Analysis of the Field Reliability of Repairable Systems

David Trindade, Ph.D.
Distinguished Engineer
Sun Microsystems, Inc.

INFORMS International July 2007 Puerto Rico

Measuring the Field Reliability of Repairable Systems

Objectives

Highlight **limitations** of *MTBF* as a reliability measure for repairable systems

Show some simple but powerful alternative **methods** called ***time dependent reliability (TDR)*** for analysis and modeling of repairable system data

Illustrate - using actual field examples - how *TDR* techniques provide valuable **insights** into important reliability issues.

Definition of a Repairable System

A system is ***repairable*** if, following a failure at time ***t***, the system is **restorable to satisfactory operation** by any action.

Examples include servers, computers, automobiles, airplanes, locomotives, major appliances, utilities, air-conditioners, and networks. Besides ***failure times***, other measures of interest may involve cost, downtime, resources used, etc.

Most of the reliability literature is concerned with the analysis of non-repairable systems, such as integrated circuit components.

System Ages

System age is the elapsed time starting at installation turn-on. Age measures the ***total running hours*** from time zero. Also called power-on hours (POH) or operating hours.

Carefully distinguish *age* from *times between failures*, which are called the *interarrival times*.

Typical Reliability Measure: *MTBF*

Commonly a summary statistic such as the **average or mean time between failures** (*MTBF*) is used to specify or measure reliability.

MTBF usage implies that the **times between failures** constitute a **renewal process**, that is, **independent** observations coming from a **single exponential distribution**, with a **constant failure rate**, resulting in what is called a *homogeneous Poisson process (HPP).*

In a renewal process, there is **no trend**. After repair, the system is **as good as new.**

Without verification of these assumptions, such a measure as *MTBF* can be misleading.

MTBF Estimation

For a **single system** or **group of systems** operating over some time period, *MTBF* is estimated by

$$MTBF = \frac{Total\ Operating\ Hours\ on\ Systems}{Total\ Number\ of\ Failures}$$

MTBF is simple to calculate for any period, say a month, because actual system *ages* are not considered.

This estimate assumes **all systems, system hours**, and **all failures are equivalent.** There is no distinction among system ages, early life, constant, or wearout modes of failures during the time period of interest.

MTBF is a Summary Statistic – Caution: Hides Information

Consider three systems operating for 3000 hours, each with *MTBF* of 1000 hours.

System 1 had three failures at 30, 70, 120 hours and no further failures.

System 2 had three failures at 720, 1580, and 2550 hours.

System 3 had three failures at 2780, 2850, and 2920 hours.

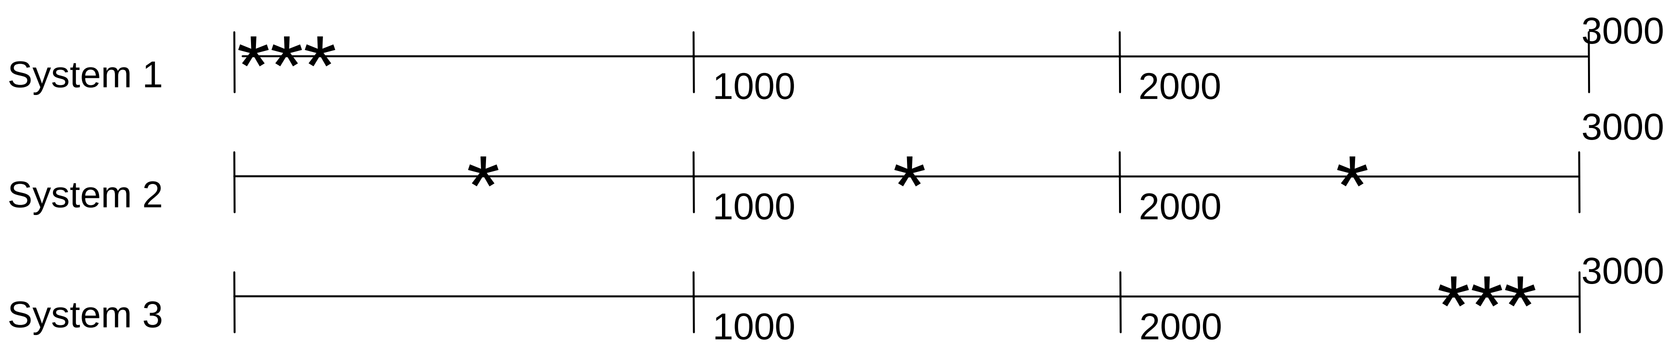

Same *MTBF* but totally different behavior.

MTBF – Dangers of Extrapolation and Misinterpretation

Expected or Typical Lifetime of a System

During the years 1996-1998, the average annual death rate in the US for children ages 5-14 was 20.8 per 100,000 resident population.

The average failure rate is thus 0.02%/yr

The *MTBF* is 4,800 years!

MTBF - Extrapolation in Qualification Activity

Manufacturers often provide *MTBF* estimates obtained by stressing ***many units*** for ***short periods*** of qualification times.

Bending paper clips example.

HPP Model Implications

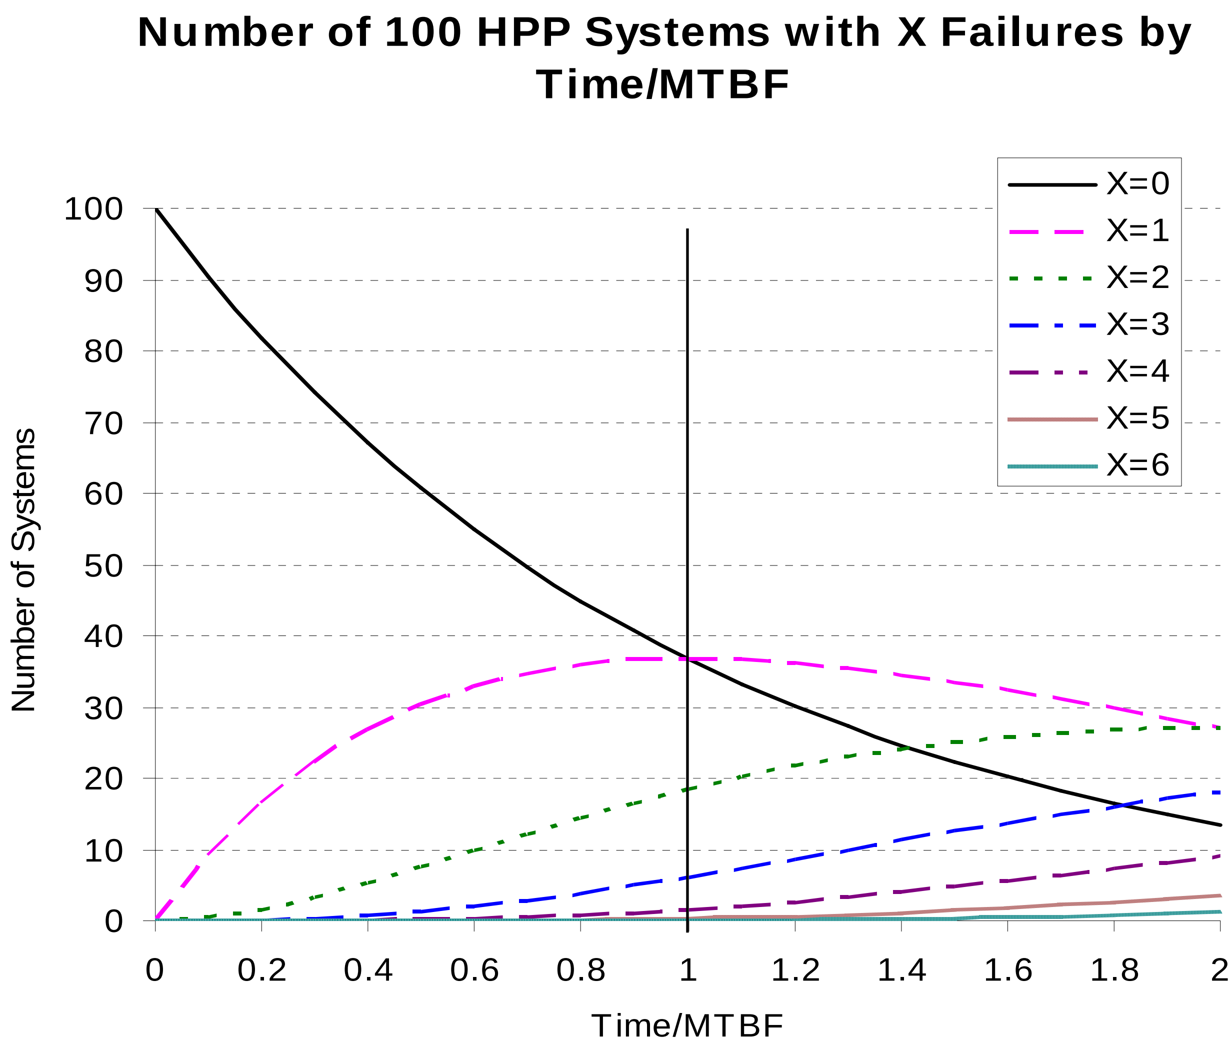

For 100 systems at Time/*MTBF* = 1, average is one failure per system but actually

X = 0 (37 systems)

X = 1 (37 systems)

X = 2 (18 systems)

X = 3 (6 systems)

X = 4 (2 systems)

What's the Point?

Using a summary statistic like the *MTBF* can be misleading and potentially ***risky*** for comprehending field behavior if we do not distinguish between stable or trending processes.

We need to analyze the **ordered** times between failures versus the system *age* to determine system reliability behavior.

Important Property of Repairable Systems

Failures occur sequentially in time.

The *sequence order of interarrival times* provides information that is very important for correct analysis.

Sequence of Times between Failures Provides Valuable Information!

If the times between successive failures are getting **longer**, then the system reliability is **improving**.

Conversely, if the times between failures are becoming **shorter**, the reliability of the system is **degrading**.

If the times show no trend (relatively **stable**), the system is neither improving or degrading, a characteristic of a *renewal process*.

Example of a Renewal Process: Simple Replacement

- Single component system: light bulb.
- Light bulb is replaced upon failure with a light bulb from the same population as the one replaced.
- Stock of spare parts all basically identical.
- Stable environment and use.
- Single distribution of failure times
  - Independent
  - Identically distributed

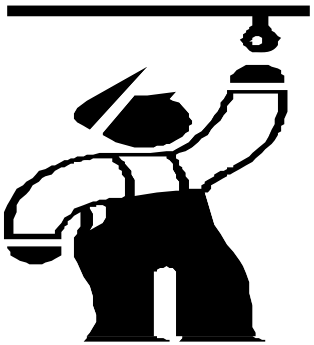

Analysis of Renewal Process

Consider a **single** system for which the **times to make repairs are negligible** compared to the failure times.

Ten failures are reported at the system ages (in hours):

106, 132, 289, 309, 352, 407, 523, 544, 611, 660.

The occurrence of repairs is

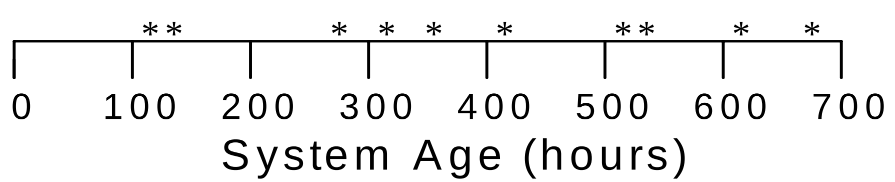

There is no obvious pattern.

Cumulative Plot for Single System

A very revealing and useful graph is the **cumulative plot**: the **cumulative number** of repairs, *N(t)*, is plotted against the system **age**, ***t***, at repair.

For the renewal data, the cumulative plot is:

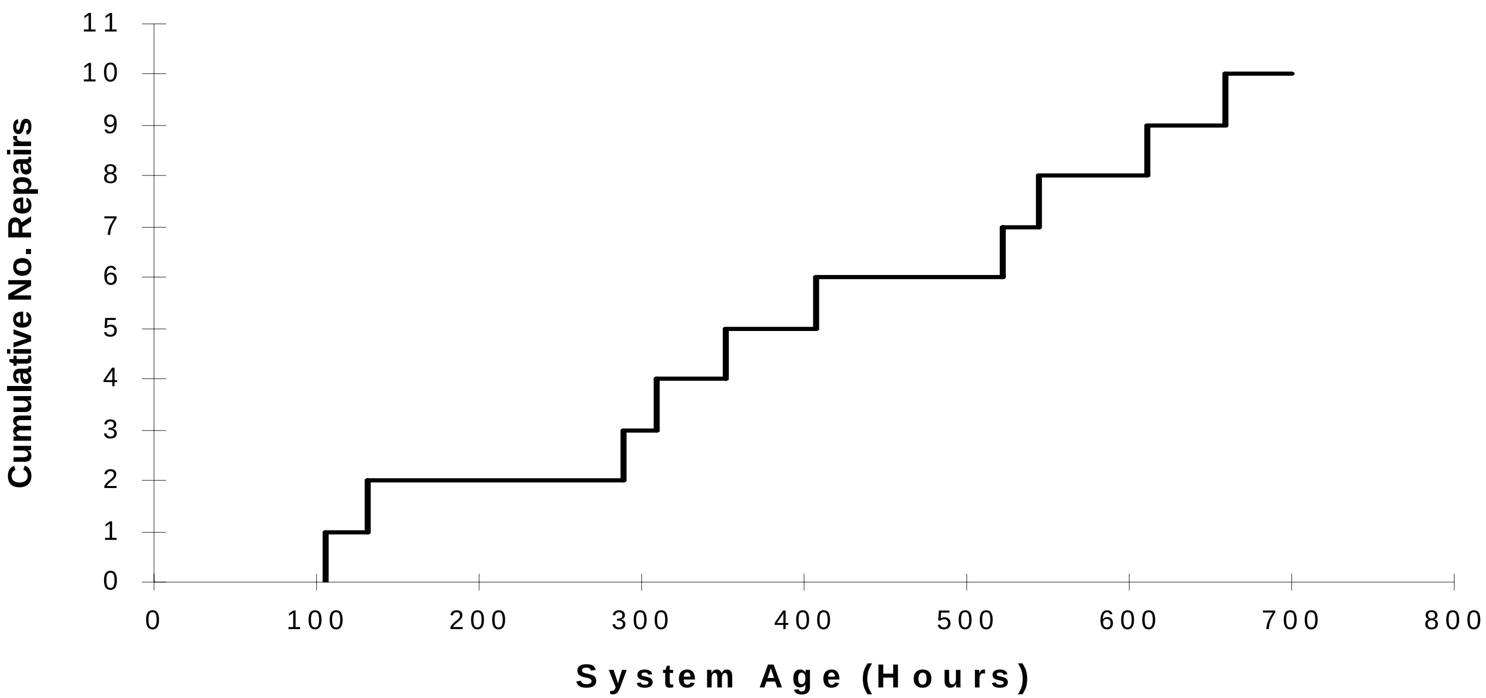

N(t) Follows a Straight Line

For a renewal process, the times between failures are i.i.d., that is, from a single population having a ***constant*** mean time between failures (MTBF).

Consequently, the cumulative plot of *N(t)* versus *t* should appear to follow a **straight line.**

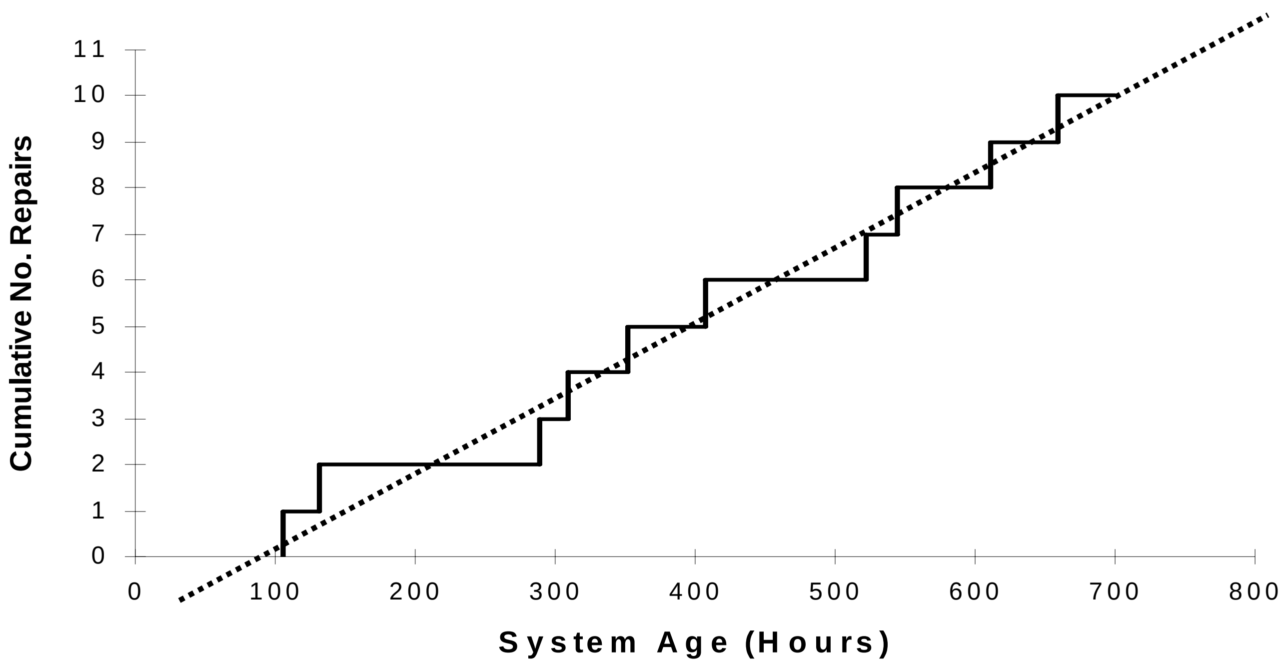

Example of a Non-Renewal Process

Consider a light bulb which is replaced upon failure but the cooling fan inside the equipment is degrading, causing a gradually rising temperature.

Times of replacement bulb failures are getting shorter.

There is not a single distribution of independent failure times (no constant *MTBF*).

To analyze system behavior properly, we must look at the occurrence order of failures .

Graphical Analysis of Non-Renewal Processes

Suppose the repairs occurred at the following times
157, 273, 379, 446, 501, 550, 593, 619, 640, 660.
The line sketch is

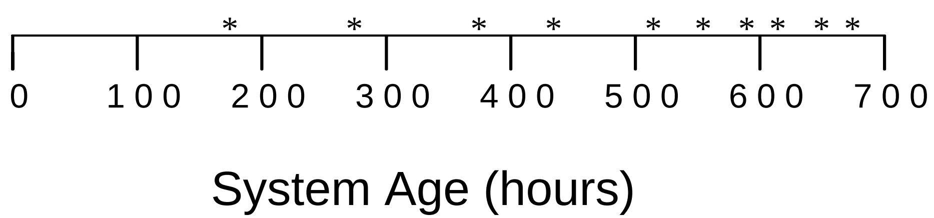

Do you see a pattern?

Cumulative Plot

The cumulative plot is shown below.

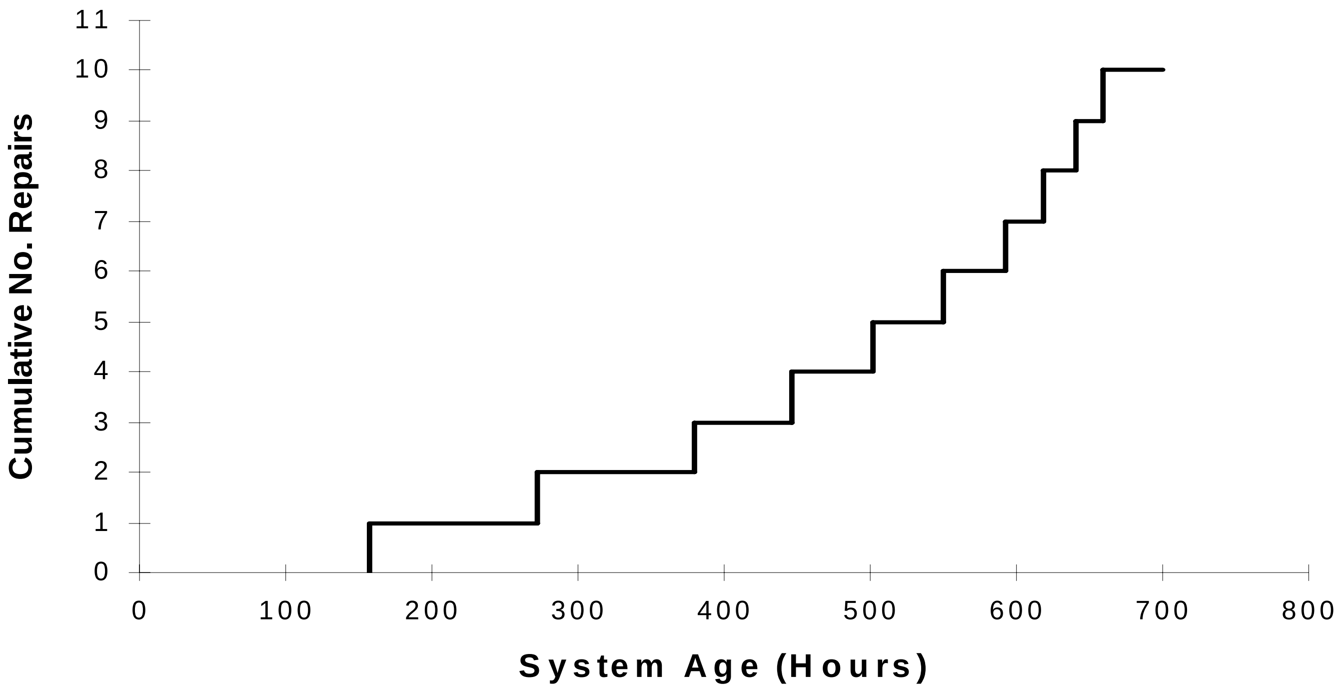

The curvature shows the frequency of repairs increasing in time, indicating system **degradation.**

Another Repairable System History

Suppose the observed consecutive repairs times were
20, 41, 67, 110, 159, 214, 281, 397, 503, 660.

A line sketch of the pattern of repairs shows:

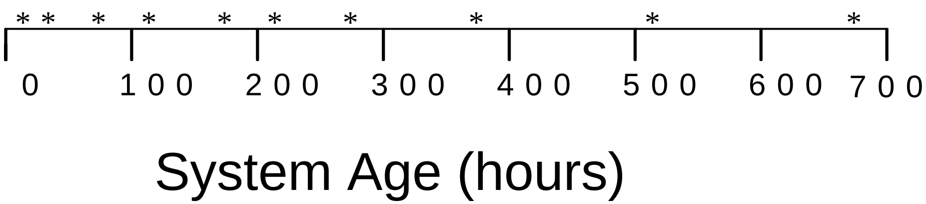

Do you see a pattern?

Cumulative Plot

The cumulative plot for this set of data is shown below.

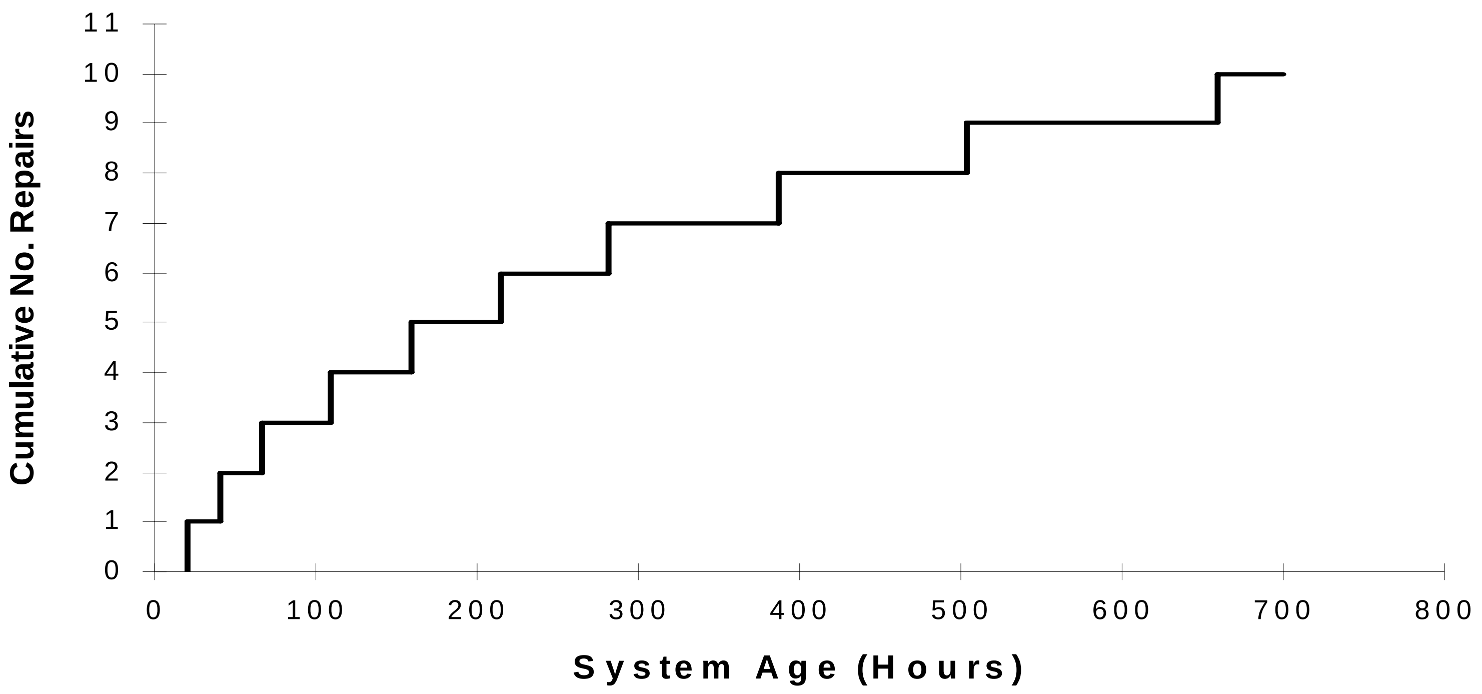

The curvature suggests a decreasing frequency of repairs, that is, an **improving** recurrence rate.

MTBF Comparisons (Interarrival Times)

Stable Renewal Process

106, 26, 157, 20, 43, 55, 116, 21, 67, 49

MTBF = 66

Degrading Process

157, 116, 106, 67, 55, 49, 43, 26, 21, 20

MTBF = 66

Improving Process

20, 21, 26, 43, 49, 55, 67, 106, 116, 157

MTBF = 66

The data are the same! Only the *order* has changed. *MTBF*s are identical! Yet, the behavior is vastly different!

MTBF Comparisons

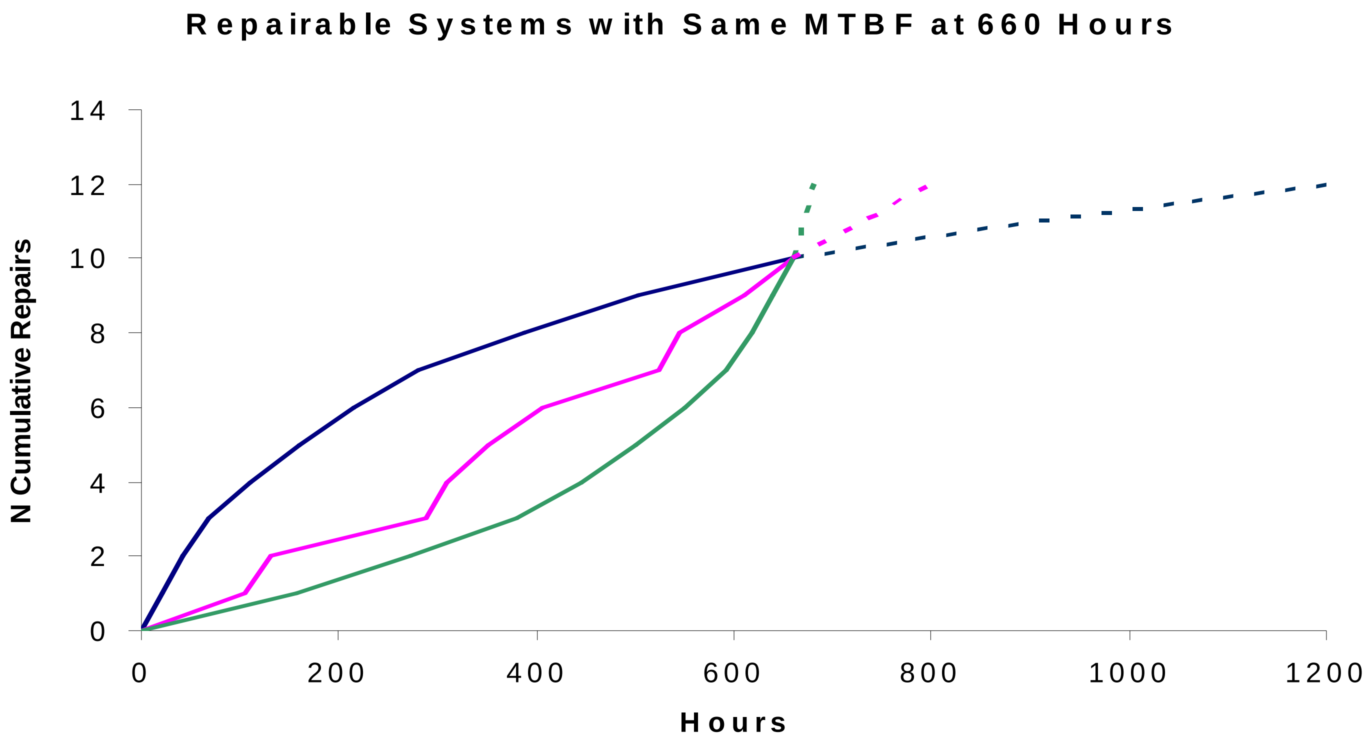

MTBF may tell us on the average where we are at some time, but *MTBF* may not reveal how we got there or where we're headed.

MTBF can be an Inadequate Measure of System Reliability

Valid only for a renewal process.

We cannot ignore potential and likely real age effects.

We need to check for validity of *HPP*.

A better and more revealing approach, with fewer assumptions, is to *analyze the data versus system age*, that is, apply *time dependent reliability (TDR) analysis*.

Analysis of a Group of Repairable Systems

Often we want to analyze the reliability behavior of **many identical or similar systems.** Because the systems are most likely installed on different dates, the system ages will vary, resulting in what we call *multicensored data.*

Example: servers installed in a datacenter at different dates throughout the year will have various field ages.

Graphical Approach to Multi-System Analysis

Consider a group of five systems, installed on different dates. Individual repair histories *N(t)* are shown as steps at each repair. All starting times are referenced to zero.

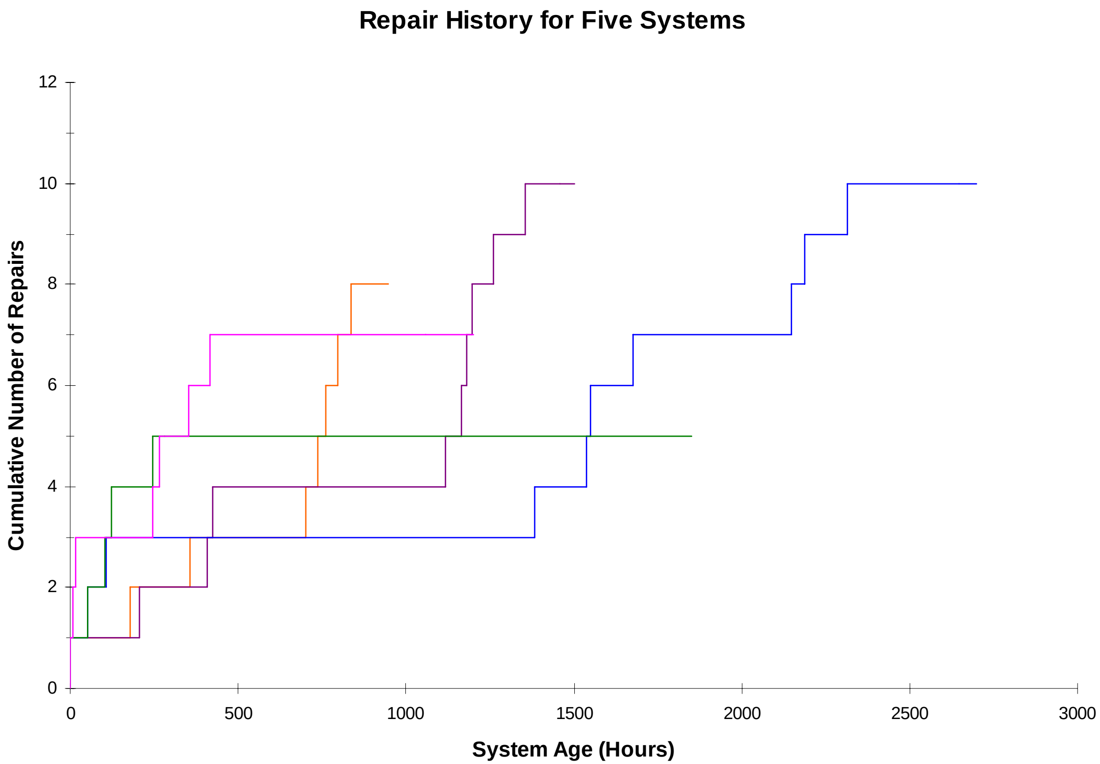

Mean Cumulative Function: *MCF*

We envision a **single** curve denoted by *M(t)* that gives the **average or mean number of repairs per system** at time ***t***. Consider a vertical slice across the individual histories. Such a curve is called the *mean cumulative function* or *MCF*.

Estimation of *M(t)* must account for the **number of systems operational (at risk) at any system age.** Simple statistical procedures can easily handle multicensored data.

MCF for Five Systems

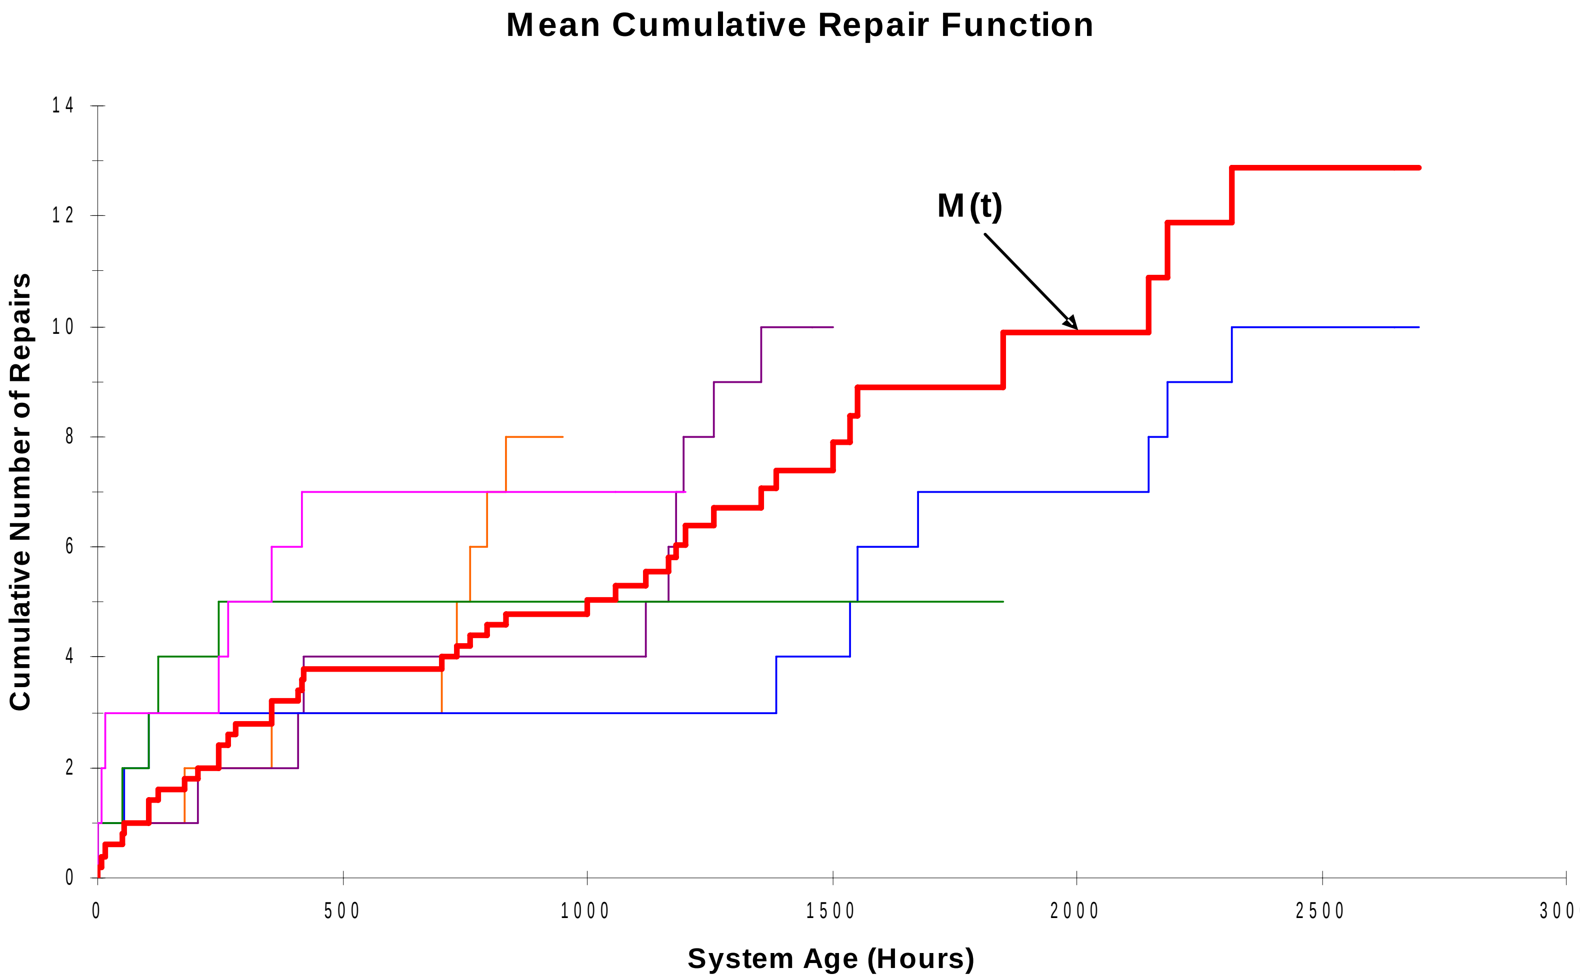

TDR Analysis – Case Studies

Customers want more than an *MTBF* measure of reliability.

What is the reliability of the systems? What should it be?

What are the causes of downtime?

What can we expect going forward?

Let's see what *TDR* can do to answer these questions.

Case Study Example: E10K Performance at Customer's Sites

Customer has five E10K systems installed in early 2001.

Customer is concerned by failures over last year that have caused downtime and impacted production.

TDR Plots for E10K Systems (with Confidence Limits on *MCF*)

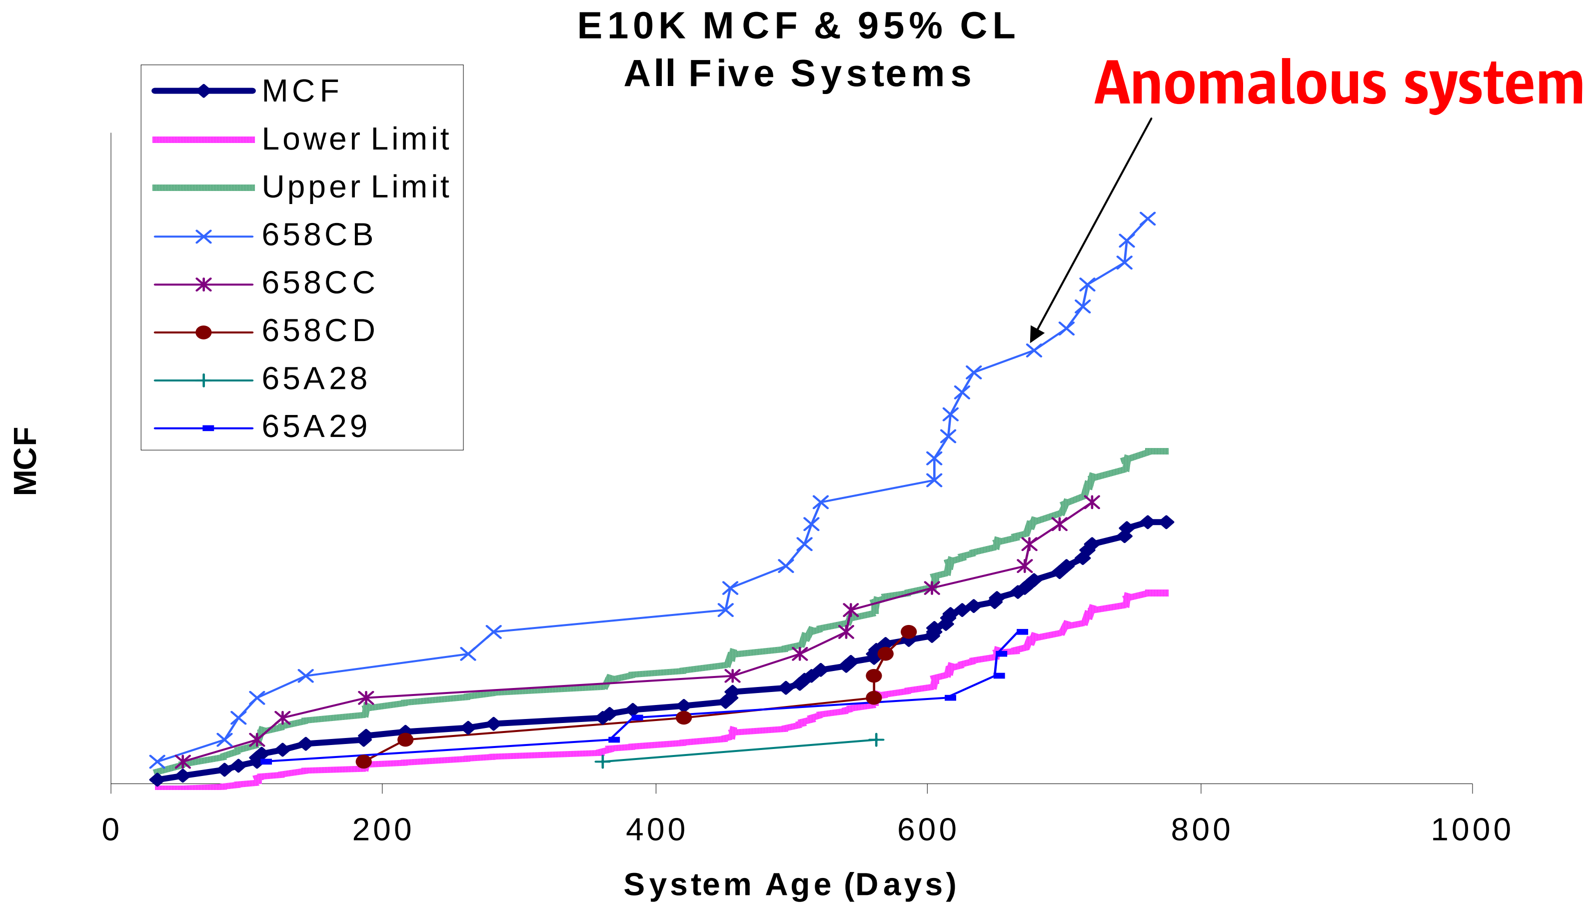

Dynamic Analysis of Anomalous System

E10K 658CB Cumulative Plot with Mode

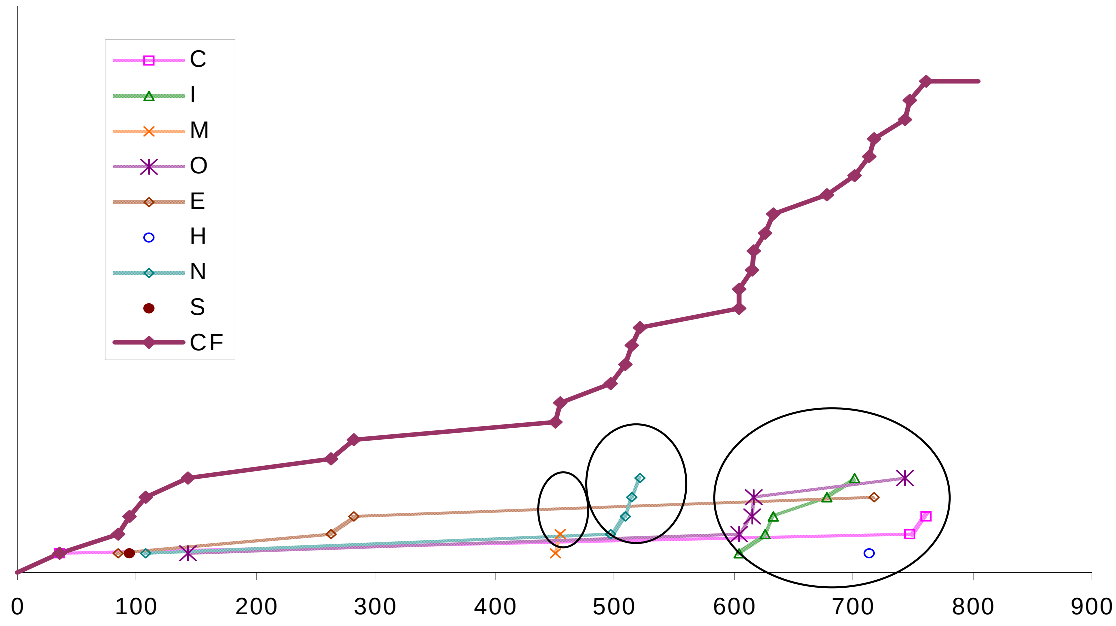

System cumulative plot with cause codes shows the *dynamic* time dependent effects of different failure types on reliability. *TDR* reveals *clustering* of failure causes.

E10K Cause Code Summary

Problems appeared when customer put systems into full production mode.

Multiple failures for same or different causes in a short time period revealed inability of service engineers to diagnose properly and repair correctly the first time.

Failure rates consequently were made artificially higher and *MTBF*s lower by repeated repair attempts for same problem. *TDR* analysis revealed the problem.

Possible *MCF* Comparisons

By platform
By customer
By vintage
By age (left and right censoring)
By calendar date
By location
By failure cause
By supplier
By technology
By payload or applications

TDR for Vintage Analysis *MCF* by Year Installed

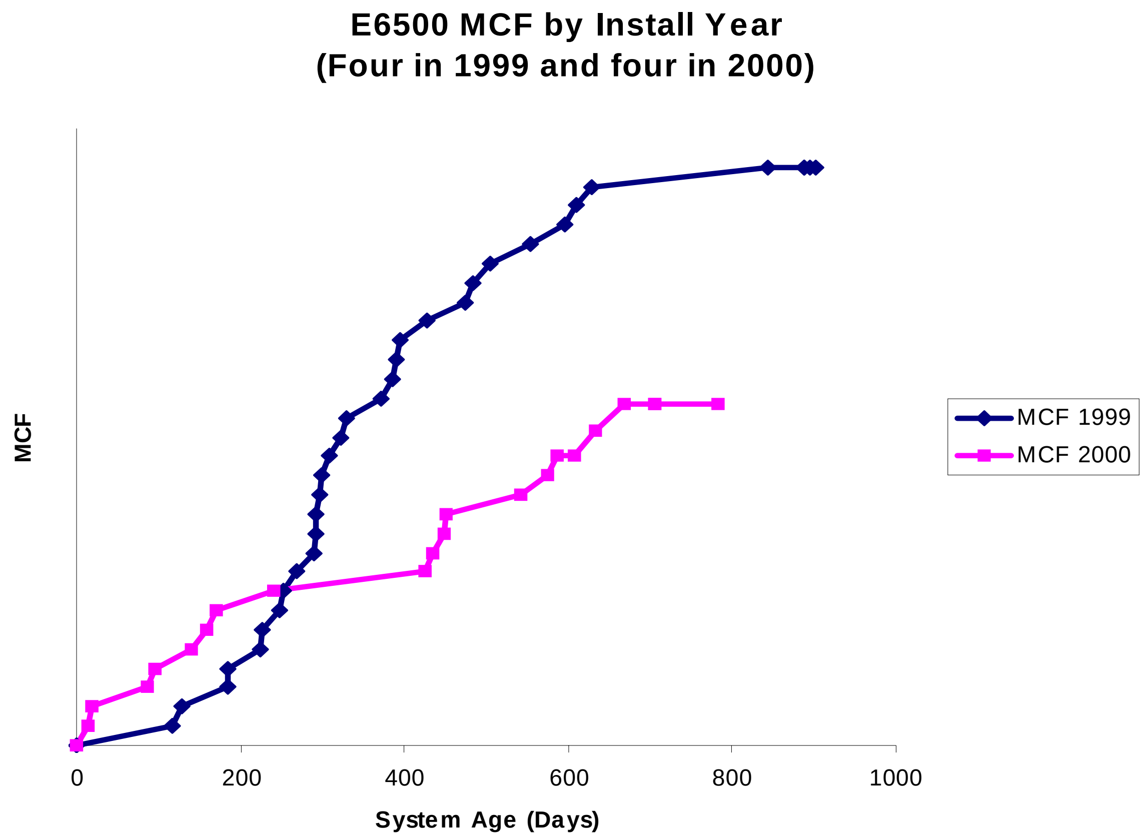

Significant difference between years.

TDR Analysis of Individual Cause Codes

Individual *MCF* plots can be constructed for each cause code to reveal failure trends.

TDR analysis of each cause code provides time related information not available through static Pareto analysis.

Cause Analysis (*MCF* of Top Four Causes by System Age)

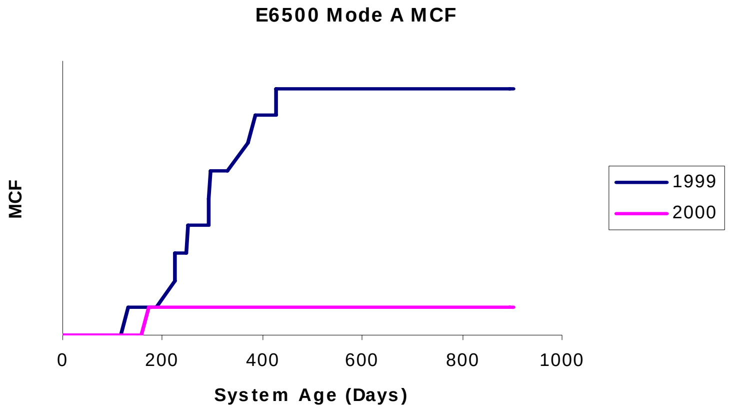

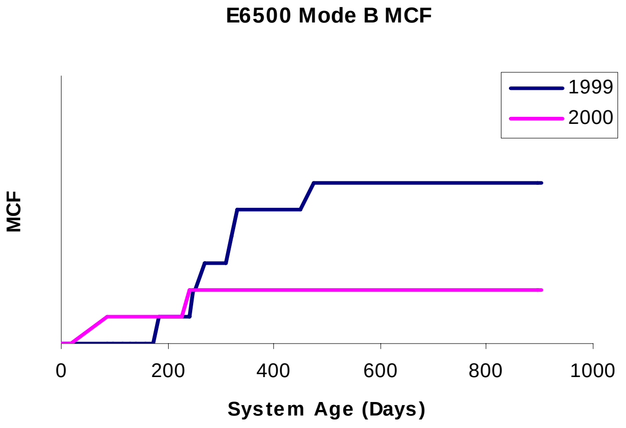

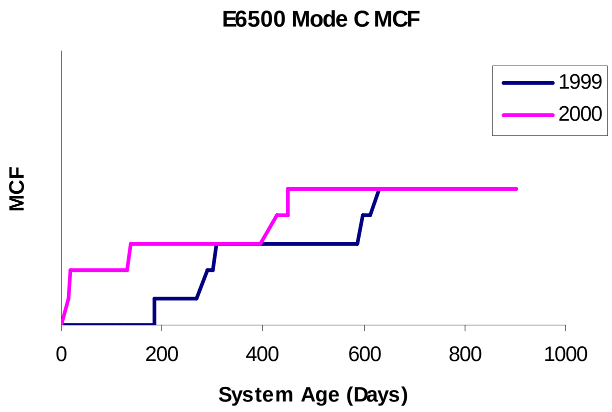

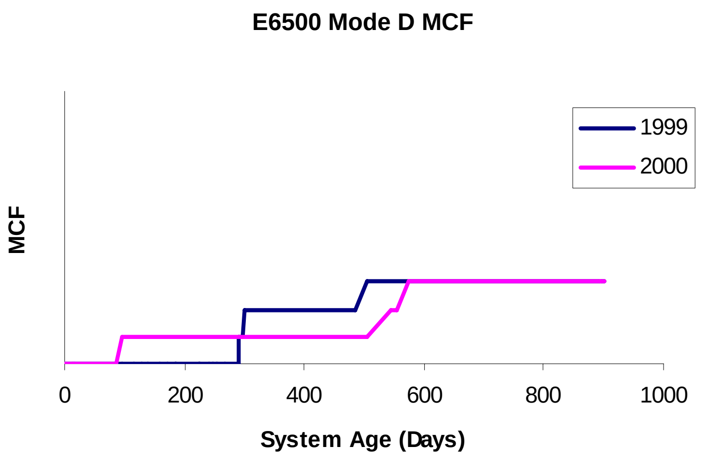

MCF: System Age Versus Calendar Date

MCF versus *system age* shows cumulative repairs per system that depend on the system operating hours.

MCF versus *calendar date* may reveal repairs across systems associated with actions (physical relocation, software patches, upgrades, etc.) during specific time periods.

Both types of plots are useful for cause analysis. Also, both types of curves can be *numerically differentiated* to obtain *recurrence rates* that can reveal interesting changes in reliability.

E10K Calendar Date Analysis

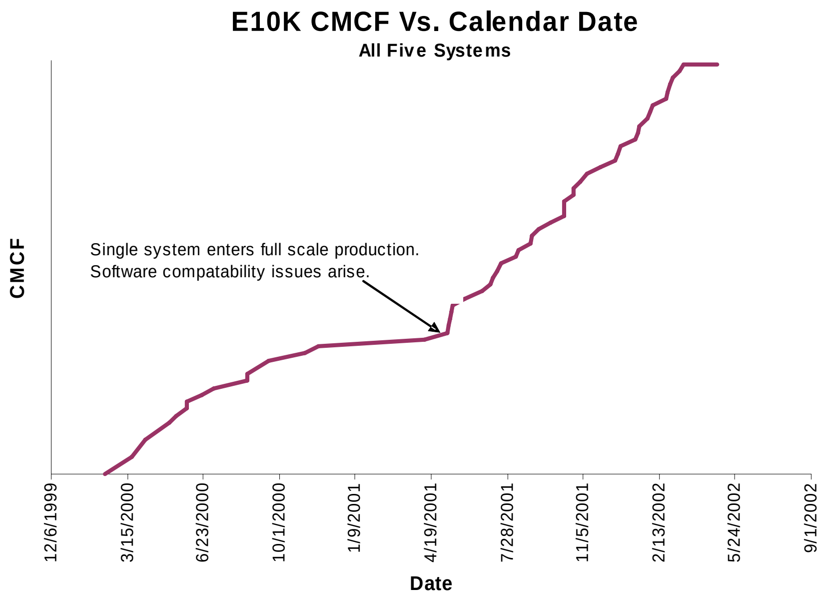

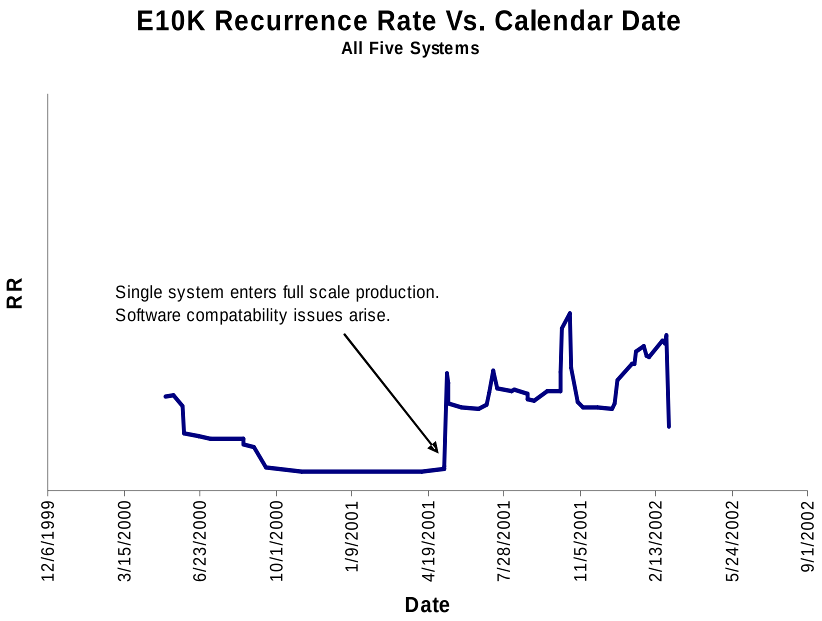

On 5/1/2001, system entered full production mode. Anomalous behavior caused by several software compatibility issues.

TDR Data Needs

For *each and every* system type (*full inventory by serial number*) at a specific customer site:

Date installed (Install)

Date data capture began (*Begin*)

Date of each failure, *if any (Failure)*

Failure cause for each failure

Description of repairs

Current or removal dates (*End*)

Configuration information

Special comments (applications, load, etc.)

Sun Product Quality Portal Example

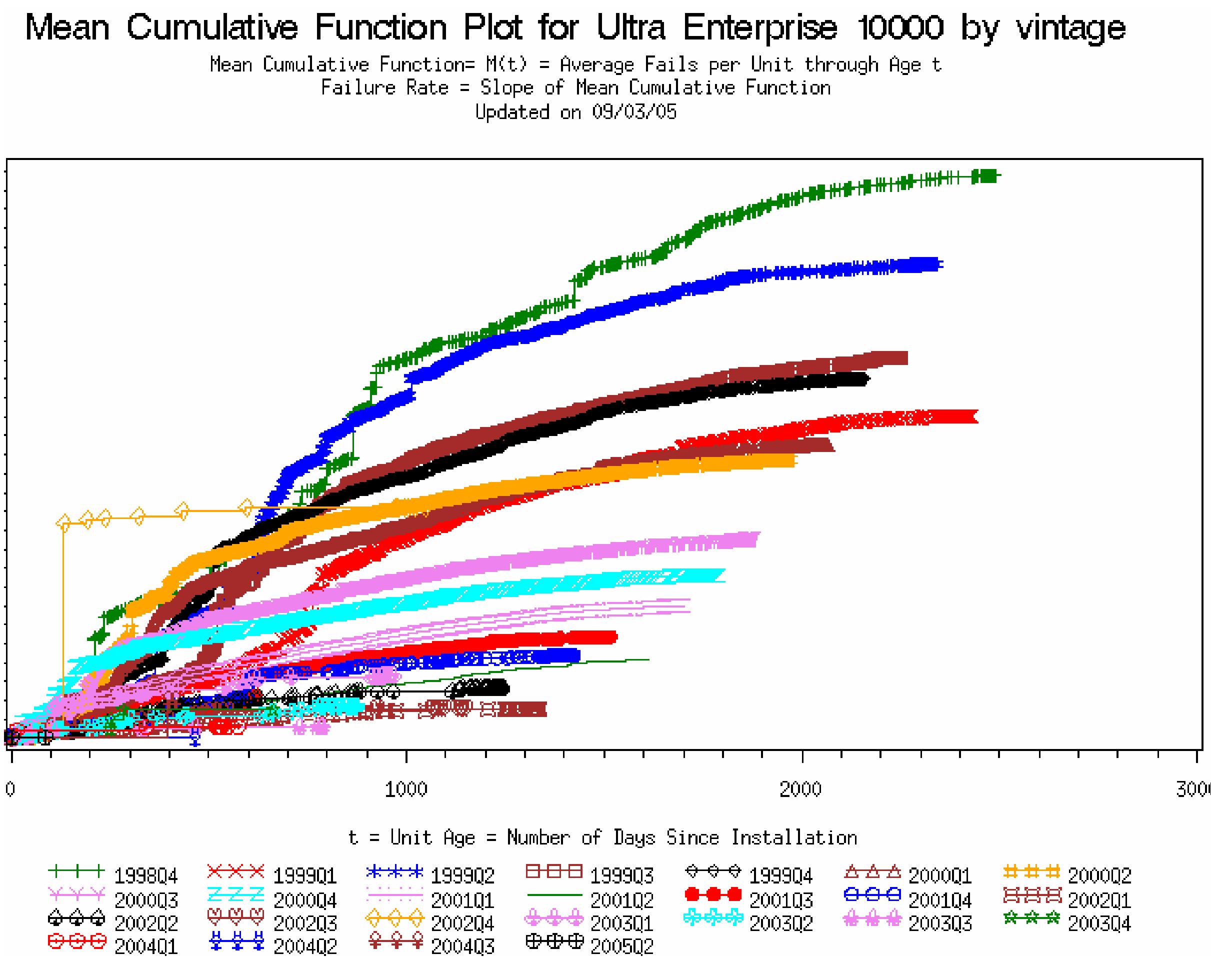

Note reliability improvement across all vintages with age. Longest MCF curves are oldest vintages. Recent vintages have the best reliability.

Summary

The *TDR* approach is a simple and effective method for the analysis of field reliability data. Models (e.g., power law) can be fit to the *MCF* curves to predict future behavior.

MTBF has significant limitations as a reliability measure for repairable systems.

There have been many positive customers experiences involving *TDR* analysis.

TDR analysis has been successfully applied to drive appropriate preventive and corrective actions at Sun with great success.

References

H. Ascher, “A set of number is NOT a data-set”, *IEEE Transactions on Reliability,* Vol 48 No. 2 pp 135-140, June 1999.

J. Usher, “ Case Study : Reliability models and misconceptions”, *Quality Engineering*, 6(2), pp 261-271, 1993

W. Nelson, *Recurrence Events Data Analysis for Product Repairs, Disease Recurrences and Other Applications*, ASA-SIAM Series in Statistics and Applied Probability, 2003.

P.A. Tobias, D. C. Trindade, *Applied Reliability*, 2nd ed., Chapman and Hall/CRC, 1995.

D. C. Trindade, Swami Nathan, Simple Plots for Monitoring the Field Reliability of Repairable Systems, *Proceedings of the Annual Reliability and Maintainability Symposium (RAMS)*, Alexandria, Virginia, 2005.

Analysis of the Field Reliability of Repairable Systems

David C. Trindade, Ph.D.
david.trindade@sun.com

Contributions by Swami Nathan, Bill Chen, Doug Ray, and Jeff Glosup.