Microelectron. Reliab. Vol. 20. pp. 205-218
Pergamon Press Ltd. 1980. Printed in Great Britain

0026-2714/80/0601-0

ESTIMATION OF THE RELIABILITY OF COMPUTER COMPONENTS FROM FIELD RENEWAL DATA

David C. Trindade* and Larry D. Haugh**
*IBM General Technology Division, Essex Junction, VT 05452, USA
**University of Vermont, Burlington, VT 05405, USA

(Received for publication 21st January 1980)

Abstract - The assessment of the performance of electronic components in actual use is important to the manufacturers of these parts for many reasons. In this paper, we discuss the complexities involved in statistically estimating the reliability of computer components from field data on systems having different operating times. We also describe the operations of a specific department of the IBM Corporation involved with measuring the field performance of locally manufacturered product. The renewal method for estimating component reliability and the corresponding assumptions are explained, and the procedure is illustrated with an application to some actual field data. In particular, the estimated cumulative distribution function, hazard rate, and lognormal and Weilbull probability plots for this field example are shown.

INTRODUCTION

The tracking of the actual performance of electronic computer components in the field, from installation to the present time, is important to component manufacturers for several reasons. This endeavor allows the producers to estimate component reliability as a function of operating time, identify failure modes through laboratory analysis of the failed items, check current product reliability models, and improve reliability projections for future technologies [1].

In this paper, we briefly describe the operations of a specific department of the IBM Corporation charged with measuring the performance of locally manufactured product in use by customers. The attention in this description will be directed not at the administrative aspects involved in such a nationwide data collection effort nor at the efficient storage and handling of such a large data base; instead, the focus will be on the statistical analysis appropriate for estimating the component reliability function in time.

Several complications arise in measuring field reliability. Since the components in a system (e.g., a computer) are replaced quickly upon failure to restore a valuable system to operation, there exists at each component site an ordinary renewal process [2]. However, for a multi-component system, unless records are kept of the site of each failed component and its replacements, it becomes impossible after the first fail to determine whether any subsequent fails occur on original components or on components which are replacements for original units. For the systems we consider this site information is not available, and we call this situation one of unidentified replacement. In addition, since systems are installed at different times, at any specific reference time of analysis there will be different total operating hours for the systems under study. If every system is referenced to a common time origin, then we have a condition referred to as multiple or progressive censoring.

We have developed for the analysis of censored renewal data [3] a nonparametric statistical approach which can estimate component reliability from the time-of-replacement data on systems of components. The method is nonparametric in that there are no restrictions on the underlying probability distribution of component lifetimes. That is, the procedure will handle normal, lognormal, exponential, Weibull, etc., type data. It was necessary to develop a new estimation approach because the existing statistical methodologies could not handle the situation of multi-censored renewal data with unidentified replacements.

DESCRIPTION OF PRODUCT AND DATA GATHERING

The IBM facility located near Burlington, Vermont, manufactures large scale integrated (LSI) circuits for use in IBM products. In particular, the parts we consider are the memory chips utilized in computers.

The chips are mounted singly or in groups of two or more on a ceramic substrate to form a package called a module. A typical module may measure twelve to twenty-five millimeters square. The modules are attached onto a planar surface to form a card. A typical card, measuring about ten by sixteen centimeters, may have thirty-two memory storage modules and a few support function modules. A system may have any number of cards. The card is the field-replaceable unit in the sense that it can be plugged into a computer system. Thus, it is the component that is replaced in the field when failures occur in chips or modules.

Upon failure of a card component, an IBM customer engineer replaces the failed unit with a new one to restore the machine to operation, generally in a matter of hours. The fact that a card has been replaced is conveyed to a central data gathering office. It is important to note here that only the date of card replacement on a given system is recorded and not the identity (via location in the machine, component serial number, or otherwise) of the individual card. Hence, except for single card systems, it becomes impossible after the first failure on a system to determine if a subsequent failure occurs on an original card or on a card that is itself a replacement for a previous failure. This situation of unidentified replacement makes the usual nonparametric procedures such as the Kaplan-Meier product limit estimator [4] inappropriate.

The primary reasons that replacement components are not identified as to serial number or machine position are simplified record keeping and improved accuracy. Serial numbers are in fact placed on certain card component types, but to require that such information be gathered would increase the time required in completing certain forms, the resistance of those involved in the task, and computer storage requirements for the final files. If the analysis is to be dependent on such serial numbers, one is then exposed to possible recording errors as well as to the age old missing data problem. In related studies where such information has been expected the percentage of missing data has been substantial. For similar reasons machine position of replaced components is not recorded. In fact, components are sometimes switched by the engineer in diagnostic checking so that the recording of position for a component may not be unique over time.

The replacement data is transferred along with system installation data to the Burlington, Vermont, Field Evaluation Department of IBM's General Technology Assurance Function. Since the systems have different installation dates, this situation results at the time of analysis in data that is progressively censored. An example of an actual censoring pattern is shown in Figure 1. This figure shows the percentage of installed components in operations at least t hours. The operating hours are called power on hours and are determined from the installation date of the system and a specific reference date for analysis.

FIGURE 1. PERCENT OF INSTALLED COMPONENTS OPERATING TO TIME t

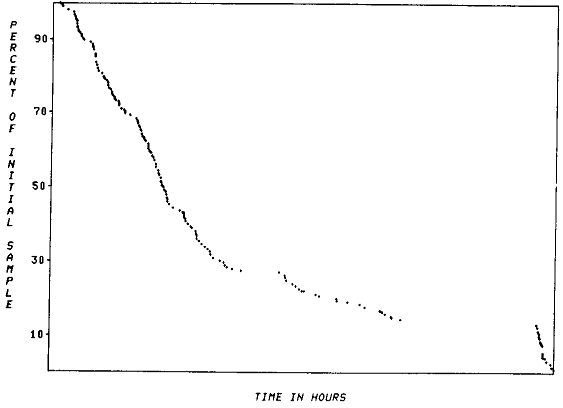

It is the mission of the Field Evaluation Department to track the performance of representative systems maintained by IBM customer engineers. The data received is placed on files which are accessed from terminals.

This department has followed about thirty different systems, each having various model types, using a computerized data base for about four years. Some of the older systems have now been tracked for about eight years. Currently the machines being tracked may contain anywhere from one card to hundreds of cards.

The programming language used for data retrieval, manipulation and analysis is APL. The programs, written by David C. Trindade, provide the implementation of the renewal approach for estimating component reliability. APL is employed because of its high power and flexibility in handling data, its interactive capabilities, and its facility in generating any desired output, whether graphical or tabular.

The advantage of maintaining such a tracking function is apparent when the alternatives are considered. Accelerated life tests are commonly used within IBM to estimate time to failure for various product parts and failure modes, particularly for research and development purposes. However, to predict removal rates for cards under actual customer usage from such experiments is difficult for a number of reasons. Those experiments generally utilize extreme temperatures and voltages not usually existing in the field, so that results are extrapolated to actual conditions. Typically the confidence intervals on percentiles of the lifetime distribution can then be very wide, even if the assumed extrapolation function is appropriate.

Also, we have to develop a model for card reliability as a function of part reliability for each of the variety of card types manufactured, as well as for the time of manufacturing because changes occur over time in the manufacturing process. Such projected card reliabilities may not be able to account for differences in machine models using the same type cards, or for the variability introduced by usual customer engineering practice in making removal decisions. Finally even if reliability projections were developed by some modelling process, one would want to have some field tracking done as a validation of those projections; and the establishment of field reliability for current product types enables more accurate projections for the reliability of new products.

MODELLING ASSUMPTIONS

In order to apply the renewal method of estimating component reliability, certain assumptions must be made about the components and the systems in which they operate. We will describe these assumptions and comment on how realistic they are for our application.

We assume that failure of a single component causes failure of the entire system. Thus, the components are considered serially connected. There are certain computers with error correcting code (ECC) for which this assumption is not true; that is the failed component can be bypassed without materially affecting machine operation. However, in the non-ECC machines of interest to us, serial connections of cards is a reasonable supposition.

We further assume that failed components are immediately replaced so that system downtime is negligible. In practice almost all system failures are corrected by the customer engineer in a matter of hours. Since the repair time is small relative to the lifetime of components, it appears safe to ignore the effect of downtime.

We also assume that the replacement components are from the same population as the failed components. Thus, there should be no reliability growth or degradation. In practice, components from one manufacturing lot can differ from those in another lot, but we suppose that there is a good mixing of components in space and time. Also lot sizes are relatively small compared to the population sizes considered; thus clustering effects should not be large. Of course, for components manufactured over a long period of time there may be improvements in the manufacturing process contributing to reliability growth. This fact is handled in the analysis by separating components into different manufacturing vintages to check for possible differences.

That components fail independently of one another is another assumption. This supposition seems to be reasonable in theory and practice. Occasionally more than one component is replaced simultaneously by a customer engineer, but this case occurs usually for reasons other than component failure. Such situations are ferreted out in the analysis and then proper adjustments are made. The frequency of such instances is small and has little effect on the data base.

It is not necessary to assume that the number of components per system is identical. A particular component may appear with greater frequency in one machine than in another. In addition, the same component may be used in more than one type of machine. However, in the data analysis, machine types are separated so that component reliability is always measured with respect to the particular system designations. Thus, it is possible to compare components in different applications.

We do not have to assume any particular time censoring pattern on systems. Since some systems have not been operating as long as others, some components will be censored or "lost" after certain running times are exceeded. The method we use allows for this progressive multi-censoring in estimating the component renewal function necessary for estimating reliability.

Component lifetimes are measured in terms of estimated running hours. However, in practice this estimate is obtained by multiplying the number of days from installation to the analysis reference date by an appropriate number, often determined from field surveys, of estimated hours usage per day, specific to a system. Thus the component response actually measured is proportional to the number of days in the machine rather than to the number of hours in a power on condition. Clearly there is some source of error introduced at this point in measuring actual lifetimes, but next to tracking actual usage hours for each system, this procedure is the best method available.

A sample of the components removed by customer engineers are failure analyzed to discover the frequency of occurrence of various failure modes. A number of these components will be then classified as NDF (No Defect Fund) because of normal operation at the time of testing. This situation can arise from either an intermittent failure mode or from a component that is truly operational. Intermittent failures can occur in practice wherein the customer engineer traces an error to a particular component, but the problem causing the error does not always manifest itself. In addition, a card which is performing properly may be occasionally pulled by a customer engineer in a problem situation which is difficult to diagnose. The hope may be that the card removed is the cause of the customer's problem although, in fact, the errors may not be due to that card directly.

Because of the counting of NDF's and other failures which may be caused by problems external to the card, (electrical overload, for example) the true reliability of the components themselves is somewhat greater than that estimated from removal data. However, results from failure analysis are employed in assessing the degree of the NDF and miscellaneous contributions to the field failure level.

SUMMARY OUTLINE OF RENEWAL ESTIMATION METHOD

The notation used will first be briefly defined, and then the method of estimating component reliability will be outlined. Let F(t) be the cumulative probability distribution for lifetime t of a particular component, with F(0) = 0. The reliability of a component is then R(t) = 1 - F(t). As failed components are immediately replaced with identical new components, we can view the set of lifetimes for components in a particular location of the machine as an ordinary renewal process. Let M(t), the component renewal function, be the expected number of components replaced through time t.

It is a well known result that M(t) and F(t) are connected via the fundamental renewal equation [5]

$$M(t) = F(t) + \int_0^t M(t-x)\, dF(X),$$

or alternatively,

$$F(t) = M(t) + \int_0^t M(x)\, dF(t-x).$$

Historically this equation has been used to determine M(t) for a particular F(t), say in setting warranties, but we shall use it to estimate F(t) by replacing M(t) with an estimated renewal function $\hat{M}(t)$.

Methods of estimating the component renewal function are fairly straightforward even with multi-censored data and different numbers of cards per system and are discussed by Trindade and Haugh [3].

The integral equation can be solved recursively by a variety of numerical deconvolution techniques which are also examined by Trindade and Haugh [3].

Research is continuing on ways to further improve the numerical deconvolution techniques in terms of estimation bias and variance. We have also reported on ways to establish approximate confidence intervals on reliability estimates [6].

ILLUSTRATION OF APPLICATION

To illustrate the data analysis programs, results for a single component type are presented. However, because of the confidentiality of the data, lifetime and hazard rate scales are necessarily deleted in all displayed plots. The plots should still be adequate to elucidate the basic points.

FIGURE 2. ESTIMATED CUMULATIVE DISTRIBUTION FUNCTION F(t)

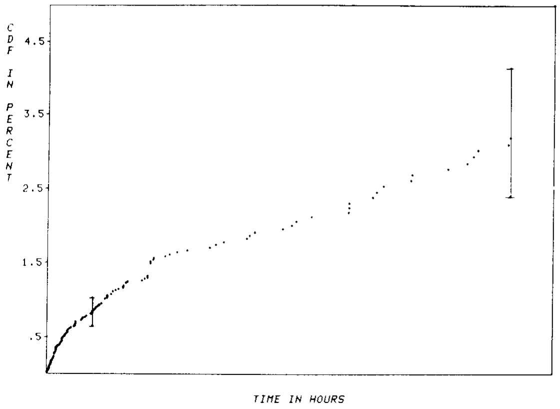

Figure 2 is a plot of the cumulative distribution (CDF) of card removals versus time, based on 137 removals made in 173 machines having a total of 8640 card positions. The CDF has been estimated by the renewal approach previously described. We have indicated the 95% confidence limits at two points in the plot. In actual practice, such a curve is compared to a product specification to establish if the components are performing above, at, or below target. One notes the small CDF values involved. In our work with field reliability, it is common to deal with percentages only to the lowest quartile, decile, or even less, depending on the system studied.

FIGURE 3. ESTIMATED HAZARD FUNCTION (INSTANTANEOUS FAILURE RATE)

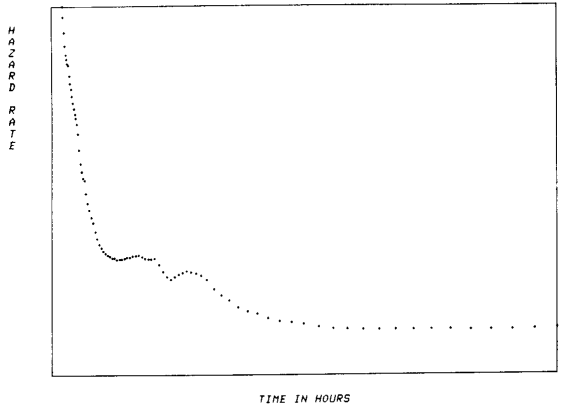

Figure 3 illustrates the plotting of the hazard rate - or instantaneous failure rate - versus time. This plot is determined by a smoothed numerical differentiation of the CDF shown in Figure 2. It is interesting to note that the behavior of this component confirms the pattern expected by reliability engineers for failure rates. We see a high initial failure rate decreasing with time until a plateau or nearly constant failure rate is achieved.

Figures 4 and 5 are lognormal and Weibull probability plots respectively of the distribution shown in Figure 2. Such plots would be used to roughly check the suitability of a particular distribution for extrapolation of reliability to future hours. Here, for example, the lognormal distribution appears

FIGURE 4. LOGNORMAL PROBABILITY PLOT

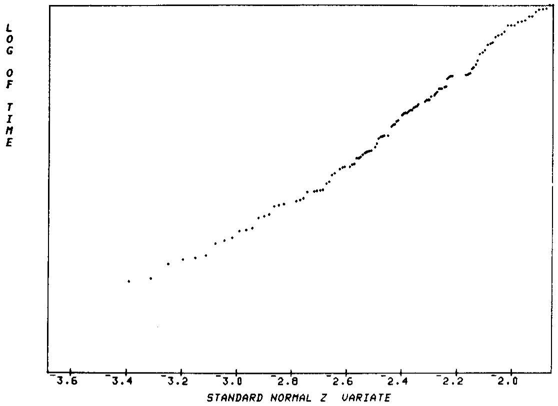

FIGURE 5. WEIBULL PROBABILITY PLOT

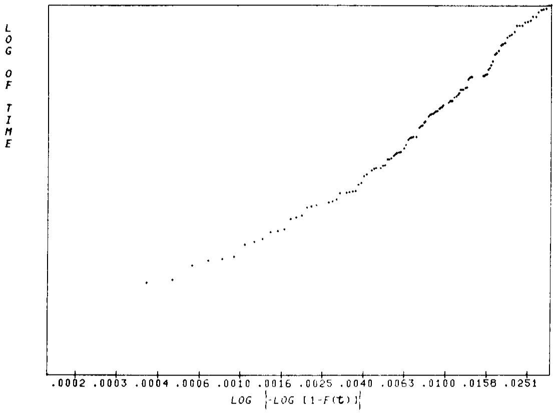

to be more nearly linear and, in fact, gives a better fit - in a least squares sense - to the data than the Weibull distribution. Indeed, it has been our experience that generally the lognormal distribution provides a better choice than the Weibull for fitting and, consequently, would be more suitable for extrapolation purposes.

Other plots, such as histograms of removal times, removal percentages per systems, etc., are also made to permit further analysis and data checking by the engineer. The APL programs permit dropping specific data points to eliminate outliers from calculation.

A report [3] describing in detail the renewal approach is available from the authors. In addition, a comparison of the renewal approach to the Kaplan-Meier method when identification of inter-renewal times is possible (as occurs in the one card system) also is obtainable [7] from the authors.

REFERENCES

1. See, for example, E. F. Platz, "Solid Logic Technology Computer Circuits - Billion Hour Reliability Data," Microelectronics and Reliability," 8, 55-59 (1969).
2. D. R. Cox, Renewal Theory, John Wiley & Sons, Inc., New York (1962).
3. D. C. Trindade and L. D. Haugh, "Nonparametric Estimation of a Lifetime Distribution Via the Renewal Function," IBM Burlington Technical Report, TR 19.0463, (Revised August 1979) and presented at Annual Meeting of American Statistical Association, Chicago (1977).
4. E. Kaplan and P. Meier, "Nonparametric Estimation from Incomplete Observations," Journal of the American Statistical Association, 53, 457-481 (1958).
5. N. R. Mann, R. E. Schafer and N. D. Singpurwalla, Methods for the Statistical Analysis of Reliability and Life Data, John Wiley & Sons, Inc., New York (1974).
6. D. C. Trindade and L. D. Haugh, "Standard Error of Non-parametric Estimators of Lifetime Distributions from Renewal Data," presented at Annual Meeting of the American Statistical Association, San Diego (1978).
7. D. C. Trindade and L. D. Haugh, "Nonparametric Estimation of Component Reliability: An Application to Field Renewal Data," presented at Annual Meeting of the American Statistical Association, Washington, D. C. (1979).