener process

ABSTRACT

This paper develops a maintenance strategy, called inspection–replacement policy, to cope with heterogeneous populations. Burn-in is the procedure by which most of the defective products in a heterogeneous population can be identified and removed prior to being placed in service. However, modern manufacturing is so well developed that a defective product is able to function for a long period of time even under aggravated operational conditions. Instead of weeding defective products out via costly burn-in tests, use can be made of them in field operation where maintaining actions will be performed to prevent early in-use failures. The inspection–replacement policy consists of an inspection, conducted in an early stage with the purpose of identifying and replacing defective products, and a preventive replacement, carried out at a later stage to prevent wear-out failures. The preventive-replacement time is dynamically determined, depending on the information obtained by the inspection. The inspection–replacement policy is compared with a joint burn-in and age-based-replacement policy to show its practicability and competence.

1. Introduction

To maintain competitive advantage, manufacturers endeavor as much as possible to produce reliable products. However, during the manufacturing process, some unavoidable manufacturing defects could be introduced, e.g. defects in the raw material, leading to a heterogeneous population of products. The heterogeneous population contains a small proportion of weak/defective products. Compared with the normal products, the weak products have a shorter mean lifetime and are prone to giving rise to early in-use failures. The early in-use failures will cause substantial costs and sometimes are hazardous. In fact, it is not uncommon to observe a heterogeneous population with two sub-populations: a weak sub-population and a normal sub-population. For example, it is widely believed that integrated circuits consist of a small proportion of weak items which have much shorter mean lifetime. The GaAs laser data set, provided by Meeker and Escobar (1998), is a typical sample consisting of a sub-group of normal devices and a sub-group of weak devices; see Tsai, Tseng, and Balakrishnan (2011). Scarf and Cavalcante (2012) considered component heterogeneity when developing an age-based maintenance policy for a single-component system; they assumed that the population of components comprises a mixture of the weak and the strong. Berrade, Scarf, Cavalcante, and Dwight (2013) developed a maintenance policy involving periodic inspections, in which inspections are subject to error; they assumed that the time a system spends in the defective state is a random variable from a mixture distribution. Recent research on heterogeneous data can be found in Cha and Finkelstein (2013), Erisoglu, Erisoglu, and Çalis (2013) and Kazmi, Aslam, Ali, and Abbas (2013), among others.

A common practice to tackle a heterogeneous population is to screen out the weak products by means of a burn-in test. Burn-in is an engineering procedure implemented at the end of the manufacturing process. In a burn-in test, all the products are subjected to harsh electrical and thermal conditions that emulate the field operational conditions. At the end of the burn-in test, only the functioning products will be shipped to customers. There has been a bulk of research on developing economical burn-in tests; see Cha and Finkelstein (2010), Pearn, Hong, and Tai (2013), Post and Bhattacharyya (2012) and Yuan and Kuo (2010), among others. With the rapid development of modern manufacturing technology, even a weak product is able to operate for a fairly long period of time under aggravated operational conditions. Weeding out weak products via traditional age-based burn-in approach is therefore ineffectual. When there is a performance characteristic whose

---

☆ This manuscript was processed by Area Editor H.-S. Jacob Tsao.
* Corresponding author. Tel.: +852 5446 5352.
E-mail address: mmzhang5-c@my.cityu.edu.hk (M. Zhang).

http://dx.doi.org/10.1016/j.cie.2014.09.001
0360-8352/

evolution is closely associated with the lifetime of the product, the condition-based burn-in approach is an attractive alternative (Xiang, Coit, & Feng, 2013; Ye, Xie, Tang, & Shen, 2012). Because a weak product most often deteriorates faster than a normal product, a condition-based burn-in test exercises all the products for a certain period of time and scraps the products with deterioration levels higher than a pre-specified cut-off level. After the burn-in test, all the products with deterioration levels lower than the cut-off level are shipped to customers. In field operation, preventive repairs are often scheduled to further improve reliability and reduce operational costs. See Ahmad and Kamaruddin (2012) for an overview of two maintenance techniques widely adopted by industrial engineers: time-based maintenance and condition-based maintenance. Nonetheless, it is well known that burn-in is costly, including such as burn-in set-up cost, burn-in operational cost and repair/scrap cost from a burn-in failure. Scilicet, by adopting the burn-in procedure, we make a trade-off between early in-use failures and a reduced yield due to the burn-in costs. It is worth noting that nowadays even a weak product is able to operate for a long term. A more judgmatic approach to tackling a heterogeneous population is to directly put all the products into field operation and replace the weak products before they fail. By virtue of appropriately scheduled inspections, early in-use failures can be mitigated, and the failure cost can be countervailed by the long-time operating income. Even if a burn-in test is able to identify most of the defective products, it will cause damage to the normal products, shortening the mean lifetime of the normal products (Cha & Finkelstein, 2011; Ye, Tang, & Xie, 2011).

To make full use of weak products and avoid impairing normal products, we develop the inspection–replacement policy for heterogeneous populations with the assumption that a suitable performance characteristic of the products is available. The inspection–replacement policy directly puts all the products in a heterogeneous population into field use without burn-in. The physical degradation of each product will be measured at time $b(>0)$. If the degradation of a product exceeds a critical threshold, then it will be treated like a weak product and replaced at time b. If the degradation does not exceed the critical threshold, then it will be treated like a normal product and will be preventively replaced at time $R(>b)$, if it survives to time R. A schematic diagram of the inspection–replacement policy is shown in Fig. 1. The purpose of the inspection is to identify and replace weak products, whereas the preventive replacement is to replace aged normal products. If the inspection and the critical threshold are well determined, most of the weak products will be identified by the inspection. It is not difficult to see that the role of the inspection is analogous to a burn-in test except that the inspection is carried out in the field operation. However, there will be some figures of merit by replacing the burn-in approach with the inspection–replacement policy. First of all, the inspection–replacement policy does not need to burn-in products, and therefore saves burn-in costs and avoids impairing normal products. Secondly, the inspection–replacement policy makes full use of weak products, and therefore is expected to yield more profits. Last but not least, the preventive-replacement time R for each product is dynamically determined, depending on the physical degradation of the product obtained by the inspection. The on-line updating technique is expected to return more cost-efficient maintenance policy.

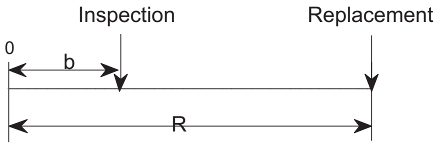

Fig. 1. Schematic of the inspection–replacement policy.

The rest of the paper is organized as follows. Section 2 details the inspection–replacement policy. The long-run cost rate function is derived. A joint burn-in and age-based-replacement policy serving as a benchmark is introduced. Section 3 draws the inspection–replacement policy on two predominant stochastic processes: the gamma process and the Wiener process. Section 4 presents numerical examples to elaborate the methodology developed in Sections 2 and 3. The GaAs laser data set is analyzed. Contradistinctive analysis and sensitivity analysis are performed by using numerical simulations. Section 5 concludes the paper.

2. Maintenance strategy & cost rate function

The product concerned stands for a component which is randomly drawn from a heterogeneous population. When put into operation, the component undergoes deterioration with the degradation process denoted by $\{X_t, t \geq 0\}$, assuming $X_0 = 0$ as with the convention. A component is considered to have failed if its deterioration level reaches a given failure threshold, denoted by $l(>0)$, which is a known constant. Define T to be a random variable representing the first hitting time of the degradation process $\{X_t, t \geq 0\}$ to the failure threshold l. The inspection–replacement policy is detailed as follows. After the installation at time 0, the component is subject to inspection at a pre-determined epoch $b(>0)$. The component will be treated like a weak component and will be replaced by a new component if it survives beyond the point, i.e. $T > b$, yet with the deterioration level X_b equating or exceeding a critical threshold $\vartheta(<l)$. The new backup component is randomly drawn from the same heterogeneous population. In the case that $T > b$ and $X_b < \vartheta$, a preventive-replacement time $R(>b)$ will be scheduled. R is the time point upon which the reliability of the component drops to a pre-determined reliability threshold δ $(0 < \delta < 1)$. Note that the preventive-replacement time R is on-line calculated, depending on the deterioration level X_b. In the case that the component survives beyond the preventive-replacement time, i.e. $X_b < \vartheta$ and $T > R$, it will be preventively replaced at time R by a new component. In the event that the component fails unexpectedly, i.e. $0 < T < b$ or $b < T \leqslant R$, it will be immediately replaced by a new component. The intention of the inspection is to screen out poor-quality components and hence to prevent early failures. The purpose of the preventive replacement is to reduce wear-out failures caused by normal components. Other mild assumptions are given as follows.

- Starting from the installation of a device, the wear trajectory is taken to have an upward trend though not necessarily monotonically increasing.
- Inspection is perfect in the sense that it reveals the true degradation level of a device and does not change the condition of the device.
- Replacement time is negligible compared to the expected lifetime of the devices. Failure is self-announcing and can be observed instantaneously.

Components with $X_b \geqslant \vartheta$ are replaced at the condition-monitoring point instantaneously, and the average cost of replacing a component is C_r. Preventive replacement of a component at time R is instantaneous and again costs C_r. The cost of each inspection is C_I, and the cost of a failure is C_f. Practical conditions define the constraints on the costs as follows: $C_I < C_r < C_f$. The decision variables of the inspection–replacement policy are the condition-monitoring epoch, b, and the corresponding critical threshold, ϑ. The reliability threshold, δ, is pre-determined by domain experts. A conservative engineer may set a high value of reliability threshold.

Let random variable V denote the length of a single replacement cycle, i.e. the time from the installation of a component to its