Testing for Time-Dependent Failures in Electronic Products

by
Michael L. Peterson

B.S. Electrical Engineering,
Iowa State University (1989)

Submitted to
the Departments of Electrical Engineering
and the Sloan School of Management
in Partial Fulfillment of the Requirements for the Degrees of

Master of Science in Management
and Master of Science (unspecified)

in conjunction with the Leaders for Manufacturing Program

at the
Massachusetts Institute of Technology
May 1994

Signature of Author _____
Department of Management Science
Department of Electrical Engineering
May 1994

Certified by _____
Anantaram Balakrishnan, Thesis Supervisor
Associate Professor of Management

Certified by _____
John G. Kassakian, Thesis Supervisor
Professor of Electrical Engineering

Accepted by _____
Jeffrey A. Barks
Associate Dean, Sloan Master's and Bachelor's Programs

Accepted by _____
Frederick R. Morgenthaler, Chairman
Department Committee on Graduate Students

Testing for Time-Dependent Failures in Electronic Products

by
Michael L. Peterson

B.S. Electrical Engineering,
Iowa State University (1989)

Submitted to the Departments of Electrical Engineering
and the Sloan School of Management
in Partial Fulfillment of the Requirements for the Degrees of

Master of Science in Management
and Master of Science (unspecified)

ABSTRACT

This thesis explores the issue of latent defects, defects which are present or are created during the manufacturing process, but which do not affect the function of a product until later in time. We present the reader with a process for analyzing the presence, impact, and importance of latent defects to final product quality, and we show how testing strategies for precipitating and detecting these types of defects may be selected. Four primary topics are considered:

(1) how to analyze the manufacturing process and use the information which is available at each of its steps to determine the importance of latent defect testing,
(2) the traditional and modern theories for latent defect behavior, and how they can be applied to an actual case study
(3) the available methods for testing for latent defects, the tradeoffs involved in each, and examples and experimentation to illustrate their application, and
(4) the factors to be considered in quantifying the costs and benefits of latent defect testing, the underlying assumptions which must be addressed, and methods for determining the net benefit of testing to the product

Experiences at Hewlett-Packard's Boise Printer Division in Boise, Idaho are used throughout the thesis to provide examples and data to support and illustrate the issues associated with latent defects.

Thesis Advisors

Dr. Anantaram Balakrishnan, Associate Professor of Management
Dr. John G. Kassakian, Professor of Electrical Engineering and Computer Science

The author gratefully acknowledges the support and resources made available to him through the MIT Leaders for Manufacturing Program, a partnership between MIT and major U.S. manufacturing companies.

Acknowledgments

The project and research which were used in creating this thesis were a culmination of input, advice, and encouragement from more people than are possible to mention in a single page here. However, I would like to especially recognize the contributions of several individuals without whom this effort would have been impossible.

Most importantly, I would like to thank the people at Hewlett-Packard for their constant willingness to share information with me, their input of ideas for further investigation, and for the time (and equipment) they freely gave to help out in this effort. Special thanks go to Sam Whitford and Roger Quick for developing the idea for this project and guiding my work, and to the members of the Test Development Group who provided technical as well as moral support during my internship. In addition, their hospitality and friendliness made my stay in Idaho a great time which I will never forget (except for the 15 minutes after my rollerblading crash at the company offsite, which I will unfortunately never be able to remember.)

Thanks also to Anant Balakrishnan and John Kassakian, my thesis advisors, for their support of my work before, during, and after the internship. Their constant input and questions were extremely helpful to me in finding new areas to explore and new questions to ask.

A very special thanks to my fiancee Michelle, who by agreeing to marry me, has already guaranteed that the internship would be successful beyond my wildest dreams. Finally, I am forever indebted to my parents for the dedication, encouragement, and constant support they have provided, which has given me the confidence to take on new challenges in my life which I never would have thought possible. I only hope that I can repay them someday by treating others with the same respect and service that they have shown to me.

Table of Contents

List of Figures and Tables

Definitions and Abbreviations

AFR	Annualized Failure Rate -- a measure of final product quality in the "field", or customer environment
ATS	Automated Test System -- a system of test fixtures and conveyor lines developed at BPR which tests boards without any human intervention
BPR	Boise Printer Division of Hewlett-Packard -- responsible for design and manufacture of small and midsize laser printers
DMD	Disk Memory Division of Hewlett-Packard -- responsible for design and manufacture of large (> 1 Gb) hard drives
ESD	Electrostatic Discharge -- an electric shock (spark) which, when passed through an electronic device, can cause severe damage or product failure
eV	Electron-Volt - unit of energy equal to 1.6×10^{-19} joules
DRAM	Dynamic Random Access Memory -- a component on the formatter boards used to store information as it is processed for use by the laser in creating a printed page
Phase 1	Experimental passive burn-in testing of formatter boards at BPR (power cycling followed by functional test)
Phase 2	Experimental monitored burn-in testing of formatter boards (power cycling with tests of each board during each cycle, followed by functional test)
H-P	Hewlett-Packard Company
IC	Integrated Circuit -- a device consisting of simple functional circuits (such as digital logic gates) connected together to create more complex functions within a single component package
MTBF	Mean Time Between Failures -- a measure of the average time between one failure of a system to the next
MTTF	Mean Time To Fail -- a measure of the average time between the beginning of product use and its *first* failure
NPR	Network Printer Division of Hewlett-Packard -- responsible for design and manufacture of large network laser printers
OEM	Original Equipment Manufacturer -- a company which sells products directly to consumers, and which may purchase subassemblies from other companies for the purpose of assembling them into these final products
PCA	Printed Circuit Assembly -- includes the printed circuit board and it related components (microprocessors, connectors, etc...)
ppm	Parts Per Million -- defined here as the total number of failures of a particular type which would be found in one million production boards, as opposed to one million *opportunities for failure*.
SIMM	Single Inline Memory Module -- a plug-in card for a circuit board such as a formatter board, which provides additional memory or test programs
SPC	Statistical Process Control -- procedures for monitoring the variation in product or process parameters, for the purpose of controlling production variability

Chapter 1 - Introduction and Motivation

1.1 Background on Latent Defects

In the production of electronics products, as with any product, there is the possibility of making mistakes or creating defects. Some of the defects may be instantly recognizable, either through visual inspection, such as a chip which is placed sideways, or through some sort of electrical testing, such as a "dead" component which is blocking the transfer of data. Others do *not* affect functionality right away, but grow over time to become a legitimate problem. These defects are *latent*, or hidden from view or inspection tests. Pynn defines latent defects for a printed circuit board as "assembly, device or operational defects which exist in a board but are not discernible until some time after the board has been manufactured."[1] To address latent defects, we must understand their causes and behaviors over time, determine how to precipitate and detect them, and decide what testing methods are most appropriate for a given situation.

This thesis examines the importance of considering latent defects as part of a testing strategy. To illustrate the application of these ideas, each chapter uses details of the "story" of tests and latent defects for printed circuit assemblies (PCAs) known as "formatter boards" used in LaserJet printers at Hewlett-Packard. These stories are based on interviews, data collection, and test experiments by the author and others at H-P's Boise Printer Division (BPR) in Boise, Idaho. We have disguised the data presented in this thesis for proprietary reasons, but we have maintained relationships within the data in order to illustrate key points.

1.2 Motivation for Production-Level Testing

Product testing at all stages of design and production is important in detecting design flaws, improving reliability, and ensuring that products are safe and in compliance with applicable regulations and specifications. In an ideal world, we could design a product so reliable that no production-level testing would be necessary. However, as today's electronic products test the

limits of product and process technologies, testing during production is helpful in detecting many types of potential problems, including the following general categories:

(1) Changes and limitations in incoming component quality. Vendors often make adjustments to their processes, either intentionally or unintentionally, which may impact product quality. In particular, they may make changes which have no immediate impact on part quality, but which may reduce the robustness of the product so that it fails after a period of time, or after being subjected to later assembly processes. In addition, many of the components used in high-tech products today are only available from a limited number of vendors. As a result, simply "choosing" to buy from only high-quality vendors is often difficult or impossible. Furthermore, achieving high quality is a demanding task in itself. Unacceptable quality for a critical component on a complex circuit board may be below 100 parts per million (0.01%). A seemingly minor process change can thus make the difference between good and bad final products.

In addition, some components may have fundamental limits as to how reliably they can currently be built. Certain product types or process technologies may not yet be capable of achieving very low levels of defects. A primary example in the case of laser printers is crystal oscillators. Repair/replacement rates during the board production process over one two-week period amounted to 11,000 ppm, and similar oscillators on other printer products failed at a rate of about 1,500 ppm during final printer assembly (after passing functional tests at the board production factory). (The nominal goal for quality of the entire formatter board during printer assembly is 2,000 ppm.). Therefore, for a board of about 150 components, this single component is using up well over its share of the allowed defect rate for the entire product. As a result of using hundreds of components in each board, many with stories similar to those described above, manufacturers often find it necessary to consider some sort of testing strategy to avoid producing and shipping defective products.

(2) Uniqueness of product designs and production/assembly processes. The processes used to assemble components into PCAs subject parts to a variety of stresses. Examples of such processes are solder baths, infrared solder ovens, attachment of connectors, and separation of circuit board panels into individual boards. Since many of these stresses may be unique to a particular company, plant, or even assembly line, it is difficult to design components to be able to withstand every possible situation that they may face. Even within a single board, the assembly process will stress different parts in different ways, which makes it even more difficult to predict and design screens to guard against defects at the incoming inspection level. Jensen writes, "On the assembly line, the components will be handled a little differently by the operators, some will be heated more than others during soldering operations, some may be mistreated electrically for a fraction of a second during testing operations, and so on, almost *ad infinitum*."[2] For example, the feeders on different chip placement machines in the formatter factory "bang around" in different ways; some of these put larger amounts of stress on certain components than others. In addition, using

components from different vendors in untested combinations can create compatibility problems such as mismatched timing.

(3) Potential magnitudes of product defects. A strong motivator for production-level testing is the possibility of class failures or defects which are present in a large number of products. If no testing is done at the upstream production stages, it could be months before a problem is detected at a final assembly site or, worse yet, in a customer's product. Repairing this problem can then result in pulling thousands of units of inventory or even initiating a product recall. Although this is certainly not desirable, it's not uncommon for PCA manufacturers to experience a short-term annualized failure rate (AFR) increase of up to tenfold[1], as a result of a single problem like this, which can in turn lead to millions of dollars in repair costs. A testing screen which catches these types of defects can be a beneficial process.

(4) New technologies and process changes. Another use of production testing can be as a safeguard for new technologies, to verify that everything functions as expected prior to shipping a product. Planned changes in manufacturing procedures can also fall into this category. Testing can be an important tool for understanding the implications of the change on product quality or reliability. For example, one of the strong motivators for this project was a change in the procedure used to assemble the PCAs into complete laser printers, as we will discuss in detail in Chapter 2.

All of the above conditions apply not only to current defects, but to latent defects as well. Latent defects can pass undetected through a static continuity test or functionality check, only to cause product failures later in time. Possible causes include intermittent contacts, cracks in traces in the board or inside components, the use of materials with different thermal properties to make connections, and the general effects of aging on components. Problems such as these may not affect the functioning of the product until the final product is used for a period of time. Therefore, they cannot be detected by conventional functional tests. The stress of applying and disconnecting power, repeated handling, or simple product use could cause these latent defects to develop into real failures which affect product quality.

[1]AFR for a specific component of a laser printer, such as a PCA, is defined as the total number of PCAs that were found to be defective in the field in a given month divided by the total number of printers that are still in some part of their one-year warranty period. A jump in AFR from 0.3% to at least 2% is possible when a major class failure occurs.

1.3 The Role of Testing in Competitive Success

Many of the critical factors for success in today's consumer markets can benefit from the use of production-level testing. Described below are some of these factors, along with the impact that an effective testing strategy can have on them.

(1) Short product design and life cycles. Time-to-market is increasingly one of the most important criteria for success for products such as automobiles, printers, disk drives, etc. In order to go from a product concept to full-scale production as quickly as possible, many tasks need to be accomplished concurrently if possible, such as product design, process design, manufacturing equipment specification, training, and so on. A thorough testing procedure can serve as a final gate for assembled products by detecting when changes have occurred. Testing operations also need to be flexible, in order to maintain their effectiveness as new products come and go through the manufacturing environment. The more that testing can do to feed back information on the success of each stage of the design and manufacturing process, the better the quality of the finished product.

(2) High technology requirements. New advances and successes in competitive electronics industries are often driven by the ability to introduce products which are at the "bleeding edge" of technology. The danger in applying new or untried technologies to new products, though, is that the production process or the components used in the product may not be robust enough at the early stages of the technology to guarantee high reliability, hence the bleeding. In these situations, it is also difficult to accurately analyze the causes of problems that occur. As a result, designers may need to experiment with various possibilities to see what factors are responsible for the problems that are being observed. In these situations, having an effective gate that can screen defective products while they are still in the factory is critical to better understanding the technology at hand. It also prevents the manufacturer from passing defects on to customers.

(3) Fast cycle time for manufacturing products. One important factor in successful manufacturing today is a low cycle time. This is especially true for factories producing a high mix of products which are in demand in the marketplace. For example, it is not unreasonable for a PCA manufacturer to be able to build a product with 200 components in under 30 minutes. To achieve a low cycle time over the long term, we need to have less work-in-process inventory, faster and easier changeovers, and a learning-capable system. A learning-capable system is one that can quickly respond to and compensate for changes in processes or problems with incoming components. In this respect, it is important for testing to be quick, complete, and specific in the problems that it finds.

1.4 Philosophy of Testing

Designing for maximum robustness and quality is a necessary part of producing quality products. If we assume that all quality problems can be resolved or eliminated through effective

design and manufacturing processes, then relying on testing would appear to be unnecessary. However, the challenges described above dictate a limit as to how tightly the quality of a product can be controlled solely through design. Proper testing provides a feedback mechanism which can aid in identifying those defects which cannot be eliminated in advance. This is especially important when our process and inputs are not completely understood, controllable, or consistent. Testing also provides a tool which aids in learning about the nature of these defects so that they *can* be eliminated. Rather than thinking of testing as being a process that ideally would never be required, it may be more helpful and accurate to view it as an effective feedback tool that ideally catches every defect which cannot be otherwise prevented.

1.5 Organization of Thesis

When we consider latent defect testing, we need to follow three general steps:

1) Identify the type of flaw we are hoping to detect.
2) Identify a precipitation method which will bring that flaw to the surface.
3) Identify a detection method for finding the resulting defect or failure.[3]

The next three chapters guide us through this process. Chapter 2 provides some background regarding the formatter boards and the laser printers produced by H-P, which will serve as the basis for applying many of the theories and learning in the following chapters. This chapter also reviews the types of data that we can use or analyze to get a better feel for what is happening with respect to latent defects in a product. Chapter 3 reviews current literature and theory on the nature of latent defects and presents various possibilities for how to think about the sources and behavior of latent defects. By comparing the actual data to the theory, we gain an appreciation for the *problems*. In Chapter 4, the types of testing which are effective in screening for latent defects are discussed, along with specific examples from various companies which demonstrate how they have been used in practice. These represent possible *solutions* for latent defects. Finally, Chapter 5 presents the factors involved in making a testing decision based on its projected costs and benefits, and considers the degree to which a tangible, quantitative financial analysis can

be used in making decisions on these types of issues. Fig. 1.1 illustrates the decision process which forms the outline of this thesis.

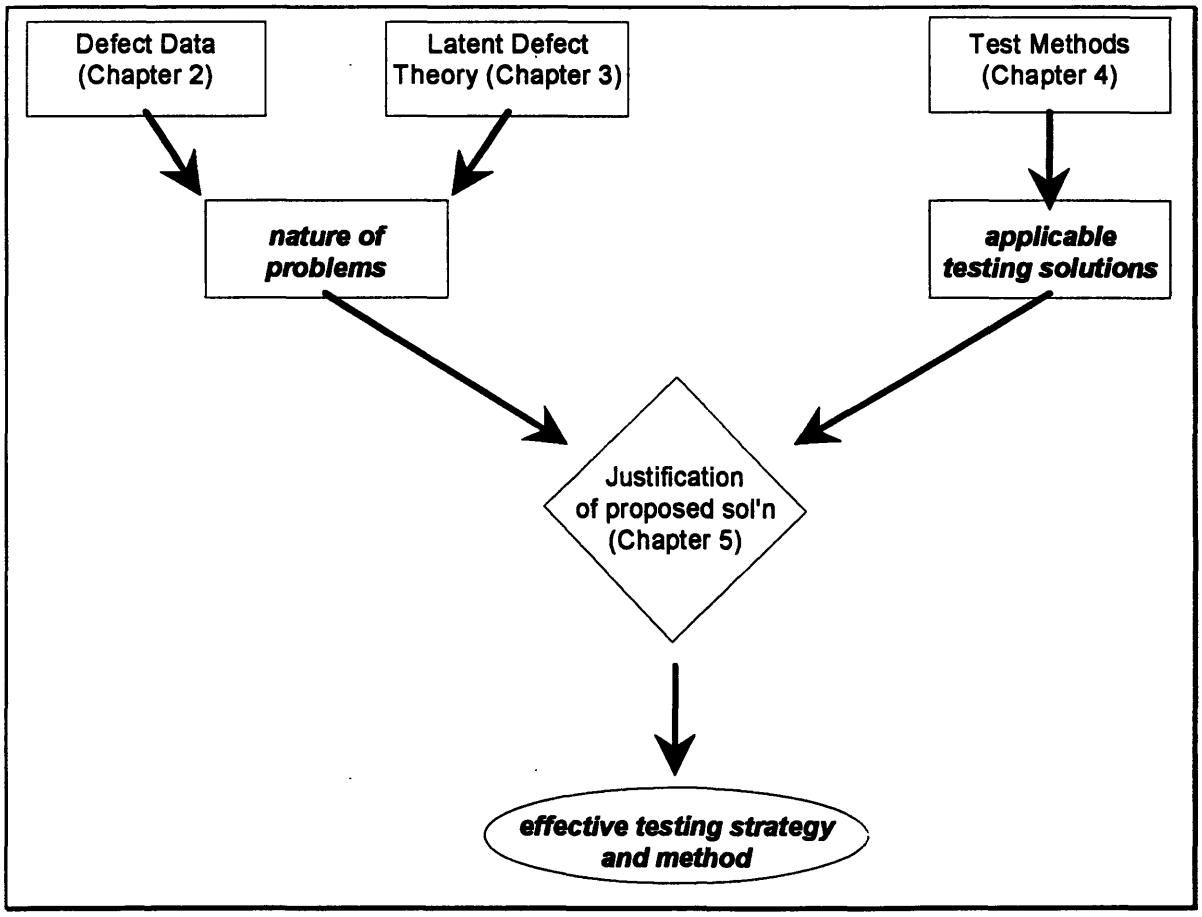

Figure 1.1 Latent defect testing decision process.

References

[1] Pynn, Craig, Strategies for Electronics Test, McGraw-Hill Book Company, 1986, p. 169.

[2] Jensen, Finn, and Niels Petersen, Burn-in: An Engineering Approach to the Design and Analysis of Burn-in Procedures, John Wiley & Sons, Chichester, Englaand, 1985, p. 83.

[3] from presentation by Gregg Hobbs at H-P, Boise, Idaho, August 2, 1993.

Chapter 2 - Sources for Information

2.1 Introduction

When making decisions about issues such as testing, we must analyze the available data in order to understand each stage of the product's assembly and use in as much detail as possible. In addition, we must assemble and compile the data in such a way that we can draw conclusions with some degree of confidence. It is fairly easy with advanced information technology to collect reams of data, but the challenge lies in determining how we use this data to make changes or decisions. The process of printed circuit assembly consists of many steps, each of which can cause or detect a potential failure. Furthermore, the defect may be severe enough to be detected at the next testing stage, or it may only cause intermittent or latent failures which require further stresses before a true product failure results. From product design through to final product use in the field, the general locations where defects may be created or may occur in PCAs include:

- design - errors in design
- design/components - due to untested or unplanned combinations of components from different/new vendors
- components - defective-on-arrival before board assembly
- components - damaged during board assembly process
- board assembly - process-caused defects such as solder shorts or opens, misplaced parts, non-optimal parameters set on process equipment, defects caused by test equipment, etc.
- damage of board during shipping to integration site
- printer assembly - damage of board during final assembly
- damage of board during shipment to dealer/end user
- damage of board during use of final product at customer site

In this chapter, we will discuss each stage of the manufacturing and final use process with respect to the types of failures that can or cannot be detected. We will also consider what types of information may be available or at least collectable. The specific case of laser printer formatter boards will be discussed at each stage to illustrate what types of data are available in reality. In order to distinguish between the possible sources for failure data, we will break the manufacturing process and use of the product into the four stages shown in Fig. 2.1.

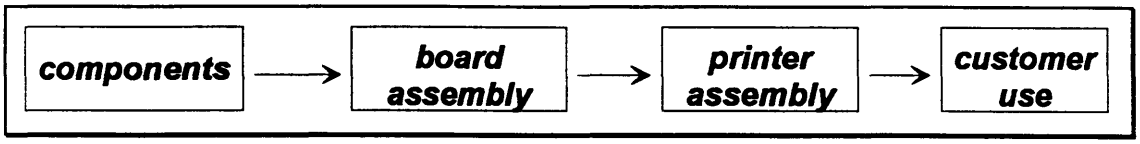

Figure 2.1 Manufacturing process overview for formatter boards.

2.2 Brief History of Laser Printers at Hewlett-Packard

Hewlett-Packard's involvement with laser printers began in 1981 when Richard Hackborn reviewed an invention from Canon in Japan and recognized its potential as a business opportunity for H-P. The invention was a desktop printer powered by a small engine capable of printing with letter quality. The printer operated by using a laser beam to charge a roller, which would in turn attract powder toner particles to it. When the roller was heated, it would melt the toner and allow it to be attracted to the paper via static electricity, creating the final copy. At the time, H-P's closest related product was a $100,000 impact printer suited for IBM mainframes, therefore this represented a major new product for the company.[1] The partnership between Canon and H-P has resulted in the sale of over ten million laser printers over the last ten years; the laser printer divisions of H-P have also become some of the corporation's most profitable operations.[1]

Today, Canon still provides the print engines for all of H-P's laser products. The design of new printers therefore requires close coordination between the two companies. The prime source of added value that H-P provides for the printer is the design and production of the formatter board, which acts as the brain of the printer. This printed circuit board, which varies in size from approximately 4" x 6" to 7" x 10", accepts input from the user's computer via a serial, parallel, or AppleTalk port, and translates the print data into a format which can be understood by the laser in the printer engine. The board itself usually contains two layers of traces, and surface mount as well as through-hole components are attached to one side. Components on the board include the

[1]There are currently two divisions of Hewlett-Packard involved in the production of laser printers: Boise Printer Division (BPR), which deals with small to midsize printers for the personal and small business users, and the Network Printer Division (NPR), which designs and builds network-sized printers, designed for multiple users in office settings.

microprocessor, an application-specific integrated circuit (ASIC) customized for each product, memory chips, data buffers, input/output connectors and other related components. Each board contains anywhere from one hundred to two hundred components, depending on the product. The following sections of this chapter describe the key issues at each stage of assembly and their relationship to the formatter boards used in laser printers.

2.3 Production Process and Available Data

In this section, we discuss the types of information which may be available at each step of the assembly process. In order to ground this discussion in reality, we will study, gather, and analyze data from the real-life example of formatter boards and laser printers from Hewlett-Packard, but we will also review possibilities for data collection and analysis which are not available for this process but which may be available for other electronic products.

2.3.1 Component Level

Companies which assemble products such as formatter boards obtain components from several different vendors. In fact, the same component may even be purchased from more than one supplier. Multiple sourcing of components is done for a variety of reasons. For example, a single supplier may be unable to supply all the parts that are needed, especially for large quantities. Other suppliers can then be available as "back-up" in case of a problem. For more common parts, multiple sourcing may be utilized as a strategy to guarantee a low cost through competition. For whatever reason multiple sourcing is done, it can cause a variety of problems for product quality. Different suppliers of supposedly identical components may use different parameters or equipment in their manufacturing processes, or may even employ entirely different technologies or designs as long as the end product still functions as required. When combined with the hundreds of other parts on a board, the interactions and dependencies among all of the components could make functionality impossible to achieve. In addition, the supplier may make

process or design changes without realizing their impact on final quality or reliability, and the customer may not know about them until products start to fail in the field. For a typical formatter board, the components can come from over thirty different vendors. Test engineers thoroughly test components from specific vendors before using them in production. However, once production begins, components are not individually tested before they are placed on a formatter board.

Component failures are detected in a variety of locations: during board assembly, during final printer assembly, or in customers' products in the field. When a component failure occurs, operators remove the component from the board and component engineers send the component to the original vendor for detailed analysis. Analyzed failures fall into three major categories: (1) no-trouble-found, (2) test failures, or (3) legitimate functional failures. No-trouble-found's (NTFs), which make up the largest category of "failures," are components for which no functional problems can be detected. Possible root causes for NTFs are a misdiagnosis of the actual failure cause when the part was originally determined to be bad, the inability to test all failure possibilities at the vendor, or an intermittent failure mode. A problem is classified as a "test failure" when it is found to be defective when analyzed, and the failure would have not been detected during normal inspection when the component was manufactured. These problems are detected by running two types of tests: (1) the part is placed in a known good formatter board and the regular formatter functional tests are duplicated, and (2) the test vectors that are run on components as they are first shipped for use are repeated. If a component then fails the first test, but passed the second, a "test failure" has occurred. These problems can be eliminated by updating the test coverage for the production parts, but at this point it is often too late to be worth making the change in production, due to the short life of any individual product. For example, for the ASIC used on a recent formatter product, analysis engineers discovered the following failure types in components returned from board assembly:

No Trouble Found	1000 - 1500 parts per million (ppm)
Test failures	400 - 900 ppm

Functional failures 200 - 300 ppm

One critical component which illustrates many of the key issues related to component quality for formatter boards is the crystal oscillator, which is responsible for all timing on the formatter. During the oscillator manufacturing process, the crystals are cut into small strips, and their frequencies are tuned by depositing silver particles on them. Capacitance is also necessary for the crystal to oscillate. Capacitance can be created within the oscillator by designing the capacitor directly into the package, or by including it as a discrete component.[2] Problems occur when loose particles fall onto the crystals after they have been tuned (either during or after manufacture of the oscillator). This can affect the frequency of oscillation to the point where the board cannot function properly. In addition, power cycling can accelerate this contamination, and result in the oscillator dying or getting stuck at the high or low voltage. Oscillators are also extremely sensitive to high temperatures involved in soldering components to the board during board assembly. If the oscillator is not able to handle the stresses of infrared or wave soldering of the entire board, special process steps are required. To assess the reliability of oscillators from new vendors, H-P tests thousands of parts at a time for component validation, looking at cycle time, rise and fall time, and frequency shifting at various stress levels of temperature and voltage. Other causes for failures include terminals within the oscillator being shorted together, delamination of epoxy used to hold the oscillator together, or electrostatic discharges (ESD) which can create an instant failure or can weaken the oscillator and shorten its useful life.

The data from component testing and analysis is most useful for determining relative magnitudes of failure modes and for discovering new failure modes which can perhaps be eliminated through production process changes. Since component engineers test only failed components from the field, it is difficult to determine the field population size for calculating field ppm rates.

---

[2]Oscillators with designed-in capacitance are more sensitive to electrostatic discharges (ESD) and are generally less reliable than those which use a discrete capacitor. In the past, H-P has used the lower reliability parts due to low availability of the more robust version.

2.3.2 Board-Level Process

When components are assembled into complete electronic products, new types of challenges and defects become important in addition to the component-level defects described above. Every possible parameter of the board assembly process can have an effect on reliability of the product, either in terms of immediate functionality or by creating a latent defect. Variations can occur in every step of the process. Potential sources of variation include the following:

- differences in amounts of solder paste
- alignment of parts by chip placement robots
- temperatures and pressures in solder ovens
- stresses from board conveyors

In addition, components operating at the margins of their timing specifications cause malfunctions at the board level if their signals are not precisely coordinated with the components around them. Since there are so many interactions and potential for problems at the board level, board assembly represents a great opportunity to learn about the types of possible problems. To be able to take advantage of this learning opportunity, we must give attention to two key areas: (1) available capacity--time and space--for testing products on the line, and (2) the ability to collect test results and compile them into formats which directly highlight opportunities for improvement. As we design a board-level testing and data collection process, we need to specifically address the following issues:[3]

1) What types of failures could be generated at each step of the test process? What types of tests would detect these?
2) At what stages of board assembly should boards be tested? Should there be multiple test stations, or would a single test at the end of the line be efficient? How much value would be added to the board between the time a failure is created and the time it is detected?
3) To what level of detail should we collect information at each test station?

---

[3]During the study of the formatter boards, our primary focus was on compiling and analyzing the existing data from the production processes, rather than redesigning the data collection process itself. Nevertheless, if the reader *has* the opportunity to redesign the testing procedures, the answers to these questions may lead to ideas for beneficial changes to the current processes.

4) How should the data be presented to stakeholders in the process? How do the information needs of senior management, process engineers, maintenance personnel, line operators, etc. differ?
5) What types of statistical analyses can be used on the data to learn more about the process? It may be a challenge to adapt tools such as statistical process control (SPC) to a product which may have dozens of potential failure modes and a variety of different characteristics which could be measured; finding an "analog" type variable which can be charted to indicate the overall quality of the product may be difficult as well, if the product being produced is digital in nature.

The accuracy of the testing equipment is also an important consideration. It is critical that the tests diagnose good product as good, and bad product as bad, as accurately as possible. This sounds like common sense, but with the increasing sophistication of electronics assembly processes, the boards going down the line are often at quality levels comparable to the ability of the test equipment to diagnose their condition. Therefore, when we consider goals of "six-sigma quality" or "zero defects," we must be sure our test equipment is significantly more reliable than our products if we want to be able to measure how well we are meeting the goals.

Formatter boards for H-P's laser printers are built on one of eleven production lines; seven are located in Boise, Idaho, and four in Bergamo, Italy. Each line is capable of handling tooling for multiple products, but changeovers are kept to a minimum. Conveyors transport the boards from station to station, with machines and robots handling all solder/flux distribution and part placements. Formatter boards are built in multiple sets on a panel (multiple boards per panel based on size: one to four boards are built on a panel), and the boards are separated from the panels and packaged at the end of the line for shipment. There is a small amount of machine vision inspection done during the board assembly as process checks, but the majority of testing is done on the completed panels just before depaneling.

Testing of the boards is done in three stages, as illustrated in Fig. 2.2. First, the on-board connections are checked by testing continuity from the connector to the circuit board. These connections mate the board with a user's computer and with the printer engine in an actual printer. Second, the in-circuit test is done by a bed-of-nails tester which makes contact with up to four

hundred points on each board to check for circuit continuity (open and short circuits). The tester also does some minimal checking to ensure that the correct components are in place and in the right orientation. In the third test stage, the functional tests are performed using a different bed-of-nails test fixture. This process allows the tests to be downloaded about one hundred times faster than if the test programs were transferred through the board's standard input/output connections. Here, test data patterns are fed in at key points in the board, and the outputs are checked and compared with the expected values. In addition, the test simulates the printing of a test page by inputting a known set of data and reading the resulting data stream which would go to the laser in the actual printer.[4] After the testing stations, a bar code reader checks the panel and verifies that all boards have passed all of the tests, an operator visually inspects the panel, and finally a punch die depanels the boards and an operator packages them.

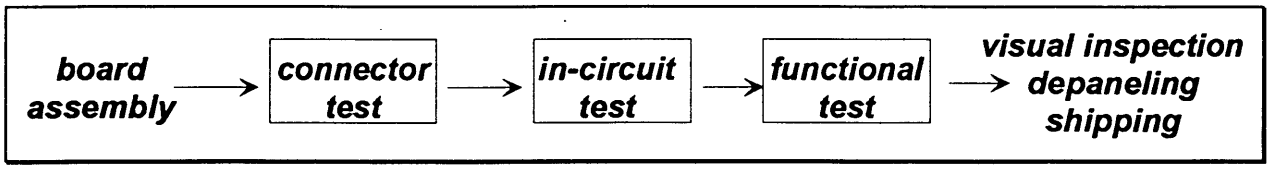

Figure 2.2 Formatter board testing process.

The actual test programs are divided into separate blocks so that it is easier to detect exactly where a failure occurred. Data stored as a result of a failure will include the particular test that failed, along with the actual and expected data values that were produced vs. what was expected from the test board. When a test failure occurs, technicians remove the panel from the production line and run more detailed tests to determine the root cause of the problem. If a legitimate cause can be determined, technicians repair the board and send it back through the line to be completely retested before shipping. The information about the repair is entered into a database by the technician and can be retrieved based on the reference designator (component

[4] In the past, each formatter was actually used to produce a page of output on a known good printer in the formatter factory, which would be visually inspected by an operator; this new test technique saves paper and avoids reliance on catching problems visually. The test system today, though, does not indicate exactly where during the print test a board fails; a real-paper print test would be more helpful in this respect.

code--oscillator, microprocessor, DRAM, etc.) and/or by cause code (dewetting, insufficient solder, etc.). Every failure has a cause code assigned to it, but sometimes a component code (reference designator) is not listed. For example, if solder balls are found on the boards but they are not near any major component, no reference designator would be recorded.

The first column of **Appendix A** contains the in-factory defect levels found on one recent formatter product for two months of production. The types of defects are broken out into two general categories: mechanical and electrical. The electrical defects include failures which can be traced to a problem in the interior of a component on the board, whereas mechanical defects include board-level or process-related defects. The labels for each row of mechanical failures describe the specific cause of failure observed in the factory. The electrical failures are classified based on the specific component which technicians determined to be defective.

Another format that an engineer at BPR developed for displaying failure data is "bullet hole charts," which graphically depict the magnitudes of many different types of failure modes. The possible defect causes and the component names are used for the x and y axes, and the areas of the circles drawn at each x-y intersection are proportional to the ppm rate for that particular fail mode. In addition, the colors of the circles are used to indicate trends over time. Green denotes improvement (smaller circle) compared to the previous week, red denotes degradation (larger circle) compared to the previous week, and a black circle indicates less than 200 ppm change from the previous week. See **Appendix B** for a sample of this charting method for one week of production.

The types of problems that appear during board testing in the formatter factory fall into three categories:

(1) Type I error (false pass). This error occurs when defective boards are allowed to pass through the testing unit. Type I errors for the formatter boards occur at the rate of about 200 ppm. Most of these are due to microprocessors; the tests at the formatter factory do not fully cover all possible functions of this component. Print tests during final assembly detect many of these missed defects. These defects commonly result in misalignments on printed test pages.

(2) Type II error (false fail). This category covers boards which are fully functional, but which fail the tests during board assembly, and occurs at a rate of under 1,000 ppm for the formatters. These errors can be due to improper alignment of the board in the tester or incomplete contact between the test fixture nails and the board itself. Foreign particles (dust, fiberglass from the printed circuit board, etc.) in the tester are a primary cause of poor contacts. Regular cleaning of the testers, automating the test process with ATS, and waiting to depanel the boards until after testing have all helped to minimize Type II error.

(3) Escapes. These are boards which are shipped, but which failed a production test or were not tested. Manual handling and loading of the boards in the testers was once a primary cause for escapes. With the recent completion of the Automatic Test System (ATS) which was designed in Boise, escaped boards are now quite rare. ATS is a series of conveyors and test fixtures which automatically pass board panels to empty test suites, test them, log their pass/fail results, and return the boards to the main line without any human intervention.

2.3.3 Final Assembly Process

Most PCAs are not sold directly to end users, but are integrated into a higher-level product, such as a disk drive, computer, stereo, or laser printer. This added level of complexity makes it possible to observe defects which relate directly to final product performance. For example, just as interactions between components with marginal characteristics cause problems at the board level, unplanned interactions or timing problems between the PCA and other parts of the final product may lead to functionality problems. In addition, testing at the final product stage provides an additional layer of security for catching failures that may have been missed in upstream tests due to intermittence or Type I errors. It may also expose components with latent defects which "grew" into legitimate failures as a result of board handling during shipping or through further use of the product.

Some of the formatter board products which were studied for this project are manufactured at H-P and shipped to one of three Canon facilities in Japan. These plants assemble the boards, Canon's print engine, and other components into complete laser printers.[5] The

[5]We analyzed available data to determine if failures were dependent on where the formatter board was manufactured (Boise, Idaho or Bergamo, Italy) or the plant where the final printer was assembled (Canon Electric or Canon Nagahama). Defect rates for each manufacturing and assembly plant were compared for crystal

assembly process takes about 30-45 minutes. During assembly, Canon applies power to the formatter board and uses it to test the printer at each stage of assembly. This results in the board being powered for a total of 25-45 minutes. Power is applied and removed from 11-20 times depending on the product.[6] When a board at Canon fails or is found to have a defect, it is sent to H-P engineers in Japan who repeat the functional tests originally performed in Boise. The engineers use functional tests and print tests in an actual printer in order to duplicate the failure which occurred on the line and learn more about the problems with the board. If nothing special is found, technicians in Boise receive the board and perform more extensive tests. For example, they perform tests using many of Canon's test programs from printer assembly. If H-P finds no problems, they send the board back to Canon to be retried. Less than 10% of the recycled boards come back to H-P a second time. If a board is returned twice, it is not shipped a third time, but it is analyzed again to try to learn more about the failures. Examples of past root failures from these boards are additional solder problems or intermittent component failures.

A change in the final assembly procedure provided the motivation for the project described in this thesis. Under the new assembly procedures, H-P purchases printer engines from Canon and integrates them with the formatters into complete printers at HP-owned integration sites in California and the Netherlands.[7] This process assembles the printer much more quickly. Total assembly time for each printer at the integration sites is 7-8 minutes, with only 3-4 minutes of power-up time (5-9 power cycles) and 3 test pages being printed, for a total of less than one minute of printing time. The disadvantage of this decrease in assembly time is the decreased power-on time and power cycling for catching additional defects and exposing latent defects. When boards fail at the integration sites, onsite H-P engineers review them and send them back to the board factory if necessary, where they undergo the same process described earlier. At the

---

oscillators and for all failures in general on a month-by-month basis for six months of assembly data for a single product; no significant pattern could be detected.

[6]Straight power-on time for the boards has varied from 15 to 30 minutes depending on the product.

[7]At the end of 1993, the process time between building a formatter board and assembling it into a printer at the integration site was estimated at 28-45 days.

integration site, engineers enter the available failure information (date, time, board serial number, error condition) into a database management system. This data can be combined with the formatter factory databases for further analyses.[8]

The returned boards from printer assembly sites fall into three general categories, and in roughly equal proportions as well: electrical malfunctions, mechanical malfunctions, or no-trouble-found. See **Appendix A** for a detailed breakdown of the types of failures which Canon's process and the newer H-P integration process detected for two similar products. We compiled this data from the board analyses performed by H-P engineers either at the final assembly site or at the board factory. We found that engineers are unable to find a root cause of failure for one-third to one-half of the functional failures during board assembly. This is primarily due to the nature of the assembly process. When a printer is not functioning during a particular test, the line worker often replaces the formatter board as a first step, since that is one of the most complex parts of the printer. If the printer then functions, the board is assumed to have been bad, when in fact the original board may not have been properly connected to the rest of the printer, or another part may have been out of place but repaired in the process of trying the new board. Since a smaller amount of testing done at the H-P assembly process than at Canon, we have less information regarding the cause of failure for the printers assembled at H-P. This is the primary reason for a larger proportion of the electrical malfunctions being analyzed as NTF from H-P (60%) than from Canon (27%).

The difference between the Canon and H-P processes could have a noticeable effect on final product quality. Since the Canon process takes more time and runs more tests, it would be more likely to catch intermittent fails and to excite latent defects to the point where they actually cause a product failure. Therefore, there may be latent defects which are excited by the Canon process but are not stressed enough by the H-P integration process to be detected; these would

---

[8]In addition to product testing at the assembly sites, BPR's quality assurance organization also performs periodic audits of completed printers. This includes stress screening and tests to ensure that all applicable regulations are being followed for noise, electromagnetic interference, etc. For each failure or defect found during this testing, engineers determine a root cause, and recommend a corrective/preventative action to manufacturing.

not be noticed until they reach the customer site. For example, the oscillator failure rates decreased sharply between the two processes--from 757 ppm at Canon to only 14 ppm at H-P. The components themselves and the board assembly process for the two boards are basically identical. Since oscillators are extremely sensitive to voltage and current, the additional power time and cycling at Canon may cause more weak components to be exposed.

2.3.4 Field Information

The goal of product testing, of course, is to improve a product's performance for the customer. As a result, field information is really the "bottom line" for how the product is doing; we must check each step of production to see how it improves final product quality. Using the intermediate measures described earlier (pristine board rates out of the formatter factory, printer assembly success or failure rates, etc.) provides information that is easier to understand and make use of locally. The tests which determine these measures are also important because they detect bad products before more value is added to them by later assembly steps. Failure data can also be helpful in determining the relative benefit of testing at different stages of production, particularly when the tests being studied are capable of detecting the same failure modes at different levels of the assembly process.

However, we need final product data in order to evaluate the entire product and production system. Basing quality judgments solely on intermediate quality measures would be fair only if the following assumptions are all true:

1) the tests are testing for ALL possible failure modes that would occur in the final product (no more, no less)
2) the detectable types of failures would not be otherwise caught at a later stage of product manufacturing
3) failures which are precipitated after completion of final assembly (due to packaging, shipment, storage, etc.) or through customer use cannot be affected or eliminated through product design or manufacturing changes

We can see from **Appendix A** that these assumptions cannot be accepted. Many of the failure modes from the printer audit in the last column are detectable at earlier stages. The same

failure modes also occur across multiple levels of assembly, which we should take to mean that no one level of testing is capable of detecting every single occurrence of certain failure modes. Field data should therefore be an important source of information. To use it effectively, however, we must understand its sources and its completeness. Described below are some of the potential problems with field information:

(1) The available data may not be complete. Many field failures and repairs are not reported to the company which produced the final product. Examples of these include:

(a) intermittent failures that a customer observes but does not report, thinking that they are insignificant or unrepairable
(b) minor malfunctions that the customer doesn't ask to have repaired because the warranty period has expired
(c) repairs that a customer does by herself
(d) repairs which are done by dealers or servicers who are not part of the company network.

Non-company repairs become particularly significant after the warranty period for the product expires. For example, H-P offers a one-year warranty with all of its laser printer products, but after that one year, a customer may find it cheaper to have repairs done elsewhere, in which case H-P will not be able to find out about or learn from the repair being done. The warranty department at BPR has estimated that only about 45% of all non-warranty repairs are done by H-P personnel.

(2) The failure data for a given repair may be misleading or incomplete. Each board replaced in the field may not represent an actual board failure. Field personnel may sometimes replace one component when another component is actually at fault. This may sound illogical, but we must consider that the primary goal of the repair person is to repair products as quickly as possible, not to perform root cause analysis on problems. In addition, H-P and other companies provide replacement components to repair people free of charge. In this case, the relative cost of one part over another would not affect a repairman's decision to replace it. In other words, there is no motivation not to wrongly replace formatters. To compensate for this in available data, we can collect a sample of boards which have been taken from customer products and determine which ones are actually "good" boards. Reducing the field data figures by this percentage will give us a better understanding of true quality levels, which would be critical for the purpose of designing testing strategies.

(3) Information on time-to-fail may not be available. If the product passes through several hands (warehouses, wholesalers, dealers) before actually reaching the customer, it may be impossible to determine when the customer actually purchased and began using the product. This is especially true if the entire distribution network is not owned by the company which produced the product. For example, H-P produces hard disk drives which

are sold to original equipment manufacturers (OEMs) who assemble them into their own units and sell them. Once the drives reach the OEM, H-P is unable to see where or when they are being used. To gauge the quality and effectiveness of their testing, H-P must settle for studying the accept/reject rates of the drives at the OEM facilities. These measured rates have been extremely volatile in the past and not very useful.

(4) The product use and environment may vary among end users. Time in the field is not a fair basis of comparison for product failures if the products are being used in different ways, especially when we are designing tests to keep failures from occurring for end users. For formatter boards, for example, use varies by product and by country. In printers which are able to go into "sleep mode", it is not necessary to turn the printers off each day; some models do not even have an on/off switch. On these models, the formatter board is constantly powered up. However, in some countries (particularly in Asia), users will still turn off their printer each night to avoid potential damage from the "brownouts" which occur in some countries as they switch power networks to adjust for changes in demand, or they will unplug the printer to conserve the energy.

Despite these potential limitations, field data is still a source for valuable information. As we design analysis methods for the data or draw conclusions, however, we should review factors such as those shown above to determine if we may be understating or overstating the magnitude of the problems we have discovered.

At Hewlett-Packard, dealers, repairmen, and engineers enter warranty and repair information into a companywide system database which can be accessed by H-P dealers, repairmen, and engineers. Over 50 different pieces of information may exist for a single printer in the field, including printer serial numbers, warranty start dates, repair dates, comments on the repair that was performed, and the component numbers of any parts which were replaced during the repair. If a repairman removes and replaces a formatter board during a repair, he sends that board to an H-P service center for further testing and repair. The service center is able to make a more detailed determination of the failure cause, but for a product with many parts in it like a formatter, technicians may still replace multiple parts to repair the board, without ever knowing for sure which part was at fault. H-P component engineers receive and study the removed parts, and the dealer receives the repaired board for use in a future field repair.

The most detailed data which repairmen normally record in the database is a comment such as "REPLACED FORMATTER. PRINTER WORKS OK NOW." To learn more about the actual failure causes on defective boards, H-P engineers and technicians receive the boards and perform detailed analyses. Quality engineers will perform this sort of analysis on several returned boards at the beginning of each new product's life in the field, and at other times as well.

We can use field information to answer the important question of how long printers have been used before they fail. The data analyses in this section use the following formula for "age" of the printer, and assume that this is a fair measure of the amount of use that the printer has received:

age = (repair-date) - (warranty-start-date)

When drawing conclusions from this data, we must be aware of some potential problems:

(1) The warranty may have been restarted more than once. For example, a printer can be on consignment under warranty and then resold as a trade sale, at which time the warranty is restarted for the new user.

(2) Little or no data is available regarding specific failure modes. The data only provides information on whether a formatter board has been involved in the repair or not. Therefore, time-to-fail studies for specific failure modes must be accomplished through data collection outside of the normal system.

(3) We do not know the exact amount of lag between the time a formatter board is produced and the time it begins use in a customer's product. For example, since assembly from Canon does not show which formatter serial number is installed into which printer serial number, it is difficult to exactly determine how long this time is, but it has been estimated at up to six months for some products.[2] This means that a major problem in the formatter factory which cannot be detected by production tests may not be noticed until six months later, or until six months of "bad" product is in the field; this is one reason why comprehensive production testing can be so important.

Once we understand the limitations of the available data, we can proceed to compile it in ways which help us draw conclusions. Fig. 2.3 represents one possible method for charting warranty repair data.

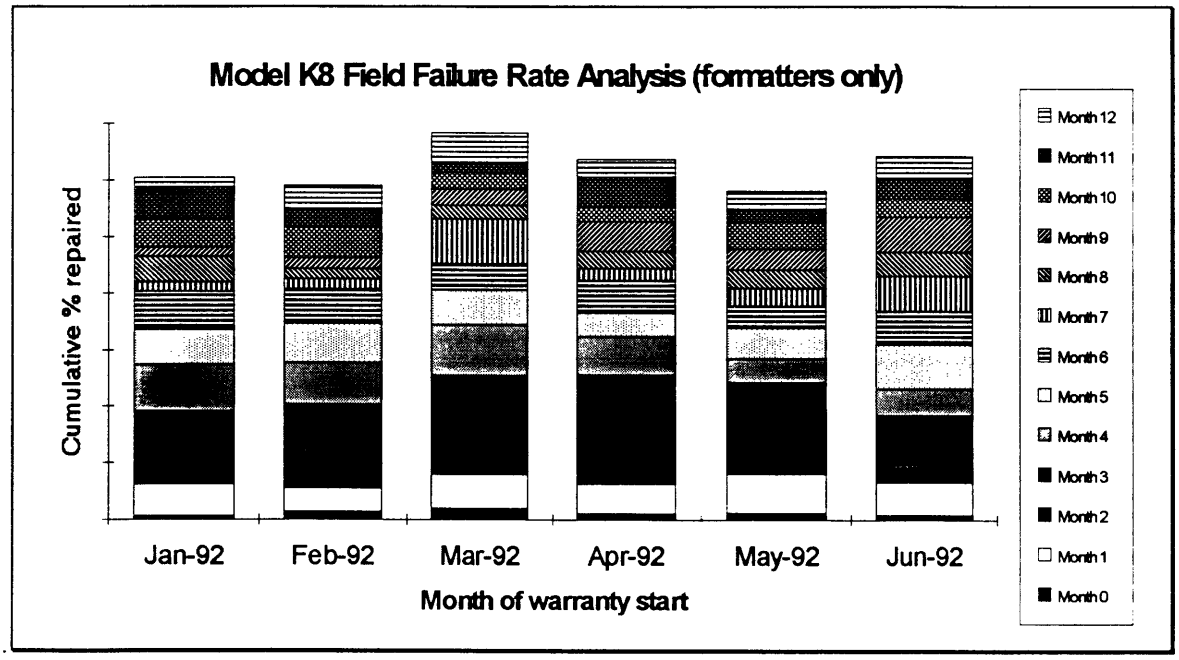

Figure 2.3 Field Failure Rate Analysis for formatter repairs.

This chart depicts every Model K8 printer which began its warranty between January 1 and June 30, 1992, and which was repaired at some point during the following twelve months. For example, the first column consists of printers which began their warranty periods during January of 1992. The bottom block, labeled as "Month 0" in the legend, depicts the percentage of all of those printers requiring a formatter-related repair within that same month. The block labeled as "Month 1" is related to the number of printers which were called in for service in February of 1992, and so forth.

We can use this graph to look at two types of trends. First, we can compare the total height of each column of data across each month to get a feel for overall product reliability changes. If the columns decrease in size for more recent warranty-start months, we could assume that more recent products tend to be more reliable. We can make the same sort of comparison for any individual month as well. Looking across the chart at the sizes of any similarly-patterned blocks indicates trends at specific points in product life. Second, by looking within each column, we can learn about the failure rates of a given population of printers over time. For example,

looking at the May 1992 column, it appears that the blocks near the bottom of the bar are larger, and that they gradually decrease in size as we look further up the bar. This indicates that the product is more likely to require repairs early on in its warranty period as opposed to later. The implications of these types of conclusions will be discussed in more detail in Chapter 3.

For comparison, Fig. 2.4 illustrates the same type of analysis described above, but for all repairs on the printer **except** for the formatter. This includes mechanical parts such as paper feeders, rollers, or cooling fans, as well as less "complicated" electrical components such as fusing assemblies or power supplies.

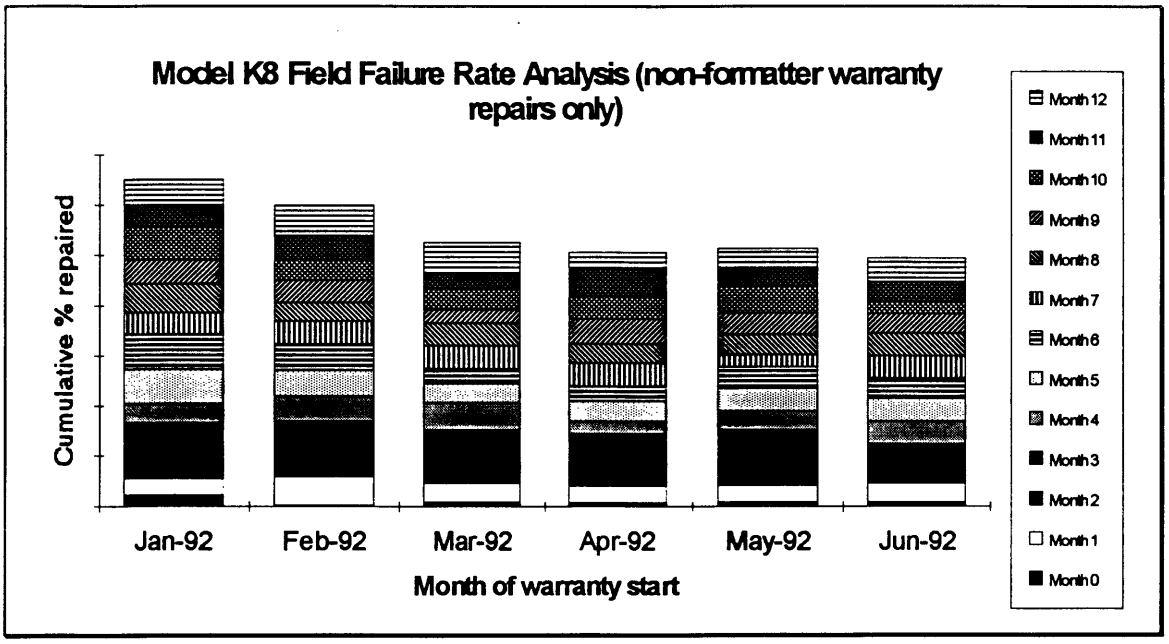

Figure 2.4 Field Failure Rate Analysis for non-formatter repairs.

In this figure, the height of each column decreases more consistently from left to right than for the formatter data in Fig. 2.3. This means that printers produced later in time tend to be more reliable with respect to their mechanical components. However, the uneven shape of Fig. 2.3 tells us that formatters do not follow this trend. In addition, the size of each block is much more consistent for non-formatter repairs than for the formatter. This means that non-formatter repairs

tend to occur at more regular intervals than the formatters. This relationship will become more obvious in the next charting technique.

Another way to look at this information is to examine the failure rate levels as a function of number of months in use, as opposed to month of start of warranty. Fig. 2.5 is an example of this. Each individual line represents the same information as one column did in the above tables. Here, we can more easily see that formatter failures for this particular product, seem to be at their maximum around 2-3 months into the warranty period, and then gradually taper off after that.[9]

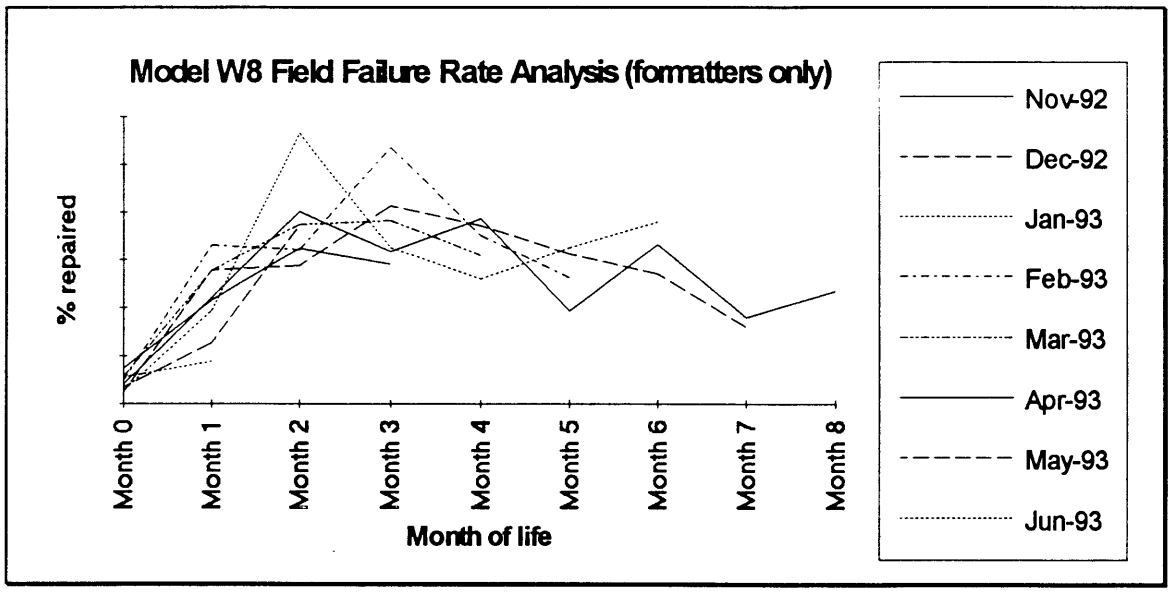

Figure 2.5 Field Failure Rate Analysis for formatter repairs by month of life.

The next three figures employ this same type of analysis, done for the same printer model depicted in Fig. 2.3 and 2.4. In these figures, the results for six different warranty-start months were averaged to produce the single line shown.[10] Note that the reliability seems to improve drastically after the twelfth month in most cases. This is due to the "warranty effect" described earlier. After the one-year warranty period is over, customers are likely to take their printers

[9]Note that some of the lines are incomplete for the later months; at the time of the analysis of this relatively recent product, there were not eight months of failure data available for the more recent months of production.
[10]Six months of data were available for each month of life referenced on the x axis; this was not the case for the product in Figure 2-3.

elsewhere to be repaired. We should therefore only use the first year of data for drawing conclusions on failure behavior for formatter boards.

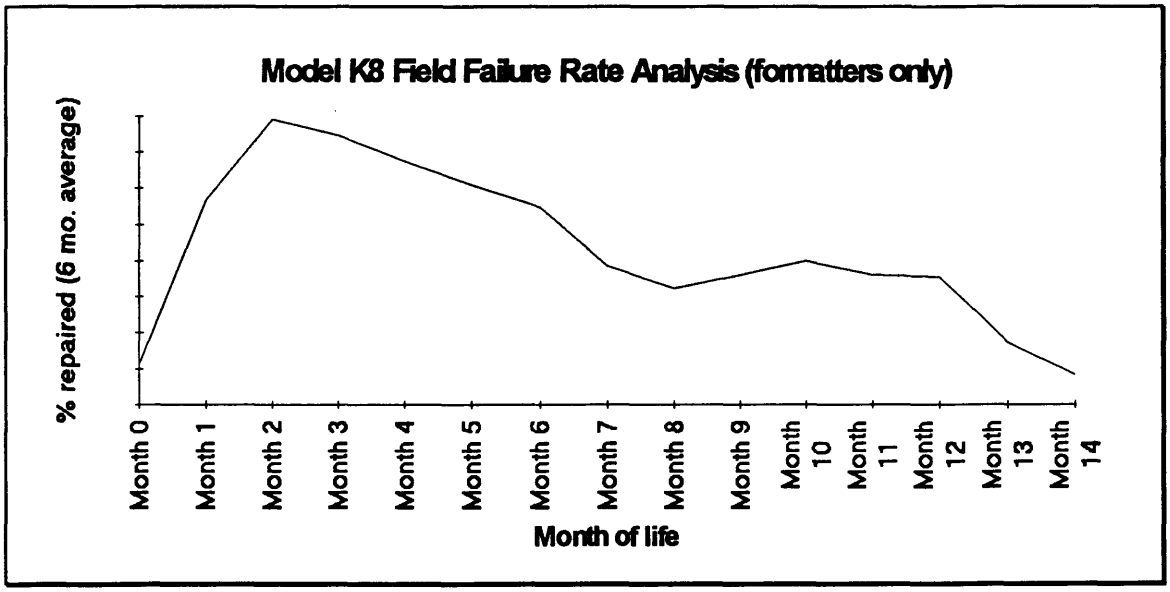

Figure 2.6 Composite Field Failure Rate Analysis for formatter repairs.

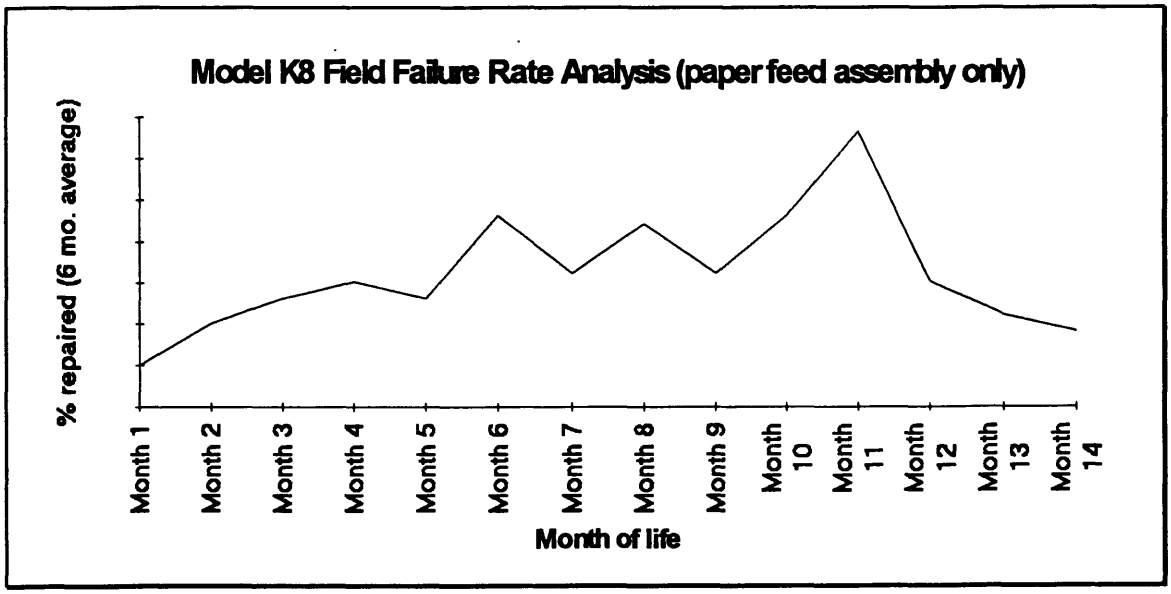

Figure 2.7 Composite Field Failure Rate Analysis for paper feed assembly repairs.[11]

[11]The paper feed assembly, as the name applies, is the mechanical subassembly of the printer responsible for transferring paper from the paper tray to the roller assembly, where the toner is actually applied to the paper.

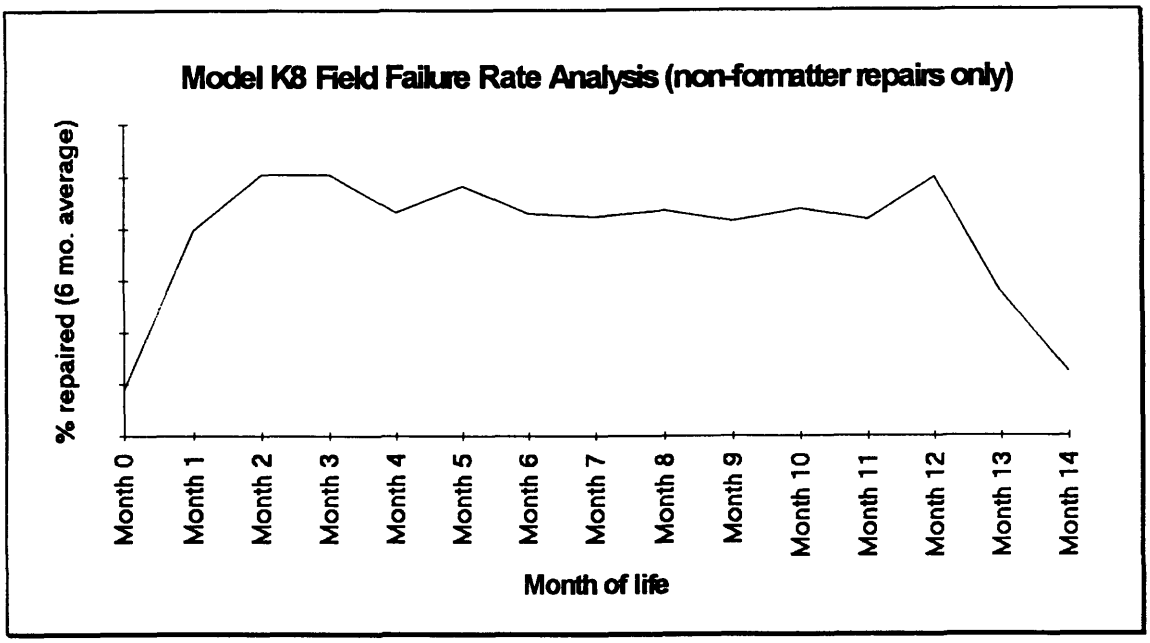

Figure 2.8 Field Failure Rate Analysis for all non-formatter repairs.

Comparing thesis analyses, we can see that formatter repairs are most likely to occur within the first few months of the warranty period. Repairs for the rest of the printer (Fig. 2.8) are equally likely to occur at any time. Finally, repairs for a purely mechanical component, such as the paper feeder in Fig. 2.7, are difficult to characterize. They seem to gradually increase with time, but in a very irregular manner. In Chapter 3, we will see how these failure shapes relate to the behavior of latent defects.

2.4 Additional Sources for Useful Information

There are many sources for information regarding the behavior of products over time. The above examples for laser printer formatter boards depict some of the possibilities and also some of the limitations of the available data. We would be able to draw more definite conclusions if the following additional types of data were available:

* *The ability to match printer serial numbers (from warranty reports) back to formatter board serial numbers.* This would make it possible to discover patterns in the boards replaced in the field. For example, failed boards may come from specific productions lines or be more likely

to be subjected to manual repairs during board assembly, which would lead us to looking at that particular assembly line or the repair process for root causes to the problems.

* *More information on board failure reasons from field repairs.* The information we normally receive from the field only tells us that the board was replaced, as opposed to the specific cause for failure. One way to learn more from each repair has been implemented at a disk drive division of H-P. There, the manufacturing facility directly receives all of the products requiring repairs. With both production and repair occurring side by side, it is relatively easy for engineers and operators to collect information about failure modes in the field and to share that learning with the production personnel.

* *Information on use conditions of products called in for repair.* Age is currently the only criteria available for analyzing the usage history of products. It may be possible for the formatter board to store information regarding number of pages printed, number of power cycles, non-fatal error conditions, etc. which would help repairmen and failure analysis people to better understand what is leading to the failures being observed. For example, some automobiles are fitted with diagnostic connectors which allow repairmen to instantly retrieve data from the car's onboard computer regarding past operating conditions and problems.

2.5 Applications and Summary

Within the production process, there are several possible locations for gathering information and testing. Fig. 2.9 summarizes the sources of information presented in this chapter.

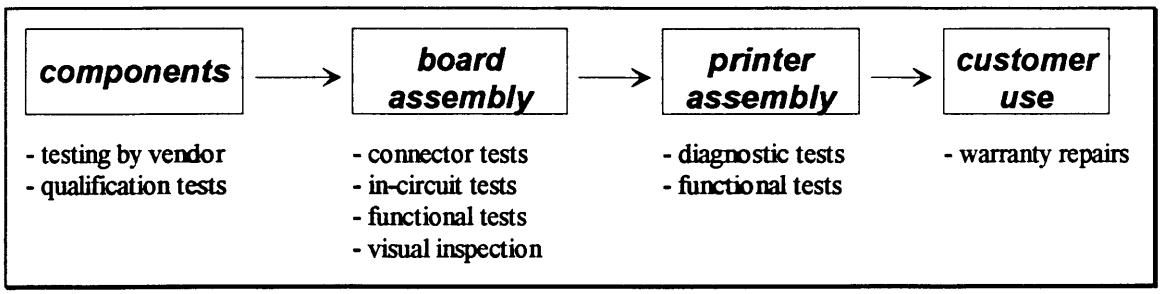

Figure 2.9 Sources for board quality information.

It is important to test at various stages of the production process. By testing a product soon after a defect is created, we avoid adding more value to an already defective product and creating useless inventory. Early testing also allows quicker feedback to the process which created the defect, making it less likely that more of the same problem will occur. On the other hand, by waiting for the product to be more fully assembled, we can test the board as it would

function in real life, and we can catch problems with interactions between subassemblies as well as latent defects in subassemblies which were triggered by assembly steps after they were first assembled. Fig. 2.10 illustrates these complementary advantages.

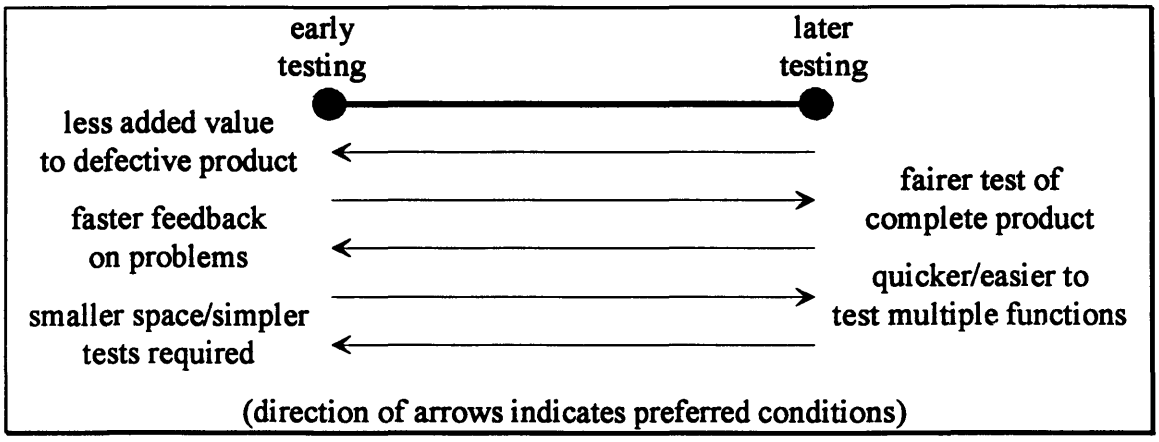

Figure 2.10 Advantages to early and late testing during manufacturing.

In addition, we can learn about latent defects in a product by looking at redundant test procedures--tests at different assembly steps which are looking for the same sorts of defects in a product. As we saw in Appendix A, the testing processes for the formatter boards contain these sorts of redundancies, in that some of the same classes of defects are detected at each stage. Comparing test information in these ways provides important clues to where defects are being generated in the process. These analyses can also serve as a check on the accuracy of upstream test stages.

This chapter emphasized the importance of collecting relevant data from throughout the entire product process. Although the goal of each individual test is not to detect latent defects, we can nevertheless learn about the behavior of latent defects in the products. This is due to the fact that the available information comes from multiple locations in the production process, data is available for many products over time, and the tests are capable of detecting similar types of failures at each step. These characteristics make it possible for us to compile this data in whays

which help us understand how failures are occurring over time. Now that we have gathered the data relative to latent defects, we can apply the theories of latent defect formation to help us fully understand the *problem* at hand. This is the purpose of Chapter 3.

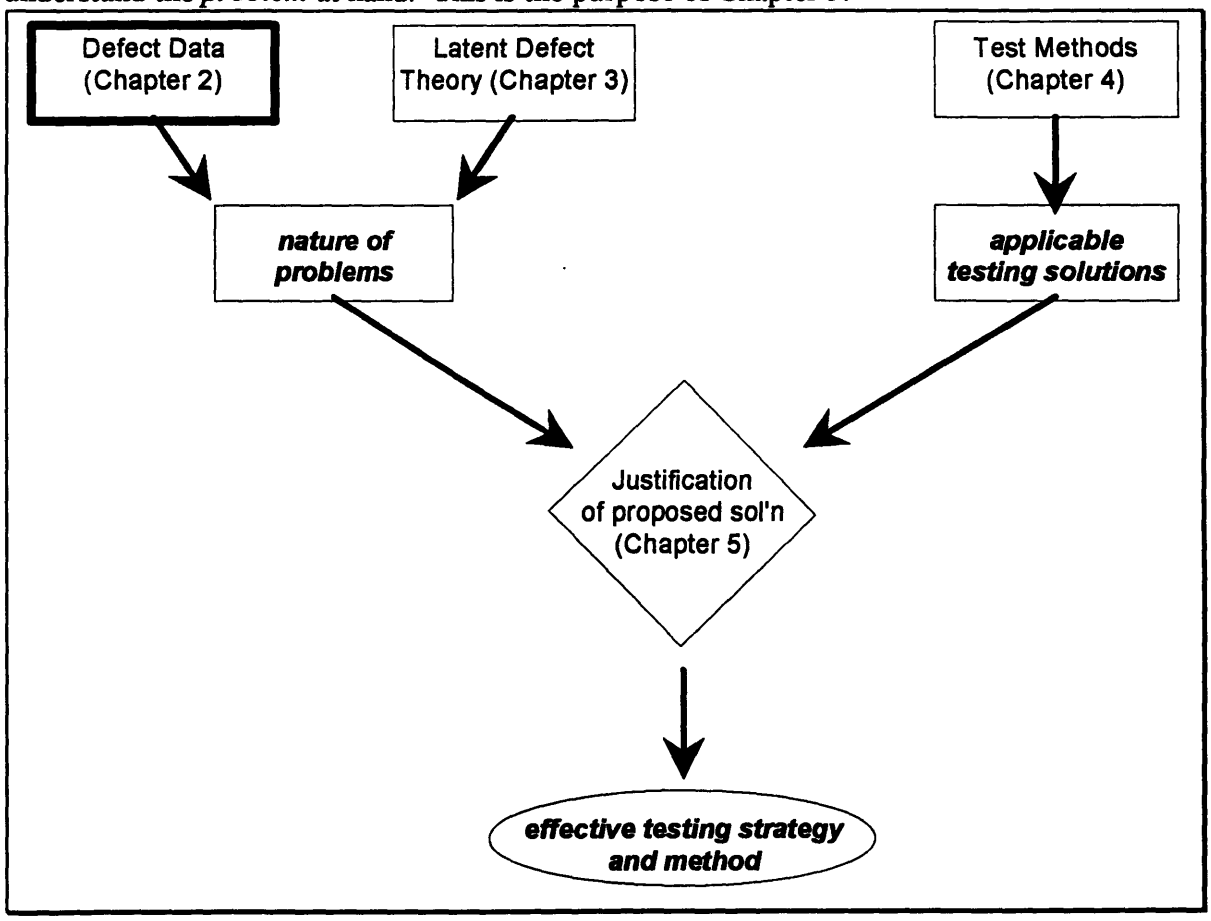

Figure 2.11 Review of latent defect testing decision process.

References

[1] Pitta, Julie, It Had To Be Done And We Did It, Forbes, April 26, 1993.

[2] Moss, Dick, Hewlett-Packard Design for Reliability Manual, 1990, p. 10-20.

Chapter 3 - Properties of Time-Dependent Failures

3.1 Introduction

In addition to understanding and gathering available data, we must also be able to recognize patterns in the data and what they mean about behaviors within the product. This chapter provides theoretical background on the causes and behaviors of time-dependent failures. We will review the concept of the bathtub curve, a generally accepted model for the reliability of electronic products over time. In addition, we will also look at some alternative theories and variations which have been proposed by various researchers. Finally, we will use the concepts outlined in this chapter to draw additional conclusions from the data which were presented for formatter boards in Chapter 2.

3.2 Root Causes and Precipitation of Latent Defects

As mentioned in Chapter 2, latent defects can originate at any stage of the board assembly process. Fig. 3.1 illustrates the variety of processes which can cause a failure mechanism to originate, possibly leading to an actual failure:[1]

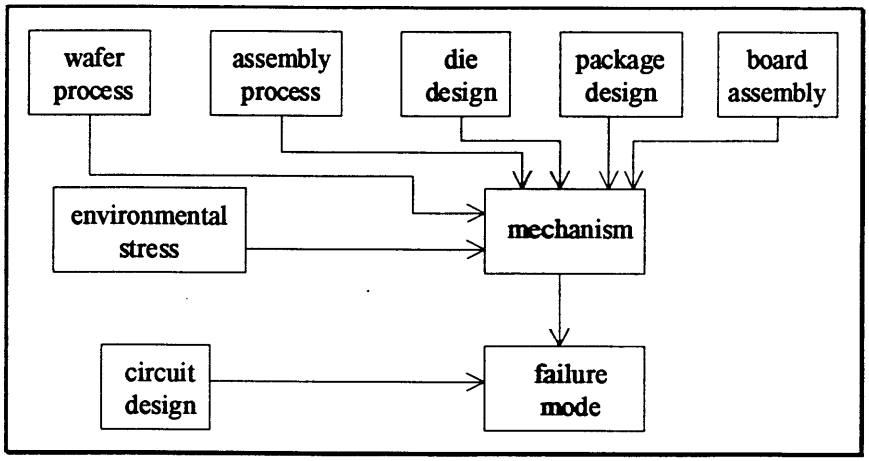

Figure 3.1 Potential sources for latent defects.

This diagram also illustrates the fact that the seriousness of a latent defect mode is dependent on the circuit design. For example, if redundancies are built into the device in order to guarantee that functionality is still possible even when circuits break down, failure modes within one piece of the product may not be as important as if the entire product relied on one specific component functioning at all times. Space shuttles and satellites are two examples of systems which make extensive use of redundancies to ensure reliability.[2]

The root causes of latent defects can be due to any of a variety of properties of materials: thermal, chemical, mechanical, electrical, and so on. Understanding how latent defects begin at their most basic level will help us in analyzing failure information. The following examples illustrate some of the basic phenomena which lead to the formation of defects:

(1) Electromigration. Electromigration is caused by electrons as they flow through an integrated circuit (IC) to create current. As these electrons flow, they create momentum which can erode away the interconnect metallization that forms the paths for the electron flow within the device, similar to a mudslide traveling down a highway and taking chunks of asphalt along with it. This can eventually lead to voids in the conductors and create open circuits. In addition, the process occurs much more rapidly if there are imperfections or contamination in the traces, causing product malfunction earlier in life.[3]

(2) Spiking. During the process of integrated circuit fabrication, layers of different materials (or differently charged materials) are deposited on top of each other to create the circuits. For example, silicon may be diffused into aluminum, aluminum may be in turn diffused into the silicon, etc... If there is a void in one of the layers or a variation in the thickness of a layer, the circuit can wear out that weak point and create an electrical short between two layers which were meant to stay separated.

(3) Dielectric breakdown. Similar to spiking in some ways, dielectric breakdown occurs as a result of defects in the semiconductor layers in an IC. It can occur at the beginning of a product's life or much later on in time, depending on the type of breakdown. In all cases, the breakdown begins when charge builds up near the interface of the silicon (which allows the flowing electrons to create current paths) and the silicon dioxide, which normally serves as an insulator to regulate the amount of electron flow through the silicon. During this charge buildup, charge carriers are inadvertently injected into the oxide. This creates localized heating, which causes the SiO_2 or Si to vaporize, in turn causing a short between the silicon substrate and the other layers of the chip.[4]

Most latent defects are fundamentally caused by contamination or a material weakness is aggravated by some sort of stress, leading to an actual defect. This applies to failures within the circuit board as well as within a specific IC. Examples at the board level include the following:

-- particle contamination on the solder pad as an IC is soldered to a board
-- the bonding of materials with different expansion/contraction rates
-- an insufficient amount of solder being used to connect a component to the board

In all these cases, latent defects can grow as a result of repeated heating and cooling, causing a broken connection. The temperature changes could be caused by the environment or through the heat generated via current flow. We can think of the root causes of latent defect-based failures being like a crack which does not initially affect product function, but which can propagate to create a real defect. If the stresses during product use or testing exceeds the inherent strength of a given material, the failure mechanism will shift from initiation to propagation. Therefore, as long as stresses stay under this value, cracks will either not appear or will stay negligibly small.[5] This model for failures can be helpful in predicting when changes in failure rates are going to occur in a system.[6]

Table 3.1 lists several types of failure mechanisms which can occur in electronic products, both at the chip level and at the board level:[7]

Failure Mechanisms	Acceleration Stresses
Fatigue crack initiation	Mechanical stresses/strains, temperature cycling, frequency
Fatigue crack propagation	Mechanical stresses, temperature cycling, frequency
Contact/interface wear (friction)	Contact force, relative sliding velocity
Diffusion	Temperature, concentration gradient
Moisture within IC package	Humidity, temperature
Corrosion	Temperature, relative humidity
Electromigration	Current density, temperature, temperature gradient
Radiation damage	Intensity of radiation
Surface charge spreading	Temperature
Stress corrosion	Mechanical stress, temperature, relative humidity

Table 3.1 Root causes of latent defects.

On the right side of the table are the types of stresses which can generate or expose the failure mechanism. Note that all of these stresses could occur during product use to turn latent defects into real defects. On the other hand, they can also be duplicated as part of test procedures to accelerate these latent defects and force them to reveal themselves before the product leaves the manufacturing location. The idea of accelerated stress testing will be discussed in more detail in Chapter 4.

3.3 The Bathtub Curve

Based on our understanding of the behavior of latent defects, we might postulate the following hypotheses:

(1) Component strength deteriorates with time (even though the rates of deterioration may well be different).
(2) Weak components deteriorate faster than strong components.[8]
(3) If printed circuit boards or devices last beyond a certain point in time, they will run without failure for a very long time.[9]

By applying general concepts such as these to analyze information, researchers have developed several models to describe failure behavior as a function of time. The most commonly accepted general model for failure behavior is that of the Weibull curve, more commonly known as the bathtub curve, due to its shape:[10]

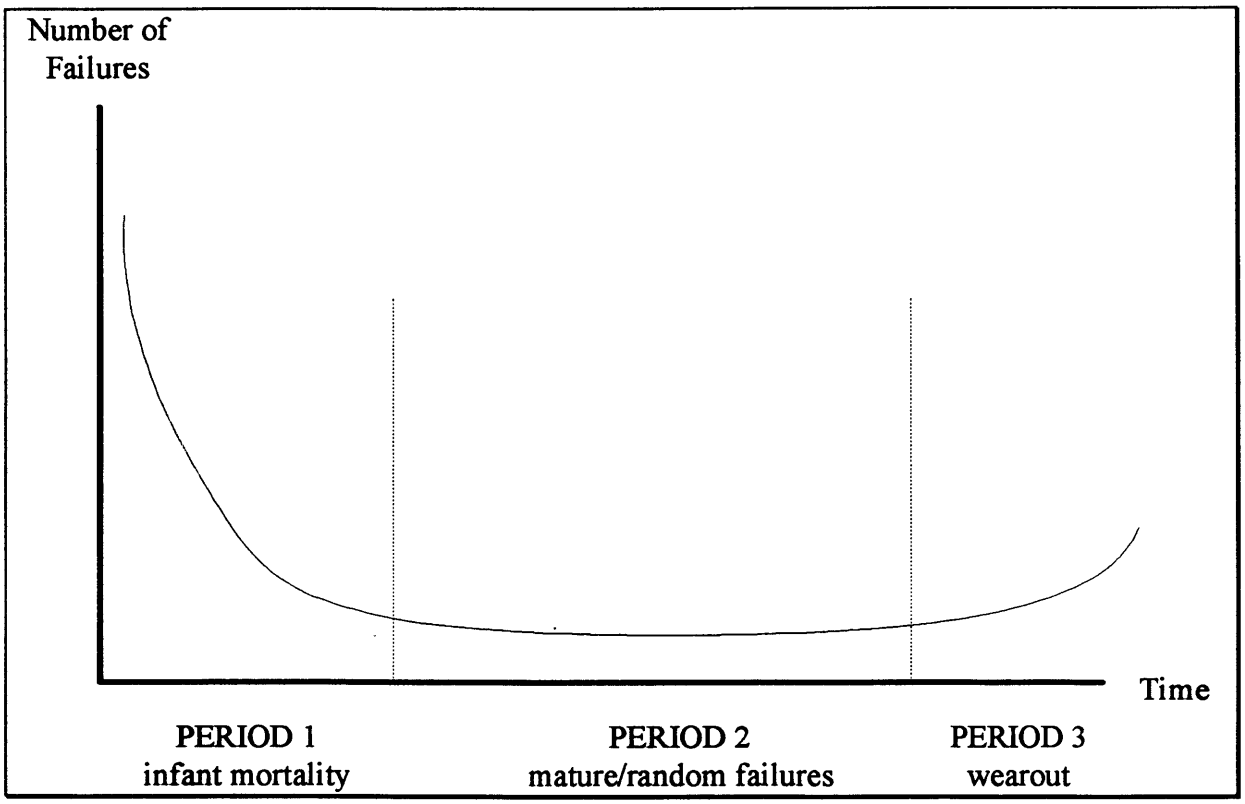

Figure 3.2 Weibull, or "bathtub," curve

As shown above, the model breaks the behavior of products into three periods. Failures during the product life fall into one of the following three categories:

(1) Infant mortality. During initial product use, the failure rate is initially relatively high, but it drops quickly as components which were initially weak or were damaged by the manufacturing process fail. The length of this period varies depending on the specific product, but for standard technology, solid-state type electronics, a year is considered typical for Period 1. For example, AT&T Components and Electronic Systems considers the infant mortality period to be the first 10,000 hours of product operation. They categorize failures which occur in this period in two ways: dead-on-arrivals, which fail immediately upon use or which never do function, and device-operation-failures, which occur after some period of operation.[11]

(2) Mature/random failures. After the infant mortality period, the failure rate settles down in Period 2 to a much more constant, lower rate for most of the product life. These failures are also called "useful life failures," since this is the period when the product would be in normal use. We assume that failures occur randomly during this period, and that they are primarily due to the inherent variation within parts or the stresses placed upon them.[1] In

[1] Although this explanation is generally accepted to be true, it's important to note that the bathtub pattern, specifically the useful life period of constant failure rate, has been identified in strictly controlled life-tests, where random fluctuations in, for example, either temperature or supply voltage do not take place. (Jensen)

other words, a product which is expected to fail during this period would be just as likely to fail after one year as, say, after ten years. It is counterproductive, however, to think of these failures as being *entirely* random. Believing in the concept of random failures can cause people to accept failures as a natural event, rather than digging into and learning from the root cause of each failure.

(3) Wearout failures. As products reach Period 3, their failure rates start to increase again, due to the inevitable degradation of the parts over time. The length of this period is often difficult to determine. Researchers estimate that a stable product subject to no random failure-type problems could last between eighty to one hundred years in the field.[12] We can conclude from this that most electronics in use today would be obsoleted, be replaced, or outlive the operating life of the rest of the system before wearout-type failures.[13]

We note that this model has been tested and verified only for electrical components and failures, rather than mechanical. In general, mechanical products have not been shown to exhibit these sorts of behaviors. For example, compare the figures of failure behaviors in the previous chapter. The plot of electrical failures (Fig. 2.3 or 2.6) shows a decreasing level of failures which we associate with infant mortality, whereas the graphs for mechanical parts or the total of all parts (Fig, 2.4, or Fig. 2.7 and 2.8) exhibit constant or fluctuating failure levels during the same time period.

3.4 Alternative Theories and Considerations for Failure Behavior

The theory of the bathtub curve is applicable to a variety of situations. On the other hand, researchers have developed modified or alternative theories which have been shown to more accurately describe particular aspects of product failure behavior for certain applications. In the following sections, we will review some of these newer concepts. None of them directly contradict the general theory of the bathtub curve. It is therefore most useful for us to view these as enhancements or clarifications of the details of potential product behavior. In addition, some of these concepts will apply more accurately to certain types of products over others. Once we understand the data and the theories, we can better explain the problems we are observing.

3.4.1 Underlying Assumptions for Failures

As stated earlier, we can think of components as possessing different levels of strength. Their relative strengths determine their length of life in the field, and their varying life spans in the field are due to their failure behaviors. Wong articulates this view more specifically in his Unified Field Failure Theory, which states that:

> "Flaws are built into a piece of equipment. Stresses, external or internal, bring the flaws to a state that causes an item to malfunction. This same process occurs from manufacturing test through field operations. Corollary: Under the same set of stresses the failure rate of an item decreases with time/cycle of stress application."[14]

In other words, failures are not entirely random. They all have some definable root cause, such as a deviation in the production process or during synthesis of a raw material. As a result, a decreasing failure rate over time is due to these pre-existing failures being precipitated during the life of the product. It follows that we should be able to use the same failure rate equations to describe failure rates during screening as well as field failure rates.

A variation on this model for levels of component strength has been put forth by Trinidade, who proposes that a product population can be grouped into "mortals" and "immortals."[15] For each possible failure mode, mortals possess the latent defect, or the potential for that failure mode, and the immortals do not possess that fatal flaw and will thus never fail for the associated defect mechanism. Trinidade states that it is not accurate to look at failure data relative to the entire population. Instead, we should make conclusions based on the percentage of "mortals" that fail during each period of time.

We can also think of the strength distribution of parts as bimodal or trimodal. A bimodal distribution consists of a freak distribution of parts with substandard strengths and a main distribution of strong parts with an acceptable distribution of strengths around a design value. With a trimodal distribution, we would also consider infant mortality failures as being a third type of grouping. These would have different, more serious failure modes than the freak

subpopulation. Examples would include cracked chips, open or nearly open bonds, mask defects, bad package seal, foreign material contaminants, bad welds, etc. Fig. 3.3 illustrates where each of these distributions would occur in time. The freaks would have the same failure mode(s) as the main population, but they would occur earlier in time due to their being inherently weaker. The bimodal concept describes the general classes of failures that can be expected. For special cases or for specific types of failure modes, though, other distributions such as the trimodal or a more complicated multimodal distribution may be more accurate.

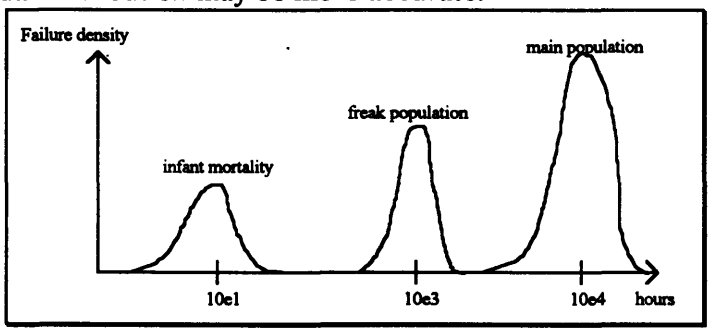

Figure 3.3 Trimodal failure distribution.

The trimodal concept dictates that effective screening for early failures would require two types of tests: (1) freak failures may be screened out by burning in components for a fairly long time (100-1000 hours) in an accelerated environment, and (2) infant mortality failures may be screened out by a short-term (10-30 hours) equipment burn-in in a normal (or slightly accelerated) working environment.[16]

3.4.2 Test Strategy Implications of Failure Theories

The three concepts for failures described above--unified theory, mortals and immortals, and modal distributions--all imply that there is really no random period for product failures. As a result, it would be effective to screen for failure modes that have a constant or even increasing failure rate over time, because we would be identifying mortals for that failure mode, and eliminating them from the population. Each failure should also provide us with information we can use to prevent other failures from occurring in the future, regardless of when that failure occurs. By contrast, the bathtub curve implies that we can effectively test for Period 1 failures

only, since every unit in the population is equally likely to fail during its useful life. Therefore, removing certain products has no beneficial effect on the failure curve over time.

These distinctions have an important impact on the testing strategy we choose to pursue. If we are able to learn from each failure, it may be unnecessary to test every single product in order to improve quality. On the other hand, if we are unable to eliminate future product failures using the information from our tests, we will be forced to test every single product, especially if our defect levels are already relatively low. Chapters 4 and 5 will discuss the impact of these distinctions in more detail.

3.4.3 Effect of Individual Failure Modes on Overall Failure Characterization

Graphs of failure curves from actual data seldom look as clean as the ideal bathtub curve in Fig. 3.2. This may be due to a relatively small quantity of available information, but it could also be a sign of particular behaviors which are worth investigating in more detail. Wong's research shows several examples where plots of failure rates over time have shown abnormal-looking "bumps" at various times, as depicted in Fig. 3.4.[17] These bumps may be due to "failures of small flaws left over from major flaw groups due to limitations in inspection and test". In other words, measurable variations in the failure rates can originate from inspection techniques designed to catch all of a particular defect type, but only down to a particular size. For example, a testing screen may detect cracks in thin film resistors larger than 0.1μm. The remaining defects, all being of similar magnitude, will then propagate and surface as defects at about the same time. Irregularities can also be due to different failure modes, each with a different time-to-fail.

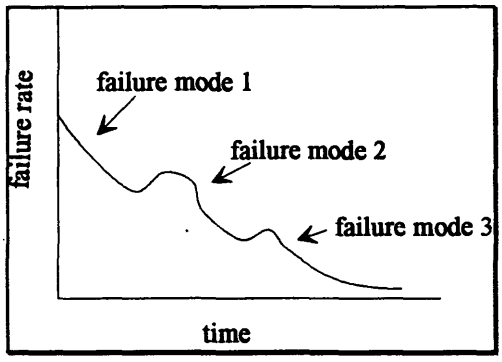

Figure 3.4 Effect of individual failure modes on overall failure plot.

To understand the source of these variances, we must examine the causes of the failures specifically during those times. For example, certain types of failure modes (crack propagation, dielectric breakdown) accelerate as the number of stress cycles or the amount of use increases. Many failure modes which behave in this way have *crack propagation* as their root cause. In this situation, each stress cycle causes tensile stress at the end of the crack. The stress is moderated by the surrounding material flowing to create a rounded end of the crack. This cycle is repeated over and over until the crack is large enough to cause a failure, such as a film resistor going out of tolerance or a solder joint to open.

It is also important to keep in mind that when we draw a single failure curve for a laser printer or similar complex device, we imply that all failures are created equal, rather than regarding some failures as more serious than others. For example, basing conclusions on the warranty data plots shown in Chapter 2 assumes that anything requiring a service call was equally serious. In some situations, this may be misleading, such as when problems occur which do not permanently affect operation of the product. A customer may not report these failures, but the manufacturer would still be interested in knowing about them, because they impact the product's quality reputation in the field.

3.5 Applications for Available Data

Once we understand the underlying behavior of latent defects, we need to determine which of the theories or assumptions are most applicable to the real-life situation under analysis. In the

remainder of this chapter, we will use information and data for the formatter boards to illustrate the issues which need to be considered at this stage.

3.5.1 Analysis of Data

There are many possible types of mathematical models which can be used for interpreting data on failed products. Examples include Weibull distributions, hazard distributions, instantaneous failure rate distributions, etc. Before we perform this sort of analysis, we must satisfy two important criteria:

(1) The accuracy of the data must be established. Analysis of the data is only as accurate as the underlying assumptions and quality of the available data. To illustrate this, consider the factors needed for a Weibull analysis, one possible method for graphing and analyzing time-dependent data. The parameters of the Weibull formula make it possible to adapt the function to a variety of life distribution patterns. It is defined in terms of time, t, as:

$$f(t) = \frac{\beta}{\eta^{\beta}} t^{\beta-1} \exp\left[-(t/\eta)^{\beta}\right] \text{ for } t \geq 0$$

The Weibull distribution is similar to a normal, exponential, or other standard type of distribution, in that it can be defined by two parameters: shape (β, exponential when $\beta = 1$), and location (η). If we wish to perform a complete Weibull analysis with field data, we must have time-to-fail data for failed parts as well as the ages of all the products in the field which have not failed as of the time of the analysis. For many electronic products, this information is hard to obtain from warranty data. It may also be inappropriate to extrapolate fair estimates for the population, based on the sample of good information which *is* available. For example, one of H-P's disk drive divisions closely tracked its field information to check the consistency of the data which was regularly collected, but no consistent pattern existed. Engineers attribute this to the fact that a bad disk drive in the field could quickly be replaced by a good one waiting in the inventory pipeline or on the customer's shelf. As a result, the customer had no incentive to immediately report the failure, making it hard for H-P to know exactly when the failure occurred. Weibull is thus more useful for analyzing the results of a controlled experiment than for analysis of warranty data.

(2) "Conflicting" data must not be allowed to dilute each other. If a limited amount of data is available for each individual failure mode, we may need to combine data for different modes to produce a statistically significant sample size. The danger in this is that the resulting average of the data will not expose failure details which may be significant. For example, if one failure mode in the data is sharply increasing over the same time period

that another mode is sharply decreasing, the resulting plot will show a constant failure rate, which would be misleading:

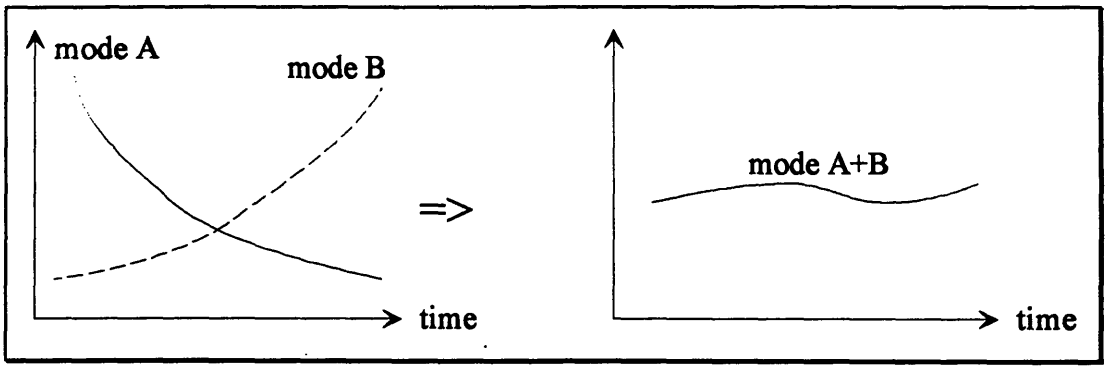

Figure 3.5 Illustration of dilution of failure modes.

We should therefore break down data into specific categories as far as possible without losing significance. Potential broad categories of classification include the following:

-- mechanical vs. electrical fail modes
-- "class failures" (episodic) vs. persistent failures
-- intrinsic failures
-- design vs. process failures

3.5.2 Relations Among Failures and Testing Procedures

Chapter 2 illustrated the types of information which are available from tests at each step of the assembly process. Based on the data we collected, we can draw the following conclusions about defects as they relate to our ability to test for them:

(1) Defects and problems do not occur at consistent levels over time. This should come as no surprise to people responsible for day-to-day firefighting in production. Defect levels at each stage will increase and decrease for periods of time, due to process problems or changes. For example, the printer assembly plant once observed a sharp increase and later decrease in quality problems which was due to an insufficient solder problem back at the formatter factory. In another case, final assembly found a sudden change in the frequency of parts missing from connectors on the formatters. The final assembly process was found to be primarily responsible for this, even though people initially assumed that the boards were arriving from the board production facility in this manner. Since problems rise and fall in severity, it is important to collect data for long periods of time to avoid

overfocusing on current, but relatively minor, issues at the expense of more severe problems.[2]

(2) Defects will impact multiple stages of production, even though tests are designed to eliminate these defects. Some defects may pass a first test, only to be caught at later stages or to make it out to the field. This can be due to intermittent types of failures or the limits of test procedures in reliably catching defects. For example, engineers at H-P studied the correlation of crystal oscillator failures at formatter assembly in Boise with final printer assembly at Canon, since these types of problems could be detected at both locations. They noted the failure rates at Canon, and boards which failed were traced back via their serial numbers to their build date in Boise. They then looked at the data relative to time. The study showed a strong match in the fail rates at Canon and Boise. We can conclude, then, that the tests at Boise cannot catch absolutely all of the problems that occur; a percentage of them escape the factory.

(3) Final quality levels need to be analyzed to determine the effectiveness of in-house testing. In testing, we wish to develop a program severe and comprehensive enough to catch every weak or failed board that can be caught, and only lets "good" boards in every respect into the field. If this were the case, as the manufacturing process improves, field failure levels would remain at a minimum attainable level, and improvement could be measured by looking at improvements in the percentage of boards that pass the complete in-house tests. The best (lowest) attainable failure level would be limited, however, by (1) the gaps in the test coverage, and (2) the fails that are recorded for illegitimate reasons. Examples of this include replacing the formatter even though it was never faulty, or a customer calling requesting repair for something unrelated to the product's function. For these reasons, it is important to do comprehensive tests of boards returned from the field. This analysis could determine the number of boards with (1) fail causes that fall outside of what the normal tests would catch, and (2) with no problems (either intermittent or hard fails). Knowing these two levels allows us to continuously improve testing procedures, predict what the field failure levels ideally could be on the current products, and make predictions about future products.

3.5.3 Testing Decisions

Successful tests must be designed for the particular types of failures to be detected and eliminated. Despite the variations on the bathtub curve discussed earlier, the common thread through all the theories is that the primary focus for latent defect testing must be for infant or early-life failures. Furthermore, based on the behavior of defects, this type of testing can only be effective if the failure modes exhibit a decreasing failure rate over time. According to Wong,

---

[2]On the other hand, though, there's limits to how much data is available, and also how much past data may be relevant to the current production process.

most of the data originally used to derive the bathtub curve was based on electron tube equipment. Therefore, different failure models could apply today since current products use newer, different technologies. For example, failure rates in today's electronic equipment are better characterized by a constantly decreasing failure rate and no period of constant failures.[3] For example, see Fig. 2.6 for the case of the formatter boards. The formatters show a constantly decreasing failure rate over time for the period studied. This supports the idea of latent defect testing as being beneficial. To further test this hypothesis, we should look at individual failure modes within the population to check that their behavior is similar to the formatter, as discussed in Section 3.5.1. Unfortunately, we did not have data at this level of detail from the formatters.

If we can simulate the stresses that the printers receive in their first few months of life, we may be able to force a majority of the warranty problems to occur in the factory rather than letting them reach the field. By contrast, the paper feed assembly in Fig. 2.7 is a strictly mechanical component which would not be expected to exhibit the same types of behavior as electronic products. This appears to be the case from our analysis. Fig. 2.8 looks at all components in the printer with the exception of the formatter; the flat slope of this failure behavior indicates that it would not be effective to do burn-in testing at the final printer level during manufacturing, especially since this constant failure behavior is indicative of random failures. For constant or increasing failures over time, we must conclude that stressing the boards would simply increase their likelihood of failing earlier in their field life, but have no beneficial impact on overall field quality levels. In conclusion, latent-type testing would be much more useful on formatters than for the entire printer.

3.6 Summary

Our goal in testing for latent defects is to catch defects in the plant rather than letting them reach customers. This means eliminating as many of the infant mortality failures as possible (or

[3]See (Wong, 1981) for examples of decreasing failure rates from product tests and data analysis.

economically feasible). Customers would thus receive products already in their useful life, where they are more reliable.

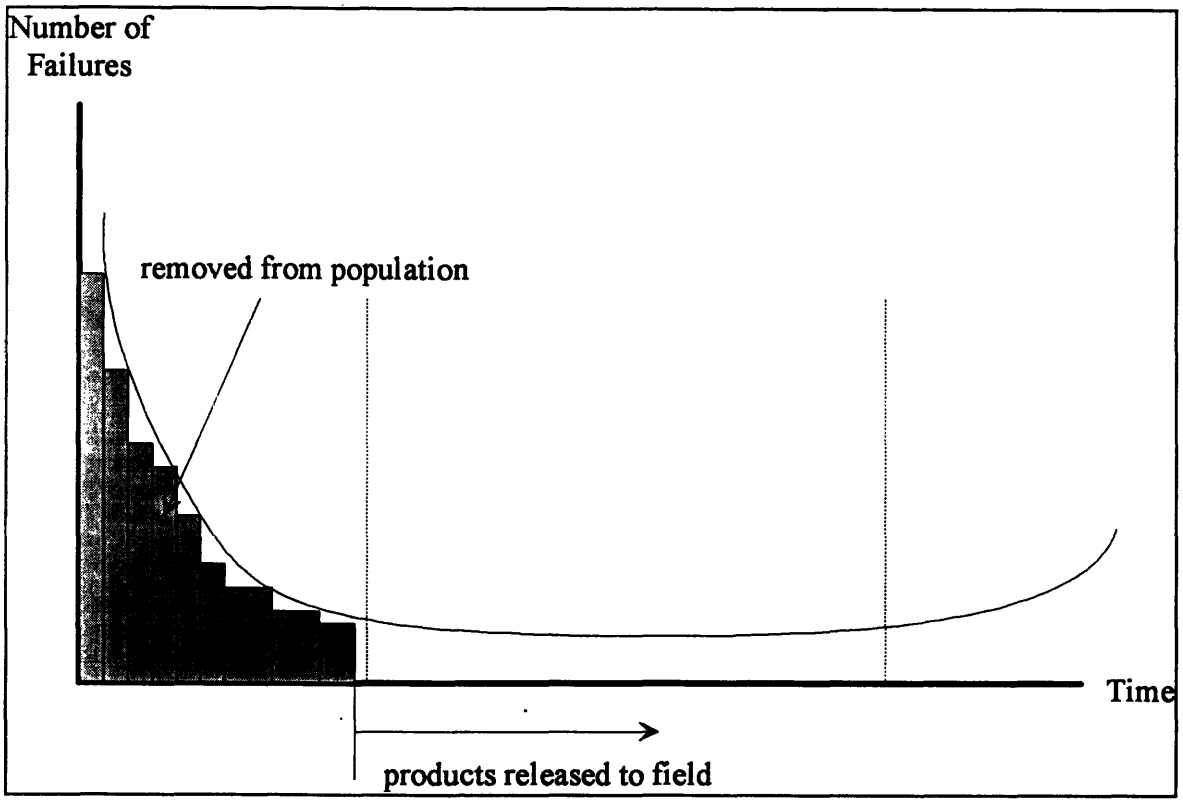

Figure 3.6 Effect of testing on products sent to customers.

The above graph displays the instantaneous failure rate at various points of time for the product. We can see that the effect of testing would be to shift the y-axis to the right by a length of time corresponding to the test which is performed. The shaded area under the curve represents the number of products we would expect to catch via the test.[18]

The better we understand the characteristics and failures for our products, the better we are able to design an effective testing method to precipitate and detect them. The following two figures summarize the types of information we have discussed in this chapter. Fig. 3.7 illustrates the information we need for understanding failure behavior of the product during its early, or

infant, stages of life. First, the *area* of the failure curve tells us the potential benefit of doing latent defect testing. The products under this curve represent infant mortality failures with specific causes that we should be able to recreate within the factory rather than letting the customer see them. Second, the *length* of the curve over time determines the amount of product life we need to simulate in order to detect these latent defects. This simulation can be done in real time or it can be accelerated, as will be discussed in the next chapter. As we implement a test and measure the results, we would expect to see a shape similar to the past failure behavior if our test is effective.[19] Third, the *composition* of the curve is critical. In most situations a failure curve is the sum of several different failure modes, each with its own unique behavior over time. As we study data, we need to understand each of these individual modes as well as the overall results. For example, as we saw for the laser printer in Fig. 2.8, failures occur evenly over time rather than decreasing. Unless this constant failure is the cumulative result of several decreasing failure rates for various subassemblies, latent defect testing at this level is not a beneficial measure.

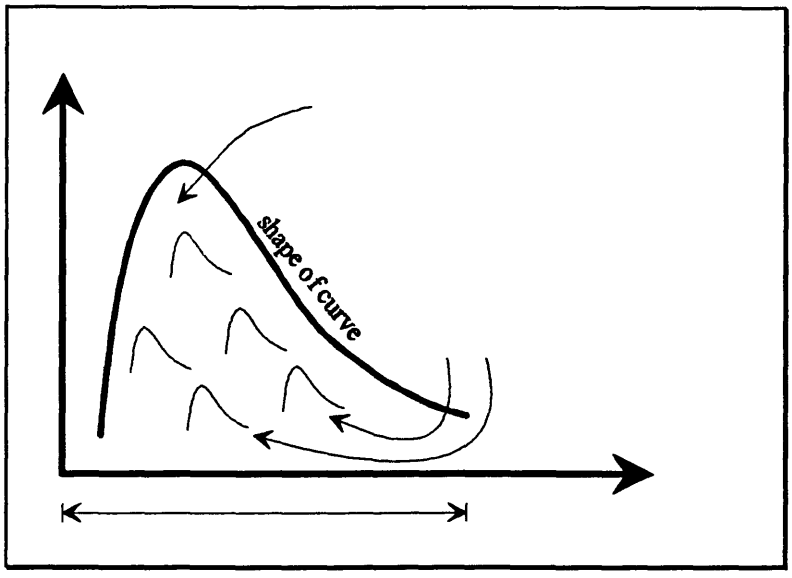

Figure 3.7 First dimension of time-dependent failure model.

The other two factors we must understand are pictured in Fig. 3.8. First we need a *foundation of understanding* about the product: the number of products being produced, the

characteristics of the final product, the past history and success of the product and process technologies which it uses, etc. As detailed information is gathered on failures, knowing this background information helps us to see consistencies in what would otherwise appear to be unrelated data. We also need to be aware of the *changing needs* for test due to the changing nature of defect modes. As we perform testing over time, we analyze and eliminate root causes of classes of defects, causing a reduction in the future levels and types of failures in the field. Taken to the extreme, it is possible for a latent defect test to negate its own effectiveness. This can happen if the test finds all of the failure modes for which it is designed, and the factory implements successful preventive measures. However, in real life, new failure modes appear as a result of process changes, new vendors, etc. For this reason, a comprehensive test can be extremely valuable in detecting surprises before they reach final products.

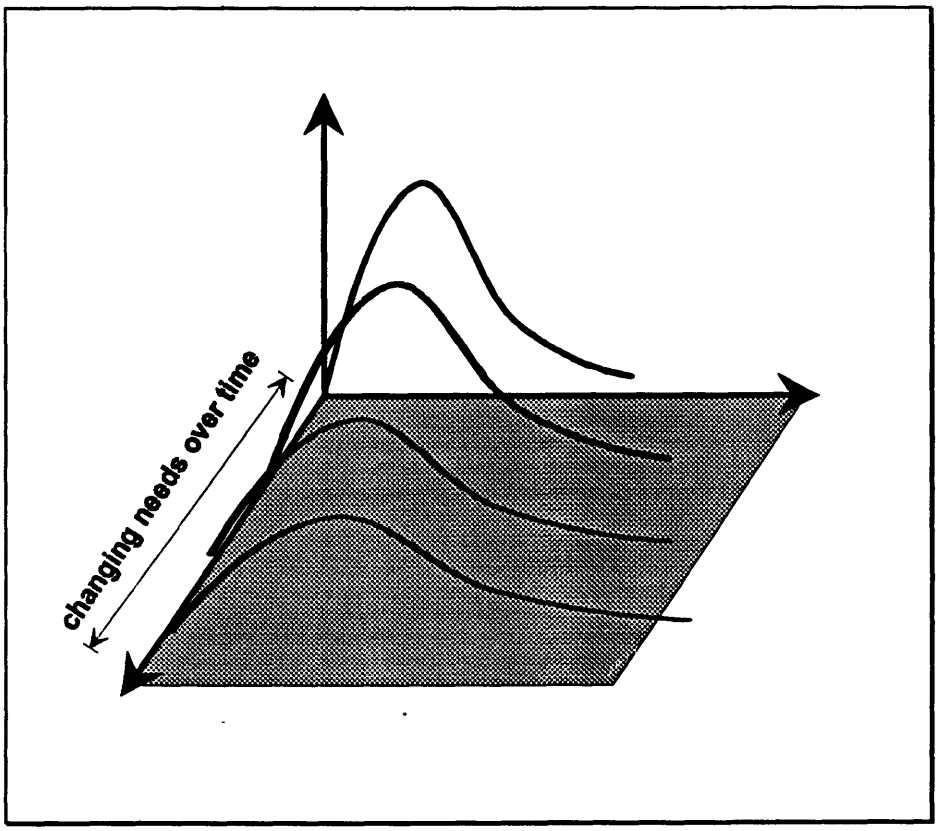

Figure 3.8 Second and third dimensions of time-dependent failure model.

In conclusion, we learn about these dimensions of failure behavior by collecting data and using available models or research to help us better understand why and how defects are formed and grow into functional failures. The bathtub curve serves as a useful framework for discussing the various aspects of product life and latent defects. Research has served to clarify and update aspects of this model, and has also given us some insight into specific applications. To apply the models we need data which is consistent, specific, and complete. Although the models lend themselves to detailed statistical analyses, we often use them also as qualitative tools for understanding the data we *do* have available, since our available data may not allow us to go into more detail with confidence.

So far, we have considered the two bodies of information which enable us to understand the *problems* related to latent defects: data and theory. Our next task is to develop a feel for the possible *solutions* which can help us eliminate latent defects in our products. This will be the focus of the next chapter.

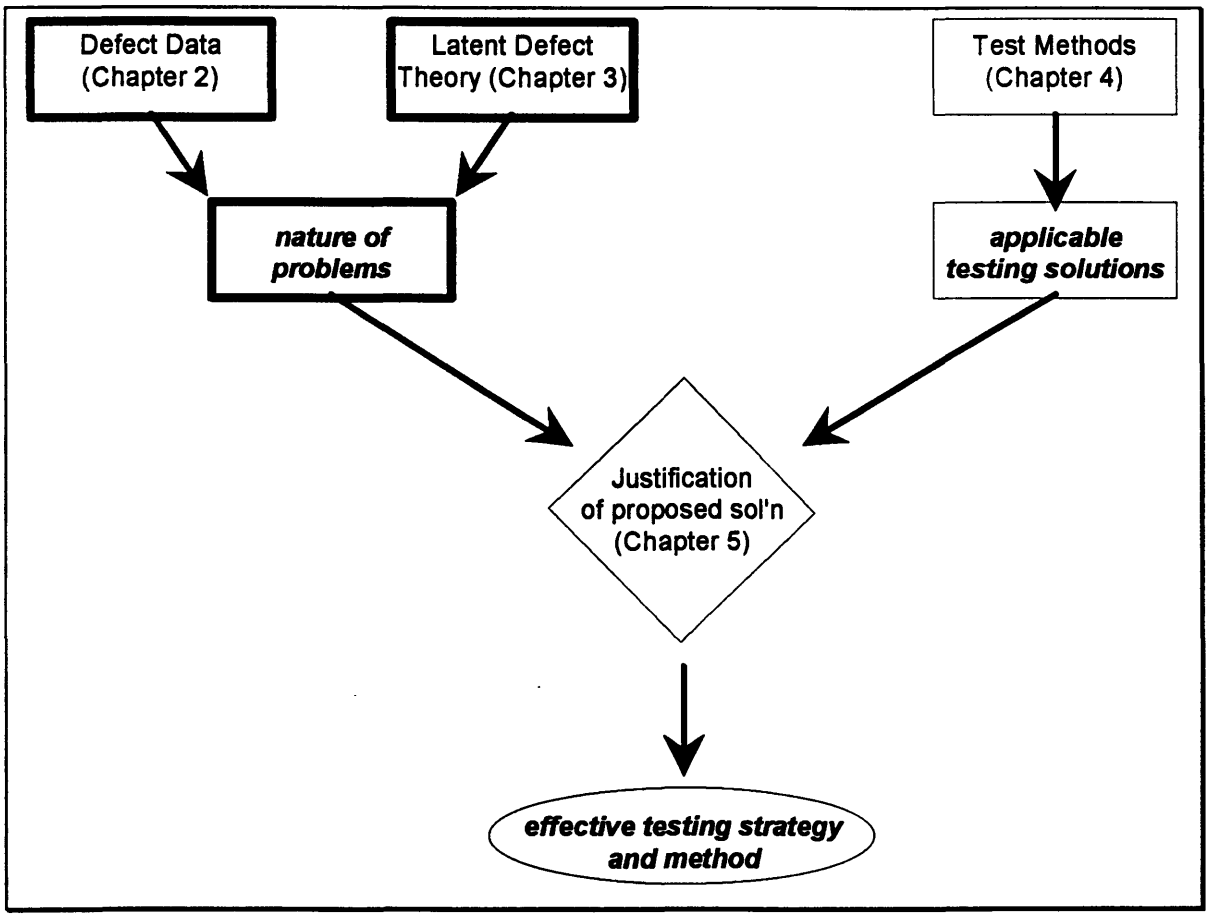

Figure 3.9 Review of latent defect testing decision process.

References

[1] Pynn, Craig, Strategies for Electronics Test, McGraw-Hill, 1986, p. 12.

[2] Wong, Kam L., "Off the Bathtub Onto the Roller-Coaster Curve," Proceedings of the Annual Reliability and Maintainability Symposium, 1988, pp. 356-363.

[3] Wolf, Stanley, Silicon Processing for the VLSI Era, Volume 2: Process Integration, Lattice Press, 1990, p. 264.

[4] Ibid., p. 113.

[5] Hu, J.M., D. Barker, A. Dasgupta, and A. Arora, "Role of Failure-Mechanism Identification in Accelerated Testing," Journal of the Institute of Environmental Sciences, July/August 1993, pp. 39-45.

[6] Crevecoeur, G. U., "A Model for the Integrity Assessment of Aging Repairable Systems," IEEE Transactions on Reliability, March 1993, pp. 148-155.

[7] Hu, p. 40, and Nelson, Wayne, Accelerated Testing: Statistical Models, Test Plans, and Data Analysis, John Wiley and Sons, 1990, p. 12.

[8] Jensen, Finn, and Niels Petersen, Burn-in: An Engineering Approach to the Design and Analysis of Burn-in Procedures, John Wiley & Sons, Chichester, England, 1985.

[9] Pynn, p. 128.

[10] Pynn, p. 129.

[11] Musolino, Larry S., and Timothy Conrad, "Early Life Reliability Monitoring of AT&T-C&ES Manufactured Electronic Components," 1987 International Test Conference, IEEE, 1987, p. 446.

[12] Wong, Kam L., "Unified Field (Failure) Theory -- Demise of the Bathtub Curve", Proceedings of the Annual Reliability and Maintainability Symposium, 1981, pp. 402-407.

[13] Ibid., p. 404.

[14] Ibid., p. 405.

[15] Trinidade, David C., "Can Burn-in Screen Wearout Mechanisms?: Reliability Modeling of Defective Subpopulations - A Case Study," 29th Annual Proceedings, IEEE Reliability Physics Symposium, pp. 260-263.

[16] Jensen, p. 18.

[17] Wong, Kam L., "Off the Bathtub Onto the Roller-Coaster Curve," 1988, pp. 356-363.

[18] Huston, Hance H., et. al., "Burn-in Effectiveness--Theory and Measurement," 1991 IEEE Reliability Physics Symposium, pp. 271-276.

[19] Jensen, p. 115.

Chapter 4 - Methods for Latent Defect Testing

4.1 Introduction

There are basically two types of defects which are sources of product failure: those which directly and consistently affect product function, and those which over time will grow in severity until they affect the function. We refer to these types of defects respectively as *time-independent* or *time-dependent defects*. The tests described in Chapter 2 for formatter boards at H-P help us understand latent defects, but each test in itself is designed to detect time-independent failures. The in-circuit and functional tests check the functionality of the board at that time, and if the board passes the tests, the board is sent on to the next process step. In this chapter, we discuss tests which can be used to precipitate and detect latent defects.

Latent defect tests account for the possibility that although a product may be fully functional at one time, it may be likely to fail when tested or used in the future, due to all the reasons for latent defects presented in Chapter 3. We will discuss these test methods in order of "severity," based on the amount of stress they may put on a product. However, most of the topics or specific information presented for each of test will apply across all of the tests being presented.

4.2 Visual Inspection

Although the literature does not explicitly address visual inspection as a method for latent defect testing, it can be treated as such. During visual inspection, human operators or machine vision equipment inspect products with the intent of finding visible defects which would not be otherwise detected. The two extremes of problems which would be observable on sight include *cosmetic defects* in the board (which will never impact functionality), or *fatal defects* which make it impossible for the board to function, such as missing components or large cracks through circuit traces. In between these two extremes are latent defects: parts or systems which are likely to fail with wear or fatigue during the product life, even though the product may be fully functional at

the time of the inspection. Examples of defects which can be visually detected include the following:

--insufficient solder (sign of potential open circuit)
--excess solder (sign of potential short circuit)
--raised or loose connectors
--misalignment or lack of components needed for downstream stages of assembly (screws, bail clips, connectors, etc...).

In most testing for latent defects, there are two steps that occur: (1) the *precipitation* of latent defects--exciting them or stressing them to the point that they actually interfere with product function, and (2) the *detection* of these failures via functional checks. Visual inspection, however, does not follow this procedure. It works by directly detecting the sources of future failures without their ever being precipitated or excited. For this reason, it is difficult to gauge the effectiveness of this type of testing, since it may never be known if a variation that "looks" like a potential problem would actually develop into a true defect in the field. To develop effective visual tests, then, it is helpful to experiment by "seeding" products with visual defects and stressing these products to the point of failure. Another difficulty with manual visual inspection is that it is difficult to consistently check all features on every board as they come down assembly lines operating with fast cycle times, and to accurately sort out all of and only the boards that truly have a potential problem. Finally, operators can easily become fatigued doing this type of work, making the inspection less effective. Automated vision systems help in improving consistency and can obviously work much longer than a human observer, but these systems are more expensive and require a large amount of initial set-up for training in what to watch for and how.[1]

4.3 Burn-In

In visual inspection, no effort is made to precipitate latent defects so that they actually affect the functioning of the product. On the other hand, burn-in testing makes an effort to cause the defects to develop into failures. The definition of burn-in testing varies from company to company. For the purpose of discussion here, we will define it to be the operation of a product at

field use and stress levels (of power, temperature, vibration, etc.), with the intent of precipitating defects which otherwise would have occurred at a customer site.[1] When designing this sort of test, we must determine the types of defects to be detected and when we expect these defects to reveal themselves over time. Given that information , we can set up the test so that it is capable of precipitating and detecting our targeted problems. For example, power cycling occurs normally on customer products, and it creates stresses which cause product failures. If we intend to detect a failure mode which is excited by this type of stress, we could apply 20 power cycles to each board, during which the board "fully" heats up and cools down during each cycle. This would correspond to 5 weeks of use in the field where someone turns on his printer each workday and turns it off before he goes home.

For burn-in as well as the other test methods in this chapter, we must make a decision as to the necessary level of monitoring we need for observing the performance of the product. We can think of the possibilities as lying along the continuum shown in Fig. 4.1.

operation of product with no function checks	⇒	monitoring of output only	⇒	continuous product testing & recording of results

Figure 4.1 Possibilities for levels of monitoring of burn-in testing.

As we move to the right, we receive more information about the product in a given amount of time, but the test equipment for gathering this data becomes increasingly complex and expensive.[2] In the next two sections, we will distinguish between the types of test at each end of this spectrum.

4.3.1 Passive vs. Monitored Burn-in

We define *passive burn-in* or non-monitored testing (at the left end of Fig. 4.1) as testing in which the status of the product is not checked until the burn-in period is complete. One

[1]Others such as Pynn define burn-in to also include some sort of stress on the product in order to speed up the process of accelerating defects. Since we will discuss the use of stress in greater detail in the next section on accelerated stress testing, we will limit the topics here to ambient conditions or relatively minor acceleration.

primary disadvantage of non-monitored testing is the possibility that defects may be precipitated but not detected. For example, putting boards through eight hours of power cycling without fully checking their performance afterwards would only help boards to fail at a downstream process sooner than they would have otherwise. At the very least, a full functional test is necessary after burn-in.

When performing non-monitored testing, we hope that all product malfunctions would be detectable by the final burn-in test. Unfortunately, though, there are at least three basic types of defect behaviors which may not be detected in this way:

(1) intermittent -- the product works sometimes but fails other times, somewhat randomly, and thus may not fail when the single test is applied to it at the end of the burn-in

(2) healer -- the product fails constantly or intermittently at a particular stress level, but then begins working normally at all temperatures for an extended period of time

(3) constant -- the product fails consistently at some point outside the average operating conditions, but still within its designed expectations [3]

On the other hand, *monitored burn-in* periodically or continuously checks the status of the product as power or other stresses are being applied. These tests can be run using the bed-of-nails testers described earlier, or they can also be transferred to the boards via plug-in cards known as Single In-line Memory Modules, or SIMMs, into expansion slots on the board. Another possibility is to input test programs through the serial or parallel connections which normally link the formatter board to the printer. Although these processes requires more elaborate setups and a more detailed data monitoring and compilation system, many times more data can be gathered than with a single functional check at the end of the test. Another advantage to monitored burn-in is that there is no need for a post test, or the time for final test would at least be greatly reduced if the tests have already been run via the monitoring. This can save on test time and on the amount of necessary handling of the product.[4]

4.3.2 Examples of Burn-In Testing

Much of the learning about the benefits and challenges of burn-in testing has taken place during actual trials or implementation of burn-in procedures. The experiences and methods of several manufacturers with burn-in testing are described below.

(1) Laser printer assembly. One example of testing which fits our definition of burn-in is the testing process used on one of the production printer models at H-P in Boise, Idaho. Once these printers are completely assembled, they are put through a test cycle which lasts from ten to thirty minutes. These tests check the functions of the printer via addressing tests, memory tests, and actual printing of test pages from the printer. These test pages are visually inspected by the test operator. One failure mode which had been a problem in the past was related to weak power supplies. To screen for these, test engineers added a power cycling procedure to the test cycle. Each printer is turned on and off three times each for 0.5 seconds, one second, and 1.5 seconds for a total of nine power cycles, in addition to two other power cycles during the process. This test has been quite effective in isolating the marginal power supplies before they leave the factory.

(2) Intermittent failures in formatter boards. During tests on a recent H-P formatter board design, engineers placed boards in test chambers and continually monitored them as various stresses were placed on them. Dozens of failures and failed tests were observed as the tests were in progress, but when the boards were removed from the test chamber and tested in a normal use environment, none of the failures could be duplicated. Since the test equipment was closely inspected and found to still be fully functional, the conclusion was made that a large portion of these failures were intermittent in nature and therefore difficult to analyze in any detail. Intermittence is an important type of latent defect, but it is also extremely difficult to analyze when it occurs.

(3) Short-time effectiveness of burn-in. On a past workstation product, H-P found that burn-in was effective in detecting many potential power supply failures. Once this power supply was redesigned based on the test findings, the need for burn-in was negated. This is an important consideration in justifying the long-term implementation of a burn-in process. It is common that burn-in can prove itself effective in the short term, but resulting design improvements have long-term benefits which result in reducing the effectiveness of the burn-in.

(4) Continuous evaluation of test results. One final example of burn-in at AT&T illustrates the importance of constantly analyzing the results of the testing, as opposed to simply designing the test and letting it run without periodic reviews or monitoring. The product under examination here was a high-density, switch mode rectifier product, designed for use in a battery-backed power system. For the burn-in process, each product was put through continuous tests of various subsystems (controllers, individual rectifiers, display panel, etc.) and the results were plotted and reviewed on a monthly basis. When evaluating the effectiveness of the burn-in, engineers looked at three criteria. First, the

quantity of failures was examined. If there was a substantial number of failed products, they increased the test time to ensure product quality. Next, the *mean-time-to-fail* was examined. This was seen as the best single indicator of necessary burn-in time. Finally, any *late failures* received special attention. These were defined as failures which occurred within the last 25% of the burn-in time. For example, some failures can only be detected at the end of testing. A meter failure was an example of such a failure, since it would only be detected when the operator ran the final test sequence. In a case such as this, further study was done to determine when the failures occur. Before AT&T implemented these procedures, the field reliability of these products was predicted to be around 100,000 hours MTBF; after a year of doing burn-in, the reliability was approaching 1,000,000 hours MTBF.[5]

4.3.3 Design of Formatter Experiment

In Chapter 2, we discussed that a change was in progress regarding how the formatter boards were assembled into printers and tested as part of final printer assembly. The "old" process at Canon's facilities involved a longer test time and over twice as many power cycles as the "new" process at the H-P integration sites. Based on our available information on defect rates for the assembly procedures, our hypothesis was that *the decrease in power cycling caused less defects to be precipitated in the factory, in turn causing more valid latent defects to occur in the field.* In order to better understand the impact that this process change could have on final product quality, we designed a burn-in test for formatter boards. To test our hypothesis, we planned to duplicate at least the difference between the two final assembly procedures, retest the boards, and determine from the resulting failures how important this test time could be in terms of final product quality. We also hoped to better understand if the board failures at printer assembly are due to the power cycling that they receive, or if other factors such as board transportation or assembly procedures may be playing a role. Performing the burn-in as a controlled experiment would help us to isolate these factors from each other.

Representatives from the formatter board test group, the board analysis group, and the formatter factory met and discussed the following criteria for the test:

(1) Type of stresses. We agreed to limit the test procedure on each board to strictly power cycling. This was the best simulation of the final printer test processes that we were studying, and it was also relatively easy to design fixtures and cabling to make this

possible and reliable for large quantities of boards. Given that we were unsure how more severe stresses would affect the long-term quality of our particular final product, we agreed that power cycling would be a "safe" test for now.

(2) Passive vs. monitored burn-in. We would have learned the most from a monitored burn-in experiment, therefore we would have preferred to take that approach. However, obtaining the necessary equipment and designing the customized hardware and software required more time that was available during the internship. As a result, we decided to proceed with passive testing for the first phase of the testing in order to gather initial data. We developed the monitored testing capability as a second phase of the test, as time was available to work on it. (Note: For the rest of this experiment, *Phase 1* will refer to the passive burn-in testing, and *Phase 2* denotes the planned monitored testing.) We designed Phase 2 to duplicate Phase 1 as closely as possible, with the only exception being that additional connections would be needed for the monitored test programs. This would make it valid to compare data across both phases.

In addition to the passive burn-in, we performed visual checks on the boards as we tested them. When reviewing the boards, operators checked for the same types of problems that are detected in the regular factory visual check procedures: insufficient or excess solder, loose connectors, misaligned parts, etc.

(3) Duration of test. Comparing the processes for the two final assembly procedures, we found a difference between the two tests of ten to fifteen power cycles, therefore we wanted to do at least that many power cycles on each test. We agreed to test the boards for a total of twenty-four power cycles, where each power cycle consisted of power (+5V DC) being applied for five minutes, and then zero volts for another five minutes. This was designed to give each board time to fully warm up and go through whatever internal start-up tests it could, and to fully cool down again before repeating the process.

4) Quantity of boards to be tested. Based on information regarding failure rates from the assembly sites, we estimated that a power cycling test procedure should be expected to detect failures at a maximum rate of about 1,500 ppm, or one out of every 667 boards.[2] As a result, several thousand boards would need to be tested in order to have any confidence in the results. We therefore set ourselves a goal of power cycling and testing at least 6,000 boards in order to obtain some meaningful data.

At this point, we also agreed on what sort of conclusions would be reasonable to make based on the outcomes of the testing. We felt that it was important to lay out some assumptions, before testing began, rather than waiting until the data was taken and letting our conclusions be affected by knowing the actual results. We found that a failure rate of five to twelve boards out of 6,000 would allow us to state with 90% confidence that 1,500

[2]Note in Appendix A that the total ppm rates for boards at the assembly sites is significantly higher than 1,500 ppm. However, some of these types of failures could not be expected to be precipitated with a power cycling test. For example, most of the "mechanical" type failures would be excited by other types of stresses (vibration, heat, etc...) than by regular electrical powering, therefore the estimated results of the test were set as a fraction of the total electrical defects previously observed at assembly sites.

ppm could be the population defect rate for the types of defects we would be able to detect. We proposed three conclusions we could make based on the test results:

Failed boards in 6,000 board sample	Conclusion and action plan
0-4	Test is not strenuous enough, or 1,500 ppm fail estimate is high; *revise test to gain more information* (or discontinue), *review data results to date, and review phase 2 testing for the possibility of using additional stresses*
5-12	1,500 ppm is within 90% confidence bounds; test verifies our expectations of the number of latent defects we should be able to detect--*OK to discontinue test, or to revise test to gather different data if desired; testing is effective in detecting latent defects .:. determine if this test method is feasible and justifiable in production*
13 or more	Failure rate is above 1,500 ppm with at least 90% confidence; *stop test & review failed boards and determine if test may be excessive* (unlikely), *if outside factors are causing defects (handling, ESD, etc...), or if the failed boards represent class failure modes which need to be addressed and resolved; postpone start of Phase 2 testing until this is resolved*

5) Test fixturing and layout. Based on current and voltage requirements and power supply limitations, we found that we could provide power to one hundred formatter boards at a time. Three connections would need to be made to each board: power (+5V DC), ground, and a reset signal, which allows the board to begin running its startup procedure.

We connected the signals to the board via the 26-pin connection which normally mates the board to the print engine. One additional ground connection also had to be made to the serial connection at the back of the board, in order to provide a consistent ground to the entire board.[3] To provide these connections, we built "connector tails" which consisted of the following connector components which would mate with the engine connection on the board: a "handle" for the connector made out of breadboard to provide strain relief for the wires and to give the operator a place to grasp the connector, six wires (two for +5V, three for ground, one for the reset signal), and terminals to allow it to be attached to terminal blocks on the fixture. Fig. 4.2 illustrates how these connections were made to each board.

[3]The ground plane for the board was broken into two sections to minimize the interference between data communication lines and other signals on the board.

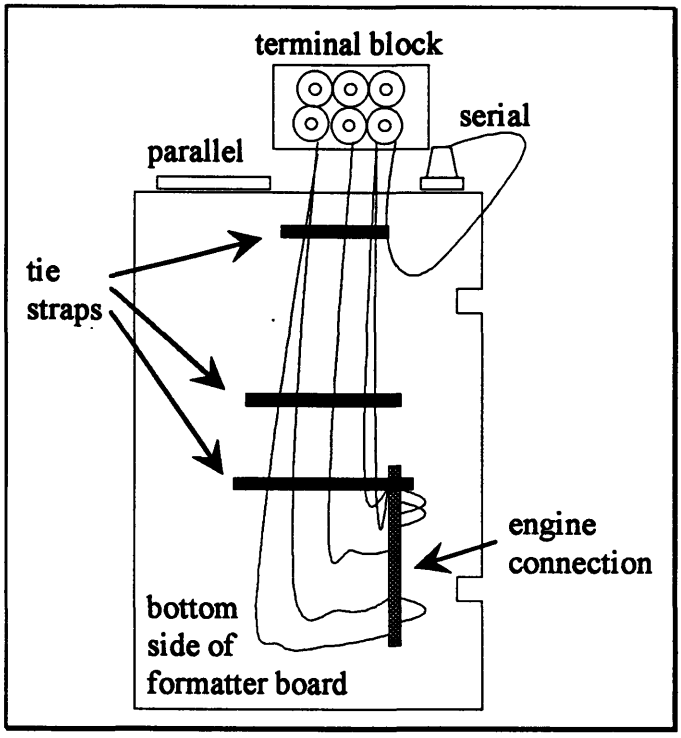

Figure 4.2 Layout of board connections.

We attached the terminal blocks to the back of shelving units fitted with ESD foam and connections to protect against electrostatic shocks. Four shelving units contained five shelves each, and each shelf contained five boards laid out as shown below.

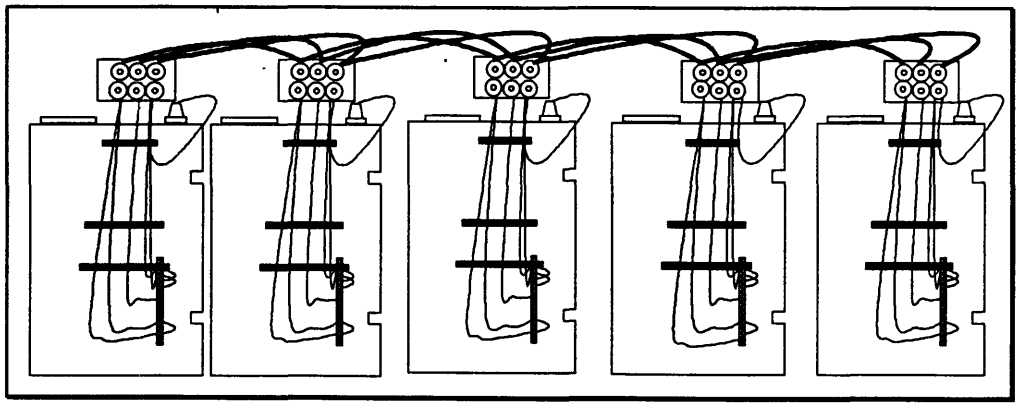

Figure 4.3 Layout of boards in test fixture for Phase 1.

We connected the wires for each shelf to the power supply (for the +5V and ground wires) or a reset board capable of sending a signal to all one hundred boards at once. Finally, we connected the power supply back to a UNIX workstation programmed to turn the power supply on and off at the desired intervals.

The last part of the setup consisted of two functional testers, which used a bed of nails to perform all the normal production tests on the boards. We programmed the testers and the workstation so that an operator could insert two boards, close the testers, and the test programs would automatically run. After the test, the screen on the workstation would notify the operator if the boards passed or failed.

6) Test procedure. Operators from the board assembly factory worked in shifts around the clock to test the boards. The following procedure was used:

1) get completed boards from the factory
2) place boards face down in racks, five to a shelf (100 boards total), and plug in engine and ground connectors
3) start power cycling program on computer

WAIT FOR COMPLETION OF POWER CYCLING

4) at end of power cycling, brush off functional tester (to avoid poor contacts), remove each board, visually inspect the board, and test using the two functional testers
5) repack all boards which "pass" functional test, put all failed boards in designated containers for later analysis
6) return repacked boards to line

Each day, we retested any boards that failed test during the previous day. If the board passed the functional test this second time, we held it for a final check with the Phase 2 testing. If the board failed repeatedly, experienced technicians analyzed the board to determine a root cause of the failed tests. If the problem could be narrowed down to a specific component, that component would be replaced after viewing the board under a microscope to ensure that the failure was actually inside the component, as opposed to a board-level defect such as insufficient or excess solder. Once a technician replaced the component, we retested the board on the functional tester to guarantee that it now works, and component engineers shipped the component to its vendor for more detailed analysis.

7) Disposition of tested boards. We repackaged all boards which passed the first functional test and shipped them to printer assembly sites. We kept all boards which failed their first functional test or failed visual inspection for later analysis. The serial numbers of all tested boards, good and bad, were also recorded in a database for later access. The database tracked the time of the test, the result of the test, and the particular test at which a failure was recorded.

4.3.4 Results of Formatter Experiment

After approximately one month of testing, we obtained the intermediate results shown in Table 4.1. The first category of defects, "functional fail - visual OK," denotes boards which failed

the functional test after the power cycling, but which had no visual defects. "Functional OK - visual fail" boards failed the operators' visual inspection, but still passed the functional test. The one board which was a "functional fail - visual fail" was a special case. The board consistently failed the functional test, and on further inspection, we noticed that a resistor pack component was missing on the board. When we checked the data record for the board, we found that it had been logged as a failed board in the factory. However, it had never been repaired and ended up being accidentally shipped. Therefore, although we could not credit the power cycling with catching this particular failure, it was still an interesting event from a visual inspection and quality standpoint.

The last two categories consist of boards which failed the functional test the first time, and also failed at least one other time, making it much more likely that there was a legitimate latent defect or problem in the board. The "semirepeatable" boards would fail the test several times, but before a specific component or failure cause could be analyzed, the board began working again and the failure could not be duplicated. The "repeatable" failures consistently failed each functional test so that a component could be determined as being the source of the problem. At that point, the component was replaced, and the board would again pass the tests.

	Number	%	ppm rate
Boards tested	5,700		
Functional fail - visual OK	53	0.93	9,300
Functional OK - visual fail	9	0.16	1,600
Functional fail - visual fail	1	0.02	175
Semi-repeatable failures (>1)	2	0.04	350
Repeatable failures	3	0.05	525
Reasons for visual fails			
insufficient solder	2		
missing screw	2		
solder balls & filled via	2		
bent leg	1		
excess solder	1		

exposed copper/cuts on board	1
missing resistor pack & pad	1
raised connector	1
	11 defects on 10 boards
Reasons for repeatable failures	
bad ROM	1
bad DRAM	1
missing resistor pack & pad	1 (note: this board never did pass tests in the plant)
	3

Table 4.1 Intermediate results of passive burn-in testing.

From this information, we determined that it was beneficial to continue the Phase 1 testing. The testing had found four boards (semirepeatable and repeatable) with definite weaknesses resulting from the testing, and some portion of the "functional fail - visual OK" boards were likely sources of intermittent problems which our equipment was not able to further analyze. From here, we continued testing boards until the end of the author's internship; the final totals at that point were as follows:

	Number	%	ppm rate
Boards tested	24,060		
Functional fail - visual OK	111	0.46	4,610
Functional OK - visual fail	31	0.13	1,290
Functional fail - visual fail	1	0.004	42
Semi-repeatable failures (>1)	3	0.01	125
Repeatable failures	13	0.05	540

Table 4.2 Final results of passive burn-in testing.

One of our conclusions from this testing is that the controlled power cycling performed as expected, based on the available data from integration sites. It was able to detect failure modes similar to what has been seen in downstream steps in the past. Unfortunately, even though our sample size is relatively large, the overall high quality of the products made it impossible to collect enough data to draw conclusions about individual failure modes, or to perform traditional statistical analyses, such as Weibull distributions, of the data with a significant level of confidence.

The following chapters of this thesis present further conclusions related to implementation and justification of this sort of testing.

In addition, the types of defects we observed were consistent with the types of defects normally occurring at downstream processes. With only thirteen repeatable failures, we could not compare magnitudes of each failure type with any confidence, but the defective components--DRAMs, microprocessors, ASICs--matched up with the most common failure modes out of printer assembly. One notable exception, however, was the oscillator. Oscillator malfunctions or timing problems are often intermittent and are also difficult to isolate. The timing mismatches they generate make it appear as if other components are at fault. The printer-level tests are able to look at all the results of the formatter's operations at once. Therefore, it is possible to isolate the oscillator as being the source of the problem.

Our testing also provided strong (albeit frustrating!) evidence of the potential confusion that intermittent failures can cause, and also reinforced the importance of being able to monitor the state of the products as they are being tested. Dozens of boards which failed the functional testing once, but passed it the second time. If we knew how many of these boards functioned properly during the entire four hours of power cycling, it would be possible to gauge the reliability of the testers with much more confidence. In addition, the semirepeatable failures also represent a lack of complete information as well. Monitored testing could have potentially provided enough information to determine specific root causes for these intermittent failures. Finally, none of the components isolated as the sources of defects on the "repeatable fail" boards failed in component-level testing by their vendors. Of all the information received back from vendors to date, nobody was able to provide additional information regarding the source of the failure.

4.4 Accelerated Stress Testing

A major drawback of burn-in testing is the time that it requires. For example, the Phase 1 experiment of power cycling each board 24 times can be considered to be roughly equivalent to a

user turning their printer on and off once a day for a month. In some respects, we can equate a month in the user's environment to four hours in a controlled test at the factory. There are clearly some benefits to this, but if we simulate even more of the product's life in a shorter period of time, we could gain information more quickly, and also have time to test for a wider variety of failure modes. This concept of speeding up the product's progress with respect to its failure behavior, is known as *accelerated stress testing*. The goal of this testing is to accelerate latent defects so that they occur in the plant rather than at the customer.[6] In this section, we discuss the types of stresses used in this sort of testing and the effects they have on different types of failure modes. We also discuss the potential uses for accelerated testing in production, and review the experience of various companies and divisions in their work in this field.

4.4.1 Types of Stresses

As was shown in Table 3.1, there are a wide variety of potential root causes or conditions which can result in an actual defect. In addition, for every type of basic failure condition or imperfection, there is some sort of stress which excites that failure and causes it to grow. For example, some of the basic types of stresses inherent in the operation of an electronic product are the following:

--	voltage	--	current
--	heat	--	moisture
--	temperature	--	power
--	shock and vibration	--	timing variations

Each of these stresses can excite different types of failure mechanisms, such as electromigration, differential expansion, or loosening of connectors.

A critical consideration in designing an accelerated test procedure is understanding what sorts of stresses or combinations of stresses excite the failure modes of interest. Table 4.3 provides a fairly comprehensive listing of the general classes of stresses, what sorts of failure mechanisms they can trigger, and what types of components are most susceptible to these sorts of stresses.[7]

Stress Type	Failure mechanisms	Component Types
high voltage	dielectric breakdown	capacitors oxide layers
high current	electromigration fusing(melting)	thin metal films conductors
heat (can be caused by continuous power)	chemical reaction intermetallics ionic migration	batteries electrolytic capacitors interconnections semiconductors (NVRAMs, DRAMs, etc...)
moisture	seal leakage dendrite growth corrosion ionic migration	polymer seals metals (Ag, Cu, Sn) thin metal films semiconductors insulators
temperature cycles	differential expansion condensation	polymer encapsulation wire bonds die attach cavity packages high voltage circuits
power cycles	power supply failure enlarged temperature cycles	capacitors diodes triacs & SCRs oscillators
shock and vibration	loosening fatigue intermittents	fasteners interconnections wire bonds die attach cavity packages
timing	digital race intermittents	digital logic

Table 4.3 Correlation of stresses to failure modes.

4.4.2 Methods and Effects of Acceleration and Defect Precipitation

Here are two basic rules which we need to follow as we look to accelerated testing techniques:

(1) "the failure modes observed in the accelerated environment must be the same as those observed under conditions of use, and
(2) it should be possible with a reasonable degree of assurance to extrapolate from the accelerated environment to the conditions of use."[8]

If we combine this with Trinidade's view of mortal and immortal portions of the product population, then we can think of exposing defects with accelerated stresses as shown in Fig. 4.4.

Our challenge is to determine what stress levels will expose the mortal boards without significantly degrading the immortals.

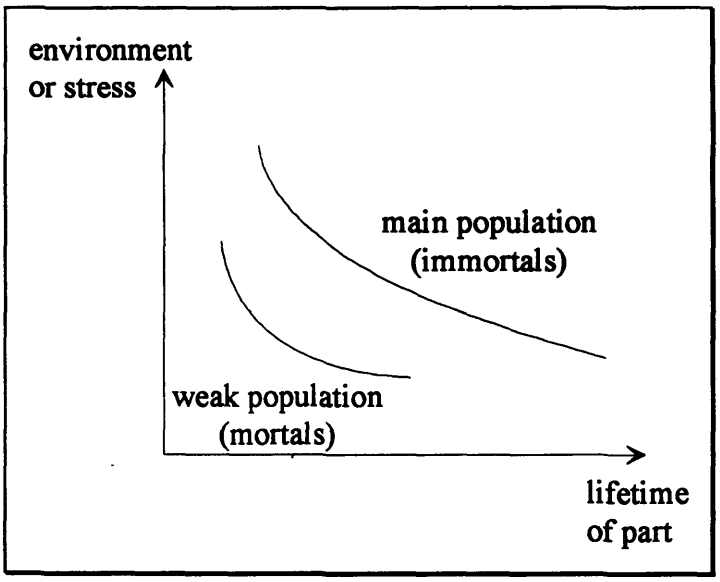

Figure 4.4 Model for effects of accelerated stress on lifetime of product population.

According to Hu et al., when designing a test, it is important to identify the critical stresses for a product at its potential failure sites. If we choose a stress to apply during testing, we must be sure to understand all the failure mechanisms which that type of stress may trigger. During an actual test, we could excite all of these mechanisms, some of which may never have a chance of occurring in the field. Our testing results (and discovered failure mechanisms) will not correspond to what occurs in the field. For example, certain types of defects can be quantified in terms of the amount of *activation energy* necessary to excite that particular failure mode, as measured in electron volts. As a result, accelerating a test by increasing the applied activation energy may not cause a desired failure mode to occur any faster, but it may instead excite completely different failure modes with higher activation energies. Temperature stresses can also lead to *changes in material properties*, which allow new failure mechanisms to develop as well.[9]

On the other hand, Hobbs and others believe that the particular stresses used to precipitate failures do not have to correspond to real conditions in the user environment. Hobbs writes that,

"product responses to the stimuli are the crucial factor; the input is not."[10] When stress screening electronics parts, he finds that nearly every failure is able to be correlated to a current or future field condition, even if the failure was generated by taking a part far beyond its design specification limits. If a component breaks, the fact that it broke before another part of the product indicates that it is the weakest link and most likely to break in the field at some point.[11] We believe that this general approach can be misleading if the test designer decides to apply any stress which generated a failure, without understanding the effects of the test he has designed. In addition, he will find it much more difficult to convince his coworkers of his findings if he cannot explain why the stress test is valid.

A wise approach to acceleration methods would take in aspects of both of Hu's and Hobbs' views. We should use available information on current failure modes to determine what stresses will be useful to run, but also experiment with new combinations of stresses or acceleration rates in order to learn about unique failure modes which customers may not have observed yet. For example, one technique often used to excite valid failures is *temperature ramping*, moving very quickly from one temperature extreme to another. Although fast temperature changes seldom occur in real life for many products, this method drastically accelerates the effects normally caused by power cycling or product use. If a test only measures product performance at constant temperatures, failure modes such as intermittent opens and shorts due to dissimilar metals or materials pulling away from each other may never be noticed. This is because once an extreme temperature is reached, the materials have time to settle back down "into position" again and the product may resume functioning. To make things more complicated, different materials react differently to thermal cycling. For brittle materials, a higher rate of change would accelerate failures, but for ductile materials like aluminum and solder, the effect is just the opposite.[12] Therefore, it is important to understand to some extent what types of failure modes we *want* to catch by using temperature ramping, but at the same time, important failure modes may appear during the testing process which we had not previously considered.

In addition to these views of accelerated testing, there are a variety of mathematical models or formulae we can use to describe specific failure phenomena. Once we find an appropriate model for defect formation, we can use it to compute an acceleration factor for relating stresses at various temperatures or vibration levels to each other, or to life under normal use conditions.[13] We must be careful, however, when using a single model to describe a product with many different failure modes. Some of the failure modes may behave differently than others under different stresses or operating conditions. A given model may work well for one set of conditions but be very inaccurate under different stresses. In the remainder of this section, we will look at three of the more general models and describe their potential use in depicting failure behavior.

(1) Arrhenius model (temperature). We can use the Arrhenius equation to model failure modes which have as their base process some type of degradation. The model calculates the hazard rate for a given failure process and a given temperature:[14]

$$\lambda_b = K \exp(-E / kT)$$

where: λ_b = the base hazard rate, or the instantaneous probability of the first and only failure for the product
E = the activation energy for the latent defect process being studied
T = absolute temperature, in degrees Kelvin
k = Boltzman's constant (8.63×10^{-5} eV/K)
K = constant

We can also use this model in terms of time-to-fail. If we know the mean-time-to-fail (MTTF) for the process at one temperature, we can determine the MTTF at another temperature with the following derivation of the above formula:

$$t1 = (t2) \exp [(E/k) (1/T1 - 1/T2)]$$

where: t1, t2 = MTTF at T1 and T2 respectively
T1, T2 = temperatures in degrees Kelvin
E = the activation energy for the latent defect process being studied
k = Boltzman's constant (8.63×10^{-5} eV/K)

For example, if we had failure data taken at elevated temperatures (in order to speed up the testing process), we could translate it back to compute the failure rate at a lower, normal operating temperature. The equation assumes a fixed activation energy for a particular failure mode regardless of the temperature or the age of the product. However, if we believe the research showing that failure rates *decrease* over time as opposed to

being constant during most of the life of the product as discussed in Chapter 3, then it follows that the activation energy should change with time and temperature as well.[15]

(2) Voltage acceleration model (voltage). Electronic products can also be stressed by applying non-nominal voltages. Voltage accelerated stress test results can also be translated to nominal voltage conditions, using a relationship similar to Arrhenius above, by applying the appropriate voltage acceleration factor. For example, semiconductor manufacturers use this relationship to analyze time-dependent dielectric breakdown failure mechanisms. The equation used to relate these two voltage levels is the following:

$$\mathbf{VAF = exp\ (C[V_S - V_0])}$$

where: VAF = the voltage acceleration factor
V_S = the applied stress voltage
V_O = the standard operation voltage
C = a constant that is a function of dielectric type, between about 0.5 and 3.0

This relationship is more effective in predicting failure modes at the component level. For a larger assembly such as a formatter board, the effect of voltage stress is much more difficult to isolate, since marginal voltages can also affect timing. As a result, interactions between components could lead to product failure, which is much more difficult to analyze and more likely to be intermittent or temporary in nature.[16]

(3) Miner's equation (general). For the Miner's equation to be used to describe components strength and deterioration, the following three conditions must be true for the part under test: "(1) damage (strength degradation) accumulates continuously in the component, (2) the component fails as soon as damage totaling a certain amount has accumulated (strength is degraded to the point of the imposed stress), and (3) the damage rate under a certain stress at a certain time depends only on how much total damage there is at that time, and not on the past history." If these are all true, we can then describe the component lifetime with the generalized Miner's equation:

$$\sum_{i=1}^{n} (t_i / L_i) = 1$$

where t_i = the actual time at stress i;
L_i = the expected lifetime at this stress condition i

This relationship can be used to design testing or to determine lifetimes as follows. Assume we are given a severe and light stress level, t_1 and t_2. We could run the part at t_1 until it breaks to determine L_1, and then we would have enough data to determine L_2. The biggest problem in using this method is proving that the third assumption is true. This condition requires us to take a very narrow view of the dynamics of the failure mode, and allows no room for factors to influence part life other than its immediate physical status.[17]

To use these relationships, we need to understand the assumptions behind them and what situations they were designed to model, so that we do not use them improperly or inaccurately.

4.4.3 Test Design and Implementation

The major difference between accelerated testing and the other test methods we have discussed is the relatively high levels of stresses we place on the products. For this reason, the challenge in test design is to develop a procedure which is harsh enough to detect infant mortality problems, but not so harsh that it has a noticeable impact on product life for the customer. One way to approach this is to design a test program which we can run on several parts and verify that it detects the necessary failure modes. To ensure that it is not *too* severe, we can run the same program several more times on the surviving parts to see if they still function. Test designers have estimated that effective tests can use anywhere from 0.25% to over 10% of the product life and still be effective.[18] Fig. 4.5 presents a flowchart for this process of production test design.[19]

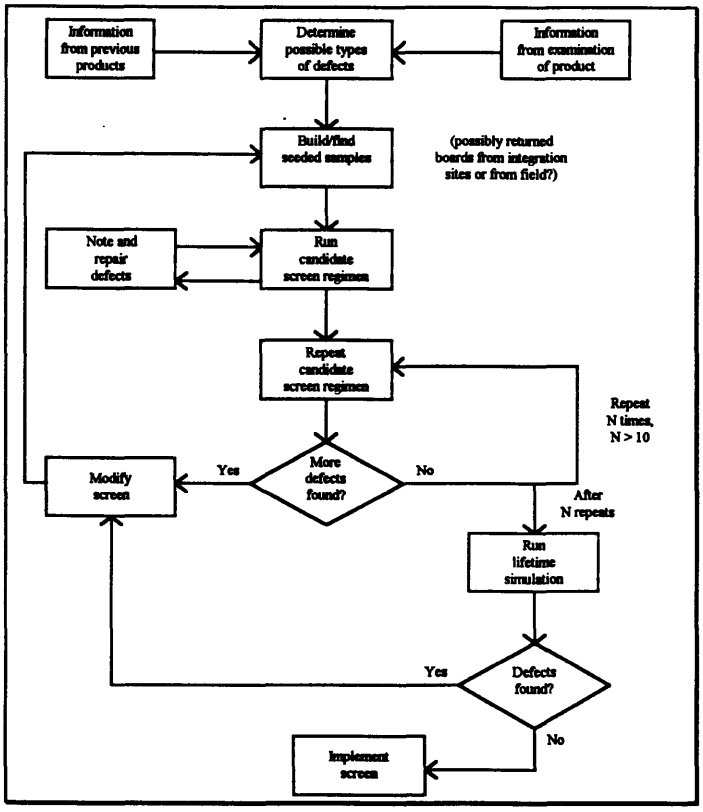

Figure 4.5 Flowchart for production-level accelerated test design.

An important point in this flowchart is the notion of using *seeded samples*, boards either known to have certain types of latent defects in them, or which have been intentionally built with weaknesses or potential problems (parts from a low quality vendor, solder with insufficient flow, misaligned components, etc.) Seeded samples make it possible to speed up the test design process, especially when the number of available samples is small or if there is limited time for using a test chamber.

Below are some specific issues or rules which need to be addressed in implementing a successful accelerated stress test:

(1) Monitor the products as stresses are being applied. During accelerated testing, unless the products are not at all robust to begin with, it is absolutely essential that the products under test be monitored as frequently as possible. Many relevant failure modes only appear in response to a particular type or combination of stresses, such as a fast negative temperature ramp or vibration during a dwell at high temperature. The failures cannot then be duplicated when the board is returned to "normal" conditions.

For example, fast temperature ramp rates (especially from high to low temperatures) are often effective in triggering failure modes on formatter boards. However, it is important to be able to spot and isolate them quickly, since they will often show up only temporarily. The test programs currently in use at BPR are able to completely test a board in about two minutes, but in that same time, the test chamber is capable of going through two entire temperature swings, from +130°C to -100°C and back again. As a result, a failure mode which reveals itself only during a particular portion of the ramp may not be seen by the test designed to catch it. To get around this problem, we instead "tested" the board under stress by connecting it to a printer outside the chamber and continuously printing actual test pages. The act of printing involves most all of the functions of the board at all times, thus if a defect appears, we will see a problem in the printed output. Once we know there is a problem to be found, more detailed tests can be run during that part of the temperature profile to gather more information.

2) Design the test procedure and sequencing such that results are explainable. Since accelerated stress testing is still a new field in many respects, there are few universal rules as to what sorts of stresses will guarantee a particular result. As a result, we need to initially control the combinations of stresses and proceed in such a manner that we can understand the impact of each change we are making to the test procedure. On a basic level, it is educational to start with a "change one factor at a time" approach. For example, start with temperature cycling on several boards, then repeat the same tests but adding vibration, then repeat all those tests but add humidity, etc. Since there will be many types of interactions which take place between certain combinations of stresses to excite failure modes, a more thorough way to approach test design is a Taguchi-style design of experiments. This will help us to better choose combinations of stresses to try, and to better analyze the resulting failure modes. The factors for this analysis would be chosen from the types of stresses which can be generated using available equipment (vibration, temperature ramp, temperature dwell, voltage stress, humidity, etc...). Test results could be measured in terms of the number of failures which occur at each permutation of stresses. For example, Sun Microsystems used Taguchi principles to compare the general effectiveness of random vibration, temperature cycling, and power cycling in detecting defects without overstressing their products.[4] The experimentation showed that vibration with continuous monitoring was extremely effective in detecting a variety of failures, and that temperature cycling and power cycling together were more effective than either stress exerted individually. Sun used the results of this investigation to aid in designing a detailed production-level test. The Taguchi approach could also be used to set the actual stress levels for each stress type during this second stage of test design.[20]

Another way to keep track of the interaction of variables is with a Schmoo plot. Originally developed for use with semiconductors, Schmoo plots chart two variables against each other to more easily show the interactions between variables. For example, Fig. 4.6 shows a Schmoo plot drawn for various combinations of product temperature and

[4]The product under test was a 8.5 x 11", double-sided, high-density, predominantly surface-mount circuit card designed for use as a central processing unit (CPU) for one of Sun's SPARC workstations.

power voltage. By looking at the plot, we can quickly see that the product is weakest at combinations of higher temperatures and lower voltages. If the desired operating region included some of this area, improvements are necessary.[21]

	4.6V	4.7V	4.8V	4.9V	5.0V	5.1V	5.2V	5.3V	5.4V	5.5V
100°C	x	x	x	x	x	o	o	o	o	o
80°C	x	x	x	x	o	o	o	o	o	o
60°C	x	x	x	o	o	o	o	o	o	o
40°C	x	x	o	o	o	o	o	o	o	o
20°C	x	o	o	o	o	o	o	o	o	o
0°C	o	o	o	o	o	o	o	o	o	o
-20°C	o	o	o	o	o	o	o	o	o	o
-40°C	o	o	o	o	o	o	o	o	o	o
-60°C	o	o	o	o	o	o	o	o	o	o
-80°C	o	o	o	o	o	o	o	o	o	o

Key: x = *failed test,* o = *passed test*

Figure 4.6 Sample Schmoo plot for viewing effects of two-dimensional stress combinations.

(3) Eliminate distractions which can be confused with legitimate product defects. It is necessary to be able to isolate true product failures from other test problems. This is especially important when testing a product which is already highly reliable. In this case, we must treat every single failure as being significant. To increase the quality of resulting data, we should take the following precautions:

--Minimize the amount of test equipment actually inside the test chamber. This equipment will receive the accumulated wear of all the products being tested, which will cause failures that have nothing to do with the product itself. Watch for hairline solder cracks in the test equipment as a sign of wear and potential failure.

--Understand the correlation between the function of the program and the operation of the actual product when running test programs, and also know what each subprogram has in common with the others. When multiple tests are showing problems during a monitoring cycle, knowing the commonalities between these tests can provide clues as to the root cause of the problem.

--Fixtures for holding the product in a test chamber should duplicate the way the product is mounted or used in the user's environment, in terms of orientation, air flow, cabling, etc... This is especially important for vibration testing, since fatigue at specific locations in the product is a function of where the stresses are being introduced into the product. In addition, the operation of the chamber itself may introduce irrelevant failure modes. Potential causes include

electric or magnetic fields generated by the switching on and off of fans or other electrical equipment in the chamber.[22]

4.4.4 Industry Examples

As in all areas of learning, we should begin an investigation into testing possibilities by taking a look at what others are doing. There have been many major successes using accelerated stress testing, but unfortunately, many companies are unwilling to share their success stories in public. Accelerated stress testing is a fairly new field. As a result, companies with a more advanced understanding of its use can create tests to improve their products' reliability over the competition, but the competition will not have enough skill in the field to match those results. The following examples are available only because the company's names are unavailable, or because the company published their initial findings and later decided to stop further publications, once they realized the unique potential that this testing can provide:

--When one generation of the IBM Proprinters was released to manufacturing, every product was put through a vibration testing screen before it went out to the field. This screen is credited with the first field failure not occurring until 9 years and 8 months after the product was sold.

--A workstation manufacturer designed a production test used on 100% of products before they were sold. Engineers found that the screen was responsible for a tenfold reduction in field failure rates.

--A television manufacturer performed highly accelerated testing in order to reduce a high level of warranty problems. After one day of testing, they identified the source of failures responsible for $2M per month in warranty costs. The solution to the problem cost only 1.3 cents per television.

--A company invested $1M in equipment, training, and development time to put together stress testing programs for both product development and for production, and kept track of the field failures for the products that came out of this program. When comparing them to what the failure rates *would* have been without the benefit of these tests, the company estimated that it saved $100M in the first three years of production of one product alone.[23]

As these examples show, the benefits of these tests can be quite impressive. In the following section, we will review a few examples of the actual stresses and tests that companies are using as part of their accelerated stress testing procedures.[5]

(1) RND Board Designers. Once prototype samples are available for a new product at RND, test engineers put hundreds of samples through combinations of the following types of tests to detect potential failure modes in the final product:
a) *Thermal shock testing* - at least 50 temperature cycles at a slew rate of 10C°/minute
b) *Temperature, frequency, and voltage margin testing* - boards are tested at various combinations of the these parameters, significantly outside of required specification ranges
c) *Passive vibration testing* - boards are tested for two hours and then checked for functionality
d) *Monitored high temperature and high humidity soak* - boards are held at a constant stress level for at least twenty hours

(2) BRD Board Shop. BRD builds small printed circuit assemblies which are assembled into final products at the same location. As the boards finish production, a sample of them are put through a test which lasts about 110 minutes. The test consists of temperature cycling, power cycling, voltage stressing, and vibration at temperature extremes. This test has been in place about five years, and has helped the company to reduce its field failure rate significantly. The tests detect problems such as delaminations inside integrated circuits, bad joints, lifted leads, memory problems, and other process-related failures such as misinserted diodes. BRD estimates that the vibration and thermal cycling makes it possible to reduce the test times for the boards by about fifty to one hundred times, while giving them the same amount of overall stresses. At the same time, they estimate that the tests reduce the board life by five to ten percent, but they are confident that the boards which *pass* the tests are actually more reliable than the boards which are not tested at all.

(3) QFD Board Builders. In addition to the online testing, a sample of the production boards also receive Board Environmental Stress Testing (BEST). This testing consists of repeated thermal cycling for at least 48 hours; eight to sixteen boards from various product lines are under test at any given time. During the cycling, functional-type tests are also run to monitor the boards. Circuit cards containing the test programs as well as several interfacing cables have to be in the chamber with the board to perform the tests, so unfortunately all of these "peripheral" items are stressed along with the product under study. As a result, then, it is hard to get meaningful data from these tests because boards often fail due to a poor connection or to the test setup, and it is difficult to determine whether the product was actually defective or if the failure was in a part of the system outside of the product.

[5]The data and company information here is necessarily vague in order to protect the competitive advantage of companies using these techniques.

(4) ABC Disk Drives. Production pieces of hard disk drives at ABC Company went through the following three types of testing:

a) *Functional tests with voltage stress.* Before circuit boards were attached to the rest of the disk drive, the boards were subjected to accelerated voltages (3-6% above nominal). The length of this test was originally 115 minutes, but it was reduced to only 60 minutes since most of the failures occurred early in the test cycle. For example, in one sample of 243 boards, almost 80% of the boards failed in the first 11 minutes of stress. Failed boards were sent back to the manufacturing facility and repaired. Unfortunately, the manufacturing site was only able to detect specific problems on 20-30% of the boards which were returned to them.

b) *Functional tests in a non-accelerated environment.* This test was originally 60 minutes in length, but studies showed that a majority of the electrical failures occurred during the first 10 minutes of the test. These tests (more comprehensive than the tests run earlier) detected a variety of component-level failures. When an electrical defect was found, technicians removed and repaired the circuit board, and reused it in another drive.

c) *Monitored power and temperature cycling.* The entire disk drive was placed in a walk-in chamber and temperature cycled at the rate of about 30C° per hour for about 24 hours. Monitored tests during this time consisted of writing data onto sectors of the disk and then checking for access times and error conditions.

(5) XYZ Disk Drives. XYZ Company went through the following procedures for their disk drives:

a) *Passive temperature cycling.* During this cycling of only the mechanical part of the drive, no functional tests were performed and no power was applied to the product. These tests required approximately eleven hours.

b) *Functional test with no stress.* The products were then put through a functional test, again for the drive itself, consisting of writing and reading data to various parts of the disks.

c) *Monitored temperature cycling.* Once the circuit board was attached to the drive, the complete assembly was functionally tested; this included checks for bad spots on media, data channel tests, head shift/seek time measurements, and also power cycling at different temperatures. These tests originally lasted for 18 hours but were reduced to 11 hours by streamlining some of the tests. In order to allow this, engineers performed months of testing to satisfy XYZ as well as its customers that this change in test time would not have a negative impact on product quality.

d) *Monitored burn-in.* Finally, two hours of final functional testing at ambient temperatures took place. The majority of this time was spent loading customer-specific firmware onto the drives.

4.4.5 Results of Formatter Experiment

To better understand the details and the possibilities for accelerated stress testing, we performed two days of accelerated stress testing ourselves, using the same type of production

boards from the burn-in testing described in Section 4.3.3. Our goals in this experiment were to learn about the following:

-- what sorts of failures would appear in the boards
-- how consistent the failures would be
-- what sorts of stresses would excite the failures
-- how the detected types of failures compared to the other burn-in testing

We used a total of four boards for the tests, and we designed two types of monitoring for checking the functionality of the boards as stresses were applied. For continuous print testing, we connected a formatter board in the chamber to a printer outside the chamber using Teflon cable for the connections, with toroids around the cable to minimize electromagnetic interference. A PC attached to the formatter provided continuous print patterns, or self test pages were continuously printed using an internal print test on the printer during vibration testing. We kept connections and cable lengths in the chamber to an absolute minimum in order to avoid failures due to the equipment rather than the board itself. We designed the second type to provide more detail about the failure modes as they occurred. We attached a SIMM memory card and a SIMM loaded with test programs to the board in the chamber, and made connections to the four input/output connections on the formatter board to monitor the output of the tests. A looping program on the PC controlled the sequencing of these tests.

See **Appendix C** for the specific stresses which we placed on each board during the tests. Two types of stress were possible with the test chamber: temperature and vibration. The chamber was capable of changing temperature extremely fast. Slew rates of over 60C° per minute were observed. The upper and lower temperature limits of the chamber were approximately +140 °C to -100°C. In addition, the chamber created vibration by using pneumatic hammers underneath the plate on which the board was mounted; these hammers struck at random times, but their intensity was controllable. The chamber was thus capable of six degree-of-freedom stressing, which implies that the board was vibrated in all possible directions even though it was

mounted in a single orientation for all of the tests. During our testing, we used the chamber to its limits and to the limits of our test equipment. We experimented with the following combinations:

-- low and high temperature dwells
-- low and high temperature dwells with vibration
-- fast upward and downward temperature ramping
-- fast upward and downward temperature ramping with vibration

During our testing, we were able to gain a better feel for the complexity of accelerated stressed. We made the following conclusions based on our trials:

(1) The boards are already extremely robust. Each board spent 2-3 hours in the chamber, and all of them were cycled close to the limits of the chamber. In every case, no permanent failures appeared that we could associate with a known good part breaking due to the stress.

(2) Failures during the tests were consistent. All four boards failed for the same error around +128°C, even with different types of stresses being applied to them before going to this temperature range.[6] Each board would also repeatedly exhibit the same failures at the same stress levels. The failures would not occur any sooner or at less strenuous levels after prolonged testing. We concluded from this that Miner's equation in Section 4.4.2 should *not* be applied to this sort of testing, since the required assumption of accumulated stresses being the sole cause of resulting fails does not appear to hold true.

(3) Each type of stress excites unique failure modes. With the limited amount of testing that was done in this experiment, it seemed that all the failures were excited by a specific type of stress; for example, the failure at +128°C would occur in the same manner regardless of the amount of vibration that was applied. This supports the potential benefits of Taguchi or seeded samples from Section 4.4.3 as reliable methods for choosing useful stress levels, in that these techniques would help us gain more specific information about the relationship between stresses and failures.

(4) More experimentation is needed. This experimentation was valuable in helping us understand what types of failures to expect, but we would need to test a larger number of boards to learn which errors are due to latent defects and which are unavoidable or inherent in the product, but would not correspond to real failures in the field. In addition, if we were designing a production-level test procedure, we would need to do more testing in order to understand how to balance stressing the board hard enough to precipitate failures with not overstressing the board, causing it to fail for the customer. Finally, we would have liked

[6]These boards are used in printers whose rated environmental temperature range is specified at 0°C to +35°C. This prompts the question of whether failures which occur so far outside of the normal use range have any relevance to normal product life, as discussed earlier in this chapter.

to test the Arrhenius equation in this setting. However, the relationship requires that we know time-to-fail at two different temperatures. Since we could only find one consistent failure mode at one temperature within the limits of the chamber, we could not test this equation.

4.5 Summary

Although a variety of methods and applications were discussed in this chapter, they all share a common purpose: the precipitation and/or detection of latent defects which would otherwise be passed on to customers. The specific methods--visual inspection, burn-in, and accelerated stress testing--possess many of the same qualities, but in varying levels. The differences among each method (and the application of each method) include the degree to which the product is monitored as it is being stressed, and the amount of stress that it exerts on the product as it precipitates failures. Table 4.4 presents a summary of the advantages and disadvantages of each of the major types of latent defect testing.

Option	Pros	Cons
no latent defect testing at all	no extra work-in-process simpler material handling	increased failure risk
any type of latent defect testing	ability to screen for problems faster class detection failure detection of component failures caused by process could catch add'l failures that are missed by downstream processes	additional cycle time (for burn-in and testing) could "encourage" lax testing by vendors connections can create false errors increased work-in-process inventory additional product handling can cause problems (dropped board, mistested boards...)
visual inspection	direct detection of unprecipitated defects low investment needed for manual inspection	slow process tedious, marginally reliable work to do manually
passive burn-in	easier fixturing low risk of excessive stress being applied	redundant testing/handling required can't isolate time-to-failure can't isolate intermittent failures may precipitate but not detect defects

monitored burn-in	reduced time needed for defect detection can measure/learn from time-to-failure (more easily determine effective run-in time) can isolate intermittent failures burn-in and testing done in same location ∴ minimal handling can reduce time needed for traditional functional testing	redundant testing required expensive/difficult fixturing required
accelerated stress testing	reduced time needed for test; defects show up faster can catch a wider variety of failures can simulate a wide variety of stresses	high equipment cost may not accurately simulate end-use environment can induce undetected mechanical damage (thus weakening the product) higher support costs (faster wear on fixtures) need to avoid excessive degradation of board

Table 4.4 Summary comparison of various testing methods.

One final tradeoff we should consider here, regardless of the specific types of test chosen, is the number of parts to be tested. Should we screen 100% of production, or is a sample sufficient? Sampling only a percentage of production can save time and space, which translates into monetary savings as well. However, for sampling to be effective, we must have the commitment and the resources to fully analyze each problem that arises. When sampling for production testing, it may take two to three months to collect enough data to be sure of a problem. In real life, however, we cannot wait for statistical significance before acting.

If a majority of the problems testable by the above methods occur in relatively large quantities when they happen, sampling may work. Examples of root causes that would fit into this category include changes in supplier quality or processes, the use of a new supplier, or changes in board assembly processes. On the other hand, if most failures are difficult to categorize or trace back to a common source, then 100% testing may be necessary. When implementing tests in production where defect rates are small, we should start out by testing 100% of all boards to get a significant sample size. After intensive failure mode analysis and

corrective action during production startup, a sampling could be sufficient for ongoing production checks.[24]

When determining what test methods would be appropriate, we should look at the available information in detail, much as we did in Chapter 3 when studying defect causes. One approach for studying this is to use the same failure modes which were collected from the available defect data in Chapter 2, and determine which modes the current processes are able to detect, and what testing methods would be most appropriate to implement in order to precipitate and detect these types of defects. See **Appendix D** for how this process can be applied to a sample of the formatter data. By scanning through the information, or by weighing each failure type relative to the measured or expected field failure rates, we can learn which testing processes would be beneficial.

So far, we have considered both the potential *problems* of latent defects and the potential *solutions* that test methods such as visual inspection, burn-in, or accelerated stress testing can provide. In the next chapter, we will complete our framework for latent defect testing by discussing the degree to which the costs and benefits of testing can be quantified for the purpose of making a final decision.

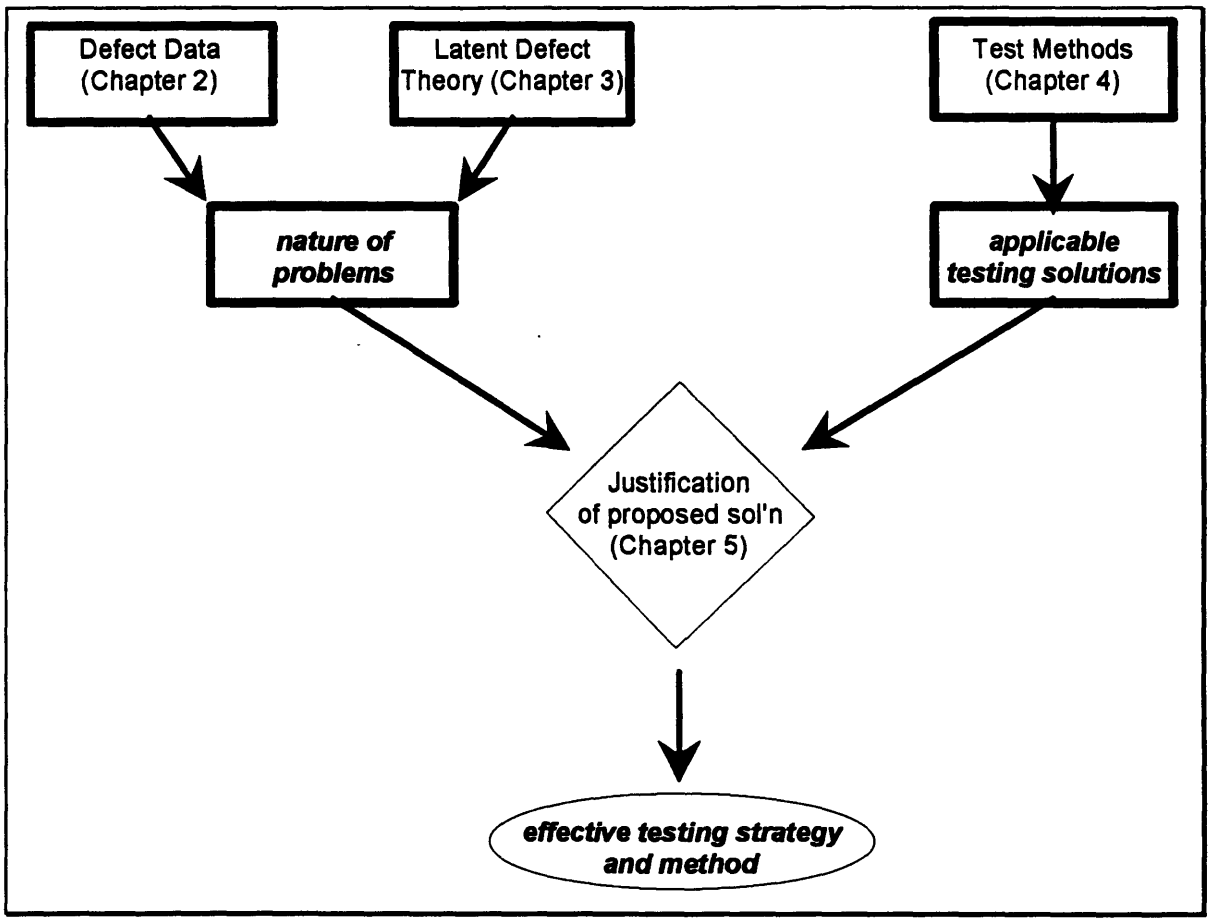

Figure 4.7 Review of latent defect testing decision process.

References

[1] Peterson, Michael, "Factors To Consider In Designing a Test Strategy for Latent Defects," Proceedings of the 1994 Hewlett-Packard Electronics Assembly Manufacturing Conference, p. 222.

[2] O'Shea, Paul, "Burn-in Changes Technique to Fit The Need," Evaluation Engineering, November 1993, pp. 117-120.

[3] Parker, T. Paul and Harrison, Gordon L., "Quality Improvement Using Environmental Stress Testing," AT&T Technical Journal, July/August 1992, p. 14.

[4] O'Shea, p. 117.

[5] Harm, Charles E., "Real-World Application of Burn-In to Improve Product Quality," Thirteenth International Telecommunications Energy Conference - INTELEC 91, 1991, p. 563.

[6] Pynn, p. 124ff.

[7] compiled from discussions with various H-P personnel (fall 1993) and from STRIFE presentation notes by Eric Andersen, 1992.

[8] Jensen, p. 75.

[9] Hu, J.M., D. Barker, A. Dasgupta, and A. Arora, "Role of Failure-Mechanism Identification in Accelerated Testing," Journal of the Institute of Environmental Sciences, July/August 1993, pp. 39-45.

[10] Hobbs, Gregg, editor/author, notes for Screening Technology: Design and Process Improvement seminar, Hobbs Engineering Corporation, 1990, p. 5.

[11] Hobbs, p. 13.

[12] Wong, Kam L., "What is Wrong With Existing Reliability Prediction Methods?," Quality and Reliability Engineering International, Volume 6, 1990, pp. 251-257.

[13] Hu et al., p. 42.

[14] O'Connor, Patrick, Practical Reliability Engineering--Second Edition, John Wiley & Sons, 1985, p. 188.

[15] Blanks, Henry, "Arrhenius and the Temperature Dependence of Non-Constant Failure Rate," Quality and Reliability Engineering International, Volume 6, 1990, pp. 259-265.

[16] Intel Corporation, Components Quality and Reliability, Intel Corporation, 1990, p. 7-1.

[17] Jensen, pp. 85-87.

[18] O'Shea, p. 117.

[19] Hobbs, Gregg, editor/author, notes for Screening Technology: Design and Process Improvement seminar, Hobbs Engineering Corporation, 1990, p. 30.

[20] Pachucki, Dennis E., "Environmental Stress Screening Experiment using the Taguchi Method," 1993 Proceedings - Institute of Environmental Sciences, IES, 1993, pp. 211-219.

[21] Parker, p. 13.

[22] O'Shea, p. 119.

[23] Hobbs, Gregg, Highly Accelerated Life Tests--HALT™, published by author, 1990, pp. 11-12.

[24] Parker, p. 12.

Chapter 5 - Financial Aspects of Latent Defect Testing

5.1 Introduction

After collecting and analyzing data and after investigating the range of possible testing methods, we must compare the *problem* to the *solution* in order to determine if latent defect testing is worth the effort. To accomplish this, we need to gather information on the financial impact of testing on the manufacturing operation, and also on the projected benefits it will generate for downstream process steps. In this chapter, we will discuss the factors which determine the costs and benefits (both tangible and intangible) involved in this decision, and we will look at several techniques for comparing the costs and benefits to each other in order to reach an overall conclusion.

5.2 Costs of Testing

The tangible expenses involved in setting up and running a testing system are similar to those for any sort of capital investment in the manufacturing environment. Testing creates cost by affecting the following parameters in manufacturing:

--disruption of other process steps
--time delay in shipping products (due to length of test)
--work-in-process inventory levels
--usage of floor space
--support costs (materials, utilities)

As a result, we must consider the following types of costs:

--holding costs for inventory waiting to be tested or under test
--opportunity cost of floor space
--utility/maintenance costs
--cost of the equipment itself[1]

[1]For example, test chambers which are capable of the accelerated stress techniques discussed in Chapter 4 can cost anywhere from $120,000 on up for the equipment and the computer support, not to mention the support systems which need to be installed into the building for electricity, noise suppression, gas disposal, etc...

These costs will vary depending on the volume of products being manufactured, the number of different products involved (especially if different fixturing requirements are needed for each product to test it), and the number or percentage of parts which will actually be tested.

Another category of costs are related to people. We must pay for personnel to not only run the test process, but also to analyze the results. The cost of these people will include not only their time, but also the support equipment they need to analyze the products and determine the sources of the problems. This includes devices such as power supplies, customized test fixtures, data logic analyzers, etc., which can be quite expensive. As a result, it will be important to invest in equipment which is flexible and multifunctional, so that it remains useful as designs or design technologies change.

The intangible costs of testing are by nature more difficult to define. Testing can increase production costs due to the following factors:

(1) Increase in cycle time. Adding test processes to a production line increases the cycle time. All else equal, this increases the amount of work in process and could also slow down the transfer of information through the build area. For example, consider a board with a visually obvious flaw in it. If visual inspection is not done until after four hours of power cycling on the boards, we would build a half shift of defective products before the problem is noticed and communicated back to the manufacturing line for corrective action. This causes longer feedback delays to production, meaning that more defective products will be built. For this reason, duplicating tests before and after the latent defect test may be necessary in order to isolate defects in a timely manner.

The importance of minimizing cycle time also implies that we should set the test time long enough to be effective, but absolutely no longer. Extra test time translates into needs for more test fixtures, more fixture space, more material tracking, more disruption to the current system, more work-in-process, etc. Each of these increases the costs of production. In addition, it is difficult to justify reducing the time of a test once it is in place. Engineers making these sorts of changes have collected months or even years of data and presented it to management and outside customers before they would approve the time reduction.

(2) Extra product handling. As more processes are added to the production process, the boards come into contact with more machines and humans. This creates additional stresses on the boards. Avoiding human contact entirely would help eliminate problems such as

dropping boards or misinserting them into test fixtures. However, the automation necessary for making this possible adds significantly to the cost of the testing.

(3) Obsolescence of test system. The flexibility of the system is critical to its cost effectiveness. For example, consider a burn-in test designed for a single product. It is likely that after a few months of production testing, the burn-in process would detect all of the different failure modes for which it is capable. If these failure modes can be prevented through corrective action, the need for the test will be negated. To continue to benefit from the investment made in the equipment, people, etc., we must be able to change the types of testing being done, or to adapt the test fixturing for multiple types of products. In other words, some additional upfront investment for items such as multiple connections to the test fixture or changeable software will allow us to spread the expense of the test out over a longer period of time.

When dealing with intangible costs and benefits, it is difficult to exactly quantify the values of each of these factors. It may be more practical for us to first evaluate the net cost or value of the testing based on the tangible factors, and then decide if the value of the intangible factors is sufficient to make the project effective, or to skew the analysis to one side or the other. We must be careful, however, not to completely ignore these factors in our analyses. For example, we may be unsure of the cost of increased cycle time, but we know it should be at least equal to the value of the inventory which is forced to wait during the production process. In this case, we should figure in at least that valuation of the inventory, rather than giving up entirely.

5.3 Benefits of Testing

The goal of latent defect testing is the precipitation and detection of latent defects which would otherwise develop into actual product failures for customers. We can define the benefits of our testing as the degree to which the tests help us accomplish this goal. As with the costs, there are tangible and intangible benefits to testing. Some of the more straightforward measures we can employ are the following:

--the value of the reduction in failure rates at each of the downstream processes as a result of the testing
--the cost of doing no latent defect testing and repairing the product at a later stage of assembly or warranty vs. testing and repairing it earlier
--the expected cost savings due to detecting a major class problem in the factory before the "bad" product could reach customers

--the difference in warranty expense between two similar products, where the primary difference is that one received latent defect testing and one did not
--the change in the "cost to use" for the customer, which includes *all* expenses incurred in operating the product: purchase price, cost of replacement parts and refiller supplies, cost of repairs[2]

Most of the benefits of testing which are difficult to quantify are related to *customer satisfaction*. We could treat equally the cost of traveling to a customer to do a warranty repair with the cost of the same repair during final assembly. However, all else the same, the customer would much rather have the product fixed before he receives it rather than after, and we should account for this in our analysis.[3] We should also consider the impact of field failures on future sales. When one customer has a bad experience with a product, he is likely to tell his friends about his problems, which decreases the probability that they would buy the same product. The point we made in Chapter 2 about the lack of information on repairs which happen after the warranty period becomes quite important here. Non-warranty repairs which the customer has to pay for can have a large impact on his satisfaction, but we will have less information about these repairs. Rough studies at H-P have shown that the real cost benefit of screening may be 5-6 times the savings in warranty repair costs. In other words, by preventing a $1 fix in the field, we could expect a $5-6 return in profitability of the product (taking into account sales volume and other factors). This implies that, if we can at least justify the costs of a test based on warranty savings alone, we *know* we have a good project, because savings due to customer satisfaction will be a bonus on top of that.

---

[2]*Cost to use* should be an important metric for customers to consider when they are purchasing a product such as a laser printer. Paying a low upfront price only to have to put up with frequent repairs is just as bad (or worse) than paying more for a quality product in the first place. Unfortunately, though, each manufacturer compiles and advertises this sort of information differently, making it difficult for customers to compare different products by this metric. (Other metrics of comparison have been more successful, such as car mileage ratings or large appliance operation costs.) The result is that customers usually end up doing their own informal analysis by talking to friends, checking warranty contracts, etc.

[3]To make this even more complicated, some manufacturers have actually succeeded in *improving* their quality reputation by the way they handle repairs or recalls of products which customers have already purchased. One example of this is the early recalls of General Motor's Saturn automobiles, in which Saturn replaced the entire car rather than replacing the faulty components, which by most standards would have been sufficient.

Another dimension to consider here is that the customers of the product are not only the end users, but also the people involved in *downstream assembly processes*. For example, as discussed in Chapter 2, some of the formatter boards produced by H-P are assembled into printers by Canon at a separate Canon-owned facility. As a result, when problems are observed at Canon with H-P's boards, this reflects poorly on H-P. The desire to keep their relationship and reputation with Canon intact makes it worth some extra expense for H-P to add extra testing processes, even if a majority of the defects for which H-P is testing would be detected by Canon's downstream processes.

Finally, the *effective life* of a testing plan has implications when considering benefits as well as in calculating costs. We need to understand how long the testing will take to justify itself, in addition to thinking about how long we can expect the tests to be effective. For example, it requires up to three months to do a complete root cause analysis on a problem found during latent defect tests, especially if intermittence is occurring or if the component needs to be sent to a vendor for final analysis. It may also take up to a year to discover all the major defect modes of an existing product, and over a year to see if the testing is having a significant impact on field failure rates. As a result, the lessons learned from latent defect testing may come too late to fix the existing product, but they can improve the quality of future products.[4] As a result, our analysis should even consider the effects on field rates of design or process changes made for future products.

[4]Before using this reasoning to justify a long-term testing investment, though, we need to be sure that the organization is capable of transferring knowledge in these ways, and that "corporate memory" is long enough that the learning will not be lost due to job transfers of key personnel or to a general short-term approach to resolving issues.

5.4 Financial Comparisons

5.4.1 General Issues Related to Test Justification

Before we look at specific methods for financial analysis, we should determine our assumptions about the product and about the testing in general. Below are three factors which will help us determine a specific approach.

(1) Consider the nature of products when setting sample sizes. The percentage of products we test has a major impact on our testing costs as well as our potential benefits. The amount of learning we obtain from each test board determines the necessary size of our test sample. On one hand, we can expect failures to have definable root causes which are present in other parts as well. In this case, it would not be necessary to test every single product to discover and correct failure modes. With this type of approach, sampling would likely be more cost effective than 100% testing. On the other hand, we may expect defects in the products to be difficult to learn from, or we believe that defects have been designed out of the product as much as is feasible. In this case, the primary focus will be on catching each failure. Determining the root cause and learning from it is still desirable, but secondary in importance. This would require us to test closer to 100% of the production in order to gain much benefit.

The way we view the nature of defects also affects how we use failure rates at each stage of production to justify testing costs. If we do root cause analysis on each defect and find that we can eliminate those types of failures, then we may be able to eventually eliminate the need for testing. In this scenario, the only way we could continue to justify the test is if new (analyzable and preventable) failure continue to arise. On the other hand, we may find that the failures we catch are difficult to analyze or prevent, or that there are frequent changes in the process (vendor quality, new vendors, production parameter changes) which can introduce new defect modes. In this situation, a one-time drop in downstream failure rates would tell us that our testing process is effective. Here, we can use the difference in failure levels (original minus the new) each period to keep justifying the testing, since the new failure modes which continually arise would bring the defect rate back up again if we stopped testing.

(2) Look at the entire system before optimizing a portion of it. We have restricted our focus in this thesis to tests involving the formatter boards for reasons of simplicity and focus. By doing this, though, we've made the implicit assumption that spending money to improve formatter failure rates has the greatest impact per dollar spent on printer quality. In real life, we should first analyze several of the leading causes of failure in the final product, in order to determine which problems could be eliminated most efficiently.

(3) Look at alternatives which exist outside of the current process. The evaluations in the next section all work within the current framework for printer production and repair: boards are assembled into printers, printers are bought by customers, repairmen come fix the

printer or the customer brings the printer into a service store to be repaired. It may be beneficial in some situations, though, to develop cost proposals based on re-engineering the existing process. If we can justify the "radical" proposal, we should focus our efforts on changing the whole system, not adjusting the current process.

For example, instead of servicing a product, some manufacturers agree to ship a customer a brand new product in exchange for the defective one. We can calculate the cost of this service policy for laser printers as shown below:

Cost to replace each printer[5]:	$950.00
Annual volume of sales for printer:	x 2,000,000
Expected % of printers needing replacement[6]:	x 0.01
TOTAL ANNUAL COST for program	$19,000,000
COST per each printer built, for program	$9.50

From this analysis, we conclude that any effort we make to improve overall printer quality should not cost more than $9.50 per printer produced (good or bad)We could more easily ship each customer a new printer at that price. This sort of analysis also makes it possible to estimate what the service operation is actually worth, and could be used as a benchmark for developing alternative systems.

5.4.2 Calculation Methods

This section presents five methods for calculating the net benefit of latent defect testing in financial terms. We will see how our conclusions from the data and theories of latent defects impact the way we determine an actual monetary value. All calculations for the rest of this chapter are based on actual data for laser printers, but the details as to how each subtotal was computed are not shown for proprietary reasons. For different situations, we would use different types of cost data based on availability and relevance. For example, the calculation shown above does not factor in the following important costs and benefits:

-- hassle of customer mailing printer back to H-P
-- dissatisfaction of customer being unable to get a repairman to do "simple" repairs
-- dissatisfaction of customer not being allowed to keep his own printer, or being forced to take a different model

[5]Includes cost of shipping a printer to and from a customer, and the cost to H-P to manufacture a new printer.
[6]The fraction of printers which receive repair calls during their first year (warranty period) and would thus be replaced in this policy.

-- cost of loss of markets where local service is desired/required/mandated
-- cost of space needed for disposal/refurbishment/analysis of returned printers
-- cost of serious problems not being detected until products reach the final customer (pipeline full of bad printers)
-- benefit of greater information available through in-house analysis of returned printers
-- cost savings from disbanding of the service network

The better we can define all possible factors for our calculations, the better decisions we will be able to make.

5.4.2.1 Warranty Cost Analysis

One method for quantifying the cost side of the cost/benefit tradeoff is to determine the amount of money currently being spent to fix defective products in the field. For example, we collected the following information on printer warranty repairs over a six month period:

Monthly cost for ALL warranty repairs	$275,000
# of printers added to warranty population each month	÷ 105,000
Average warranty cost per printer for ALL repairs	$2.62

Monthly cost for FORMATTER-related warranty repairs	$65,000
# of printers added to warranty population each month	÷ 105,000
AVERAGE warranty cost per printer for formatter repairs	$0.62

Based on this calculation, we can conclude that the maximum value of a latent defect program for formatters should be no more than $0.62 per production board (or no more than $2.62 for an effective program for the entire printer), because that is the cost for us to wait until the board is in the field to repair it. However, one important factor this calculation leaves out is the satisfaction value for a customer not seeing the repair at all. This implies that we should be willing to pay a premium to detect defects *before* they reach customers. The true threshold cost should therefore be higher than this. This analysis also assumes that a test program worth $0.62 should be capable of eliminating every failure which would occur in the field. Realistically, this is a dangerous assumption, since a test program implemented anywhere during production will be unable to screen for damage which would be caused once the printer leaves the manufacturer.

We also make the assumption here that the population of products and failures in the field over time is fairly constant in nature. If the available data for a product comes from near the beginning or the end of its sales in the field, we need to break down the monthly data in more detail to match up the printers entering the field with the times at which they require repair.

5.4.2.2 Costs At Downstream Stages

Our assumptions regarding the nature of the defects will affect the method we use to calculate the benefits of testing. The method in this section is most appropriate if we are testing every production product, as opposed to sampling. In this case, the primary focus is on catching each failure, and determining root cause and learning from it is secondary. To calculate the effectiveness of the test, we should follow these three steps:

(1) Pick a testing method and estimate how the defect ppm rates at each stage would change or decrease based on available data or trial tests.

(2) Estimate the cost of repairing a board at each stage of the production process: within the factory, at final assembly, and at the end user.

(3) Compare the cost calculated in (2) with the expected cost of performing the chosen tests to determine feasibility.

For example, consider the formatter boards we have been discussing in this thesis. Chapter 4 presented the results of visual inspection and passive burn-in on a sample of about 20,000 formatter boards. We would expect that the failed boards form our experiments represent products which would have failed at a later time. Therefore, we can estimate where these boards would have otherwise failed, and count the avoided cost of that failure as a benefit of the test. The figures below illustrate this process:

Type of testing:	visual	passive burn-in
Resulting fail rates from experiments:	1,290 ppm	665 ppm
Estimated potential defect ppm reduction at integration:	300 ppm	250 ppm
Estimated potential defect ppm reduction in the field:	500 ppm	350 ppm

Note that the full amount of the visual inspection fail rate is not distributed among the downstream stages, since it is likely that some of the boards which failed visual inspection would not correspond to an actual field failure. With better feedback from the field on failure modes, it may be possible in the future to better design visual tests which more closely correspond to true latent defects. On the other hand, since the passive burn-in provided no more stress to the boards than what they should be able to endure in the field, we can be fairly confident that most all of the failures here represent boards with latent defects which would fail at some later stage of assembly or use.

Next, we estimate the costs of repairing or replacing boards at each stage of assembly:

Average cost to repair a formatter board in the factory:	$30.00[7]
Average cost to replace a formatter board at integration:	$50.00[8]
Average cost to replace a formatter board in the field:	$250.00[9]

Over a population of 1,000,000 boards, we can multiply the defect reduction rates by the cost that would have been incurred to do those repairs to determine the total benefit of this testing:

$$\$50.00 * (300 + 250) + \$250.00 * (500 + 350) = \underline{\$240,000}$$

To understand if the testing is beneficial, we compare this result to the cost of doing the testing and of repairing the boards at the formatter factory as they fail the test. We estimate that this cost for 1,000,000 boards would be $980,000.[10] Based on these initial results, we conclude that this particular latent defect test process is not cost effective. Other possibilities to investigate would be the effectiveness of monitored testing or accelerated testing. Since these methods provide more information on relevant failure modes per tested board, they may be more likely to generate a greater defect reduction at a lower net cost per board.

---

[7]Includes costs of failure analysis and repair.
[8]Includes costs of board transport to and from assembly site, time needed to replace the board in the final printer, and the costs of failure analysis and repair of the board itself.
[9]Estimated cost to H-P for warranty repair in the field.
[10]Includes cost of fixturing, labor costs for installing, removing and visual & functional testing the boards, value of work-in-process inventory, and cost to repair failed boards.

5.4.2.3 Probability of Class Failure

The above justification method is suited to a testing method which checks 100% of all parts. On the other hand, we may find that we can learn enough from a sampling of boards to eliminate defects. This is possible when failures have root causes which can exist in other parts as well. In this case, we can determine the effectiveness of our test process by estimating the probability that the test would be able to detect and catch a major class problem, and then determining what that capture will save us in downstream failures. In our calculations, we can think of each individual defect mode as existing throughout the population as a binomial distribution. Our test effectiveness will thus depend on the probability of detecting a class defect that is present on x% of the boards based on a random sampling of them. For example, with 90% confidence, a test of 2,000 boards will detect all defects which occur in 0.115% or more of the population. With these assumptions, we can calculate the total cost of quality as the sum of the testing costs and the expected cost of the downstream failures our test will *miss*:

COST = TEST COSTS + E(FAILURES)

where TEST COSTS = (# of boards tested) * (cost of test per board)
and E(FAILURES) = (significance level) * (probability of failures escaping the test) * (total value of boards affected)

For example, based on rough studies of past formatter products, we can use the following data for a sample calculation:

# of boards to be tested each month = 2,000 boards
total cost of testing each board = $400 (value of board, test fixtures, analysis)
significance level for 90% confidence = 0.1
total value of boards affected = $75,000,000
This value is the product of the following:
monthly production = 125,000 boards/month
# of months needed to detect problem = 3 (due to inventory in pipeline)
warranty cost to repair each board = $200

We can calculate the total cost of quality as:

2,000*($400) + (1 - 0.9)(.00115)(125,000)(3)($200) = $808,625, or $6.47 per production board.

The vast majority of this cost is the cost of test, not the potential loss in the field. To reduce the cost per board, we could choose to only test the 2,000 boards every six months (every 750,000 boards) instead of every month. In that case, the total cost of quality would be:

2,000*($400) + (1 - 0.9)(.00115)(125,000)(8)($200) = $823,000, or $1.10 per production board. Note that this calculation accounts for the possibility of building an additional five months of boards with the defect due to the longer interval between tests. This causes the total cost to increase, but we are spreading this cost over six times as many boards now, making the cost per board much less.

To apply this approach with more accuracy, we should determine the "cost of test per formatter" as a function of the number of boards to be tested, since this would vary with the size of the test. By experimenting with various test plans, we can determine what values would minimize the total cost formula. To evaluate if that cost is justifiable, we must compare that optimal result to the savings which would be achieved by eliminating this the expected failure classes from the population.

5.4.2.4 "Insurance Policy" Approach

Some manufacturers have justified their latent defect tests by regarding them as protection against potential problems, without being entirely sure that these problems may be in the future. Viewing testing as insurance against major class problems involves the same conceptual view of defects we applied in the previous method: defects occur in groups due to assignable causes, rather than as individuals. Valuing testing in this way involves going through the same thought process that we would use when purchasing insurance. For example, based on anecdotal evidence. It is not unreasonable to see a class problem escape the factory which results in millions of dollars of repair costs in the field. If we had a test system which could detect these types of defect modes, the value we place on that test would be a function of the expected cost of the defect mode escaping to customers and the probability with which that would be expected to

occur. For example, if we found from our experience that there was a 5% chance each year of a major class problem which costs us over $2,000,000 to control, we should be willing to pay

$$.05 * \$2{,}000{,}000 = \underline{\$100{,}000}$$

for the annual cost of a screening system to guard against these sorts of problems.[11]

5.4.2.5 Cost Minimization Model

Finally, Jensen and Petersen propose a model for latent defect testing costs which uses information about costs and failure levels at each stage of the test and failure process. To apply this model, we must determine or estimate the following costs:

* Burn-in constant costs (BICC) - fixed costs of the test which are independent of the number of products being tested or the length of the test
* Burn-in failure costs (BIFC) - cost of a product failing during the test, including handling and repair costs
* Burn-in time-dependent costs (BITC) - costs of inventory tied up by testing process, cost to run equipment, and cost to monitor equipment and failures
* Customer repair costs (CRC) - cost to the manufacturer for repairing a consumer's product in the field
* Loss of goodwill (LG) - estimate of the intangible cost of a dissatisfied customer due to the fact that his product needed to be repaired

To calculate the total costs associated with burn-in (TCBI), we would sum these costs as shown:

$$\begin{aligned} TCBI = BICC &+ (\#\text{ of failures during testing}) * BIFC \\ &+ (\#\text{ of days each product is tested}) * BITC \\ &+ (\#\text{ of failures after test, at customer}) * (CRC + LG) \end{aligned}$$

To evaluate this total cost, we compare it to the cost of doing *no testing* (NTC) at all; the "cost" of this approach would be the extra repairs we would have to do in the field:

$$NTC = (\#\text{ of failures without testing}) * (CRC + LG)$$

---

[11]Once we have established the value of this sort of approach, we must also consider the nature of the expected defects, as discussed in Sec. 5.4.1. If we are not able to detect common types of defects or if we cannot eliminate future defects based on our knowledge of current defects, we are forced to test 100% of our products. From our research, we found that manufacturers using this sort of approach find that catching one defect allows them to learn how to prevent many more from occurring. This allows them to settle for sampling the products as opposed to 100% screening.

Once each of these totals are converted into common units, such as cost per product, we can look at the difference, TCBI - NTC, to determine the net cost of testing. If this difference is positive regardless of what value is used for the length of test, it will always be cheaper to put up with the higher failure rate in the field. If a value can be found for the length of test which makes the TCBI less than the NTC, there is a net cost benefit to performing the testing.[1]

5.5 Summary

Comparing the costs and benefits of testing requires three basic steps. First, we must gather as much detailed information as possible for tangible factors and estimate the values of intangible factors as fairly as possible. Table 5.1 review the major types of costs and benefits we should consider.

	Tangible	Intangible
Costs	inventory equipment utilities/support personnel	effect on cycle time effective lifetime of test customer satisfaction
Benefits	failure rate reduction warranty	learning customer satisfaction

Table 5.1 Summary of testing costs and benefits.

Second, before applying a decision model, we must review our assumptions about test quantities and the nature of the defects we expect to find. These will determine what sort of analysis is most applicable. Finally we use a justification model which takes into account the assumptions we have made and which is best suited to our situation, such as those detailed in this chapter. These models can be based on the following types of calculations:

* cost of alternative (re-engineered) methods for product service
* expected maximum warranty cost savings
* expected savings in downstream repairs
* the probability of detected problems occurring in a large number of boards
* the value of "insurance" to prevent class failures from reaching customers
* the fixed and variable costs of testing as a function of length of test

References

[1] Jensen, Finn, and Niels Petersen, Burn-in: An Engineering Approach to the Design and Analysis of Burn-in Procedures, John Wiley & Sons, Chichester, England, 1985, pp. 91-93.

Chapter 6 - Conclusions

Latent defect testing is an important but complex issue for the production of electronic components. Today's products and business competition have made it necessary to move beyond the time-independent testing of the past. Current types of testing are often incapable of detecting the defect modes which lead to product failures after time and stresses have taken their toll. In this thesis, we used the framework shown in Fig. 6.1 for investigating and making decisions on latent defect issues.

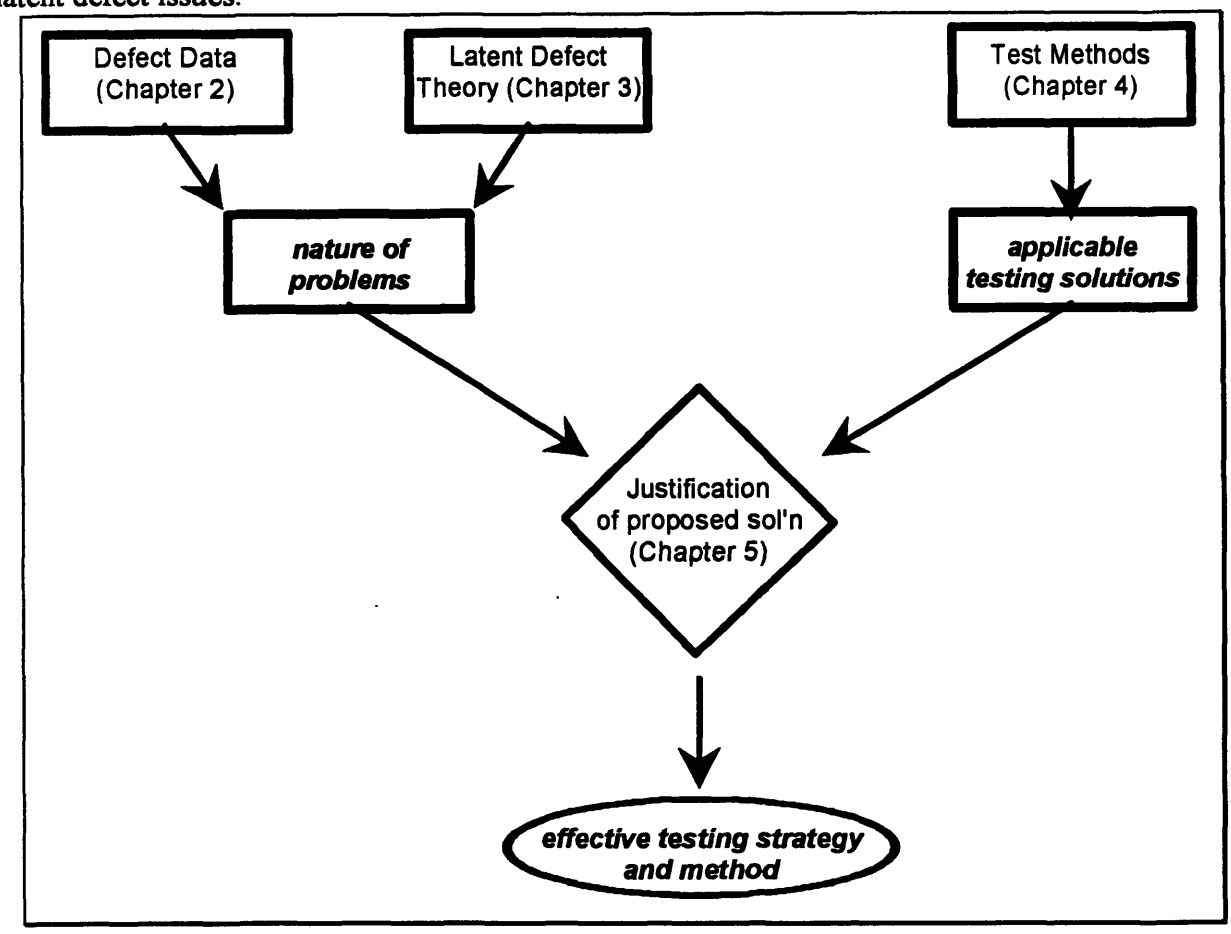

Figure 6.1 Review of latent defect testing decision process.

Throughout this thesis, we used the case of the formatter boards in H-P's laser printers to illustrate the application of our ideas and concepts. For this particular situation, our analysis has led us to the following conclusions:

(1) Board quality is limited by the accuracy of current testing methods. From analyzing the boards presented in Appendix A, we found that the test procedures throughout the production process limited the accuracy and reliability of the test data, due to the already high levels of quality in the products. Evidence of these limitations include the following:

-- false pass rates of ~200 ppm during board assembly
-- false pass rates of 800-1,000 ppm during board assembly
-- fail modes during printer assembly which can be traced back to the same failure modes occurring during that time period at board assembly
-- the relatively large number of board which fail printer assembly, but for which no failure mode can be determined

Recommendation: Test for the same failure modes during various stages of production, and continue to improve the reliability of the tests. Redundant downstream testing makes it possible to determine when failures are being missed by an upstream process. This knowledge allows engineers to redesign tests or procedures to detect these defects earlier in the process. This results in cost savings due to increased quality, and it also allows us to avoid building and adding value to products which will need to be repaired or discarded later in the process.

(2) Field data does not allow us to learn about specific failure modes. Although H-P's field data process makes it possible to look at general warranty repairs in many useful ways, it does not contain enough information about reasons for repairs or failure modes to let us draw conclusions in more detail. In addition, H-P's field data is only reliable during the one-year warranty period for the printers which makes it difficult to quantify long-term reliability. This is a significant impediment for a company like H-P whose success is due to a large part to its reputation for product quality.

Recommendation: Continue with periodic detailed studies of boards returned from the field, or gather more information during the regular field repair process. The current field information needs to be augmented to enable continuous improvement of the products. This can be done by increasing the amount of information available from every repair, and also by gathering very detailed information from samples of boards from warranty and especially non-warranty repairs. Types of information which would be helpful for root cause analysis and for future design improvements include the following:

-- number of pages printed at the time of repair
-- the contents of data registers at the time of the failure
-- the serial number of each formatter board involved in a repair
-- the nature of the failure observed by the customer (consistent, intermittent, rare)

(3) Formatter boards exhibit a decreasing failure rate over time. The bathtub curve described in Chapter 3 does not entirely fit our data for the formatter boards, but some of the concepts do apply. For example, our data does not show a constant life period or a wearout period for the formatter itself. On the other hand, the decreasing slope of the data fits the

findings of Wong and others on the nature of failures in modern electronic products. The slope also implies that a test which simulates use of the product for a given amount of time would be effective in eliminating products which would otherwise fail in the field.

(4) Passive burn-in testing is not an effective test technique for learning about formatter board failure modes. Although our testing in Chapter 4 did detect a number of detects, we could not determine a root cause for most of the failures. In addition, passive testing is unable to provide us with repeatable information on intermittent defects. This makes it impossible for us to rely on this sort of testing for anything more than a 100% screen for defective board. This impacts the type of financial analysis we are allowed to use in evaluating the test's effectiveness.

Recommendation: Investigate testing methods or analysis procedures which provide more information on each failed board, such as accelerated stress testing. The other testing methods we discussed in Chapter 4--monitored burn-in and accelerated stress testing--allow us to more consistently detect failures in boards. Therefore, they are more likely candidates for effective testing methods in production. In addition to better failure detection and precipitation, we also need to improve our techniques for analyzing the root causes of failures. We can do this by working with vendors to develop test programs which better determine failure modes within individual components. We also need to better understand the effects of each type or combination of stresses on the boards during accelerated stress testing. We also need test programs which quickly detect the resulting failures as they occur, due the intermittent nature of many of the failures. Engineers at H-P are currently studying many of these possibilities in more detail.

(5) Passive burn-in is not a cost-effective method for latent defect testing of the boards. Since we were unable to use the results of the passive burn-in experiments to learn about failure modes across other products, we must evaluate the costs and benefits of passive burn-in by assuming that failures can only be detected and eliminated individually. As a result, the cost-of-downstream-stages method in Chapter 5 best fits our situation. As we showed there, the tangible costs of the testing were approximately $740,000 more than the expected tangible benefits for an expected production level of one million boards. Our estimates of the values of the intangible costs and benefits were not enough to change this conclusion that passive burn-in would not be efficient in production.

Recommendation: Investigate other testing methods. As we discussed in the fourth point, other techniques may be more effective in helping us eliminate latent defects. Especially as we develop new products and make use of new technologies, we must continually develop new testing methods which are adapted to our current needs. For example, engineers at H-P are exploring technologies such as boundary scan, which make it possible to run detailed tests at the component level as well as at the board level. This will be a great help in analyzing the root causes of component-level failures, which was a major source of confusion in our board-level tests.

(6) The change in printer assembly processes alone does not justify passive burn-in testing.

Finally, we address the initial question which formed the basis for this project: Should we implement additional tests to account for the reduction in testing and power cycling of the printers? Through our collection of data and our experimentation with various techniques, we found that the change in detected failures at the printer assembly sites is not a sufficient cause for implementing a new testing procedure in production. Our investigation, however, has pointed to additional areas which we should continue to study and learn about:

-- the use of accelerated stress testing at other stages of the product life cycle in addition to production testing:
  -- initial design of the formatter board
  -- root cause analysis of failed boards from the field
-- gathering of field data on products assembled with the newer H-P production process, as a final check of the impact of the process change on product quality

In summary, our broad goals in this project and thesis are to understand and compare the potential *problems* and *solutions* for latent defects. In understanding the *problems* caused by latent defects, we have two sources of information available to us. First, there is the data from the production process itself. Although each step of the process may give us only time-independent information, we can assembly data from each step of the process or we can look at one test process over time to understand the time-dependent behavior of our products. Second, the theories and models for latent defect behavior help us identify trends in the information we have collected, and improve our awareness of the underlying behaviors which cause the failures we desire to prevent. Once we have this understanding of the problems, we investigate the range of possible *solutions* or testing methods: visual inspection, burn-in, and accelerated stress testing. We can choose which method may best suit our needs through our own experiments and by reviewing the experiences of others in selecting and implementing these types of tests. After we determine the problems and solutions, our final task is to compare the costs of the solution to the benefits we expect it to provide in solving latent defect problems.

Appendix A - Breakdown of Available Formatter Defect Information

Explanatory Notes

Table A.1 provides information detailed information about the types of defects which occur in formatter boards, and at what levels they occur in each stage of production. Each column contains the following information:

Type of Defect - a brief description of the defect mode

Board Level - the number of failures at the formatter factory which can be attributed to each failure mode. This data comes from the functional testing and visual inspection normally done for every board during production.
Printer Assembly - Canon - the number of failures during the printer assembly process at Canon which can be assigned to each formatter fail mode. Engineers and Canon and H-P gather this information from the Canon assembly lines on a daily basis.
Printer Assembly - H-P - the number of failures during the printer assembly process at Hewlett-Packard printer integration/assembly sites which can be attributed to each failure mode. This information is also collected from failure information on the assembly lines by H-P engineers.

Field Failure Analysis - the percentage of failed boards from the field which can be attributed to each defect mode. This data is collected periodically for each product as part of a special study of formatter boards which have been replaced in customers' printers. Engineers at H-P will collect failed boards from the service centers and analyze them in detail, rather than allowing them to go through the normal repair process which does not provide as much information back to engineering regarding the failure modes.

The *units* for the three failure levels (board, Canon, H-P) are *parts per million* boards, to allow comparisons across each stage. For example, a 3,479 ppm rate for capacitors means that out of a million boards, there were 3,479 total defects due to capacitors. Multiple failures of a single type on a single board are counted separately.

We used data from *three different* types of *boards* in creating this table. The boards are roughly similar in terms of the types and number of each part they contain. Many parts are duplicated on each board. The numbers in brackets [xxx] beside each failure mode indicate the number of each component on a K9 board. For example, the K9 board contains 35 resistors and 62 capacitors.

Parentheses (xxx) around a number indicate that those failures are also accounted for in another category and are thus not double-counted in the totals for each column. For example, the 581 ppm rate for DRAMs during board assembly is due to insufficient solders, misinserted parts, damaged parts, etc.

We use this data to draw conclusions about the following:

- the effectiveness of each test in detecting all failures of a given type
- the potential of the assembly process to create or precipitate defects
- the change in detectable failures between the Canon and H-P processes

Chapters 2-4 use the data in this table to illustrate these and other issues.

Type of Defect	Board Level (ppm)	Printer Assembly Canon (ppm)	Printer Assembly H-P (ppm)	Field Failure Analysis
MECHANICAL				
damaged parts	4,008			
missing parts	16,182			
missing nut	0	768		
missing bail	0	279	349	
missing/loose screw	(3,204)		14	
wrong nut	0	136		
misinserted part	8,389	97	14	
wrong part	2,034	47		
cracked shroud	0	43		
damaged bail	0	43		
contact problems	7,502			
insufficient solder/no solder	6,601	202	852	
excess solder	0	4		
damaged housing	0	31	70	
damaged nut	0	27		
missing connector	0	19		
damaged/defective circuit board	(5,357)	19	(14)	
bent contacts/connector pins	9,864	8	56	
foreign material	931			
raised connector	5,267			
connector pins short	0		14	
damaged ferrite bead	(1,438)		(14)	
excess solder/solder shorts	8,821			
shipping damage	0			
serial connector	(9,588)		(112)	
font connector	(5,483)		(42)	
Centronics connector	(6,795)		(782)	
SIMM connector [3]	(7,197)			
engine connector	(2,146)		(419)	
RS232 tranceiver	0			
I/O connector	(2,168)		(587)	
bar code label	(745)			
TOTAL MECHANICAL	69,599	1,723	1,369	0%
ELECTRICAL				
defective parts	6,944			
NVRAM-contents-wrong (Canon only)	0	85		
front panel keys not working (Canon only)	0	39		
oscillator	(4,887)	757	14	50%
NTF	0	372	1,470	18%
DRAM [4]	(581)	35		2%
80960	(2,459)	35		2%
ASIC (wrong part/failed test)	(6,757)	31	513	10%
buffer/communication failure	0	8	447	
EEPROM	0	8		2%
reset-IC	0	4		1%
micro-wire-short	0	4		
ROM [2]	(1,177)	4		1%
process-problem	0	4		
address/latch buffers [12]	(5,059)			
diodes [23]	(2,995)			
NVRAM	0			
UART/tranceiver	(1,304)			
resistors [35]	(3,904)			
capacitors [62]	(3,479)			
TOTAL ELECTRICAL	6,944	1,386	2,444	100%
TOTAL PPM RATE	76,543	3,109	3,813	100%
data source	LaserJet K9	LaserJet W8	LaserJet K9	LaserJet B2
sample size	134,266	257,723	71,604	300 bds. analyzed

Table A.1 Breakdown of Available Formatter Defect Information at Each Stage of Production.

Appendix B - Bullet Hole Method for Charting Factory Defects

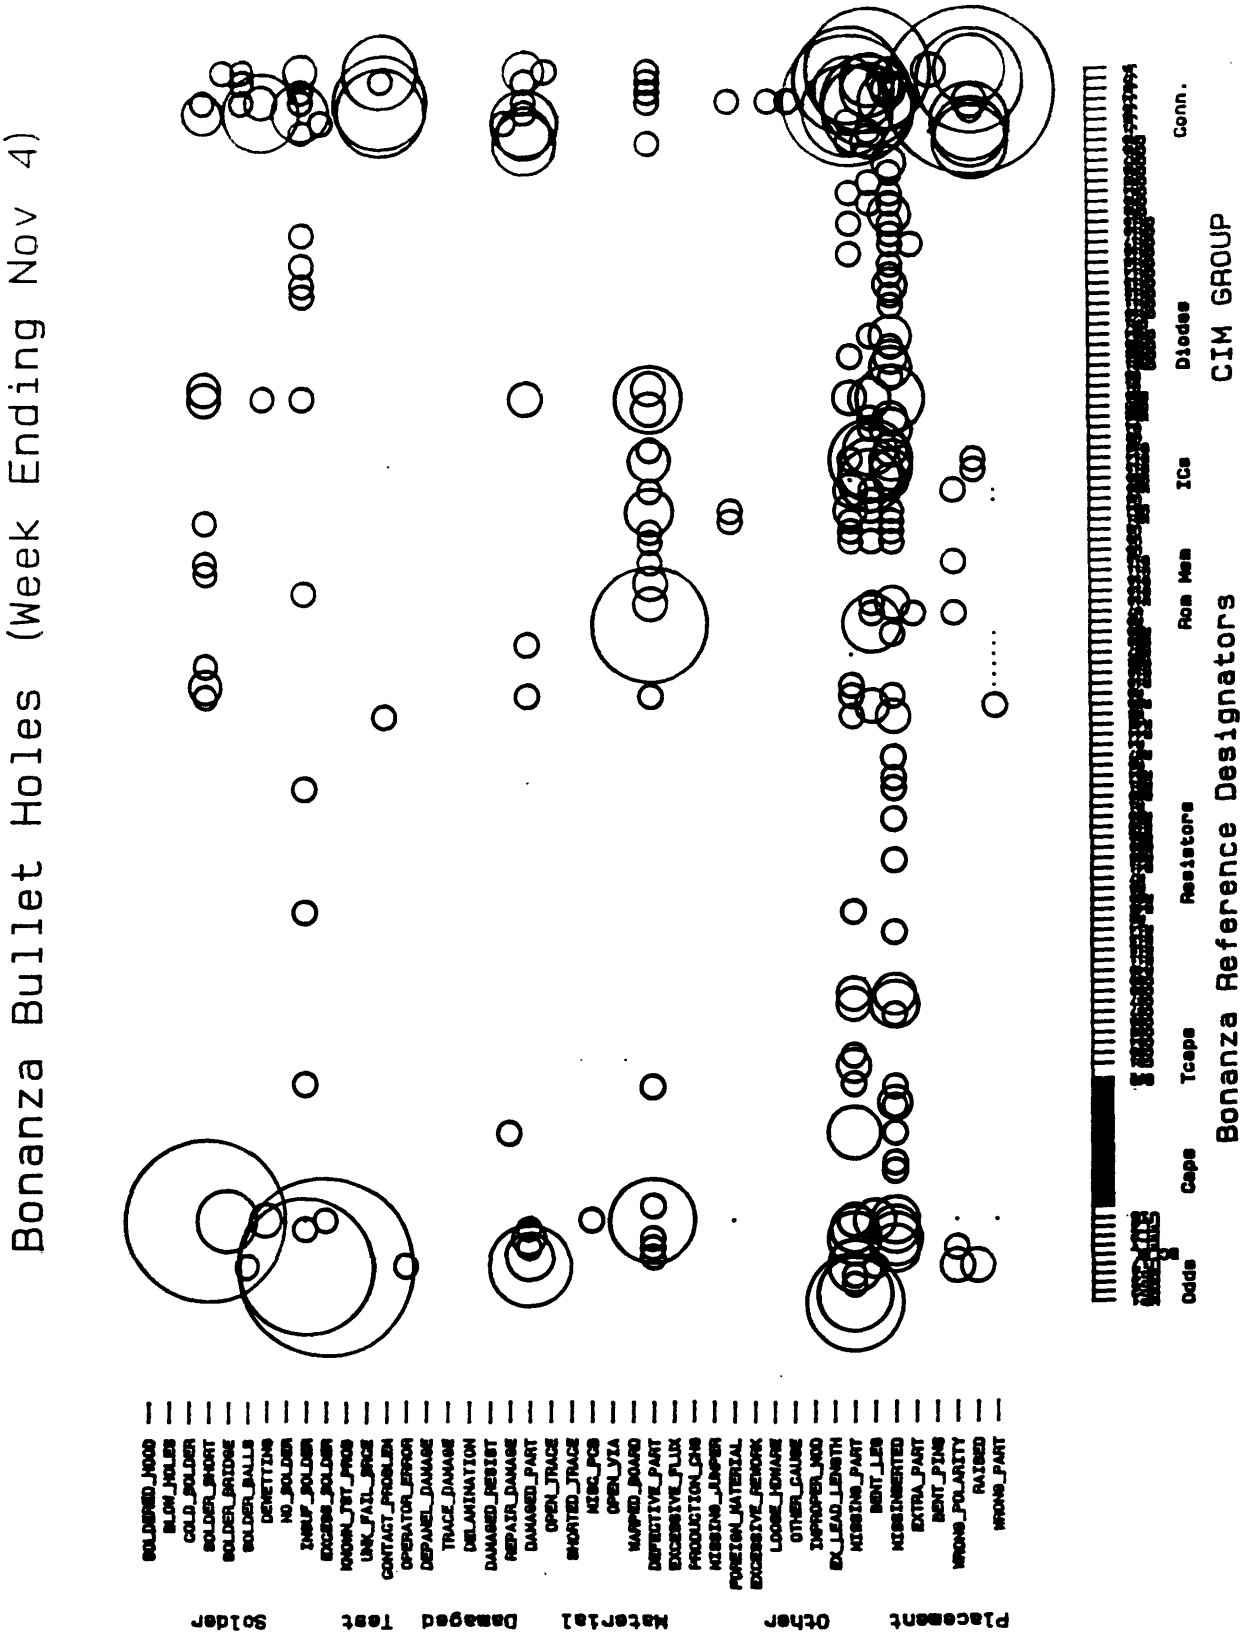

Appendix C - Details of Accerated Stress Testing Experiments

Background information for the four boards which were tested:

Serial #	Details
AJ2X	Passed passive burn-in test ~11/29
7AK7	Failed 10/7/93 9:10PM for SY_AKTO
	Failed 10/13/93 11:14AM for PRNTST_4 BAD CRC
	10/13 - formatter board was attached to a printer at PCL, and printer self-test was OK; when print test was done with board fixed in hand-crank tester at PCL, a 55 error occurred; the board also passed "terminator" testing
AJ13	Passed passive burn-in test ~11/29
6QCH	Failed 10/1/93 5:32PM for SY_AKTO
	Failed 10/4/93 10:54AM for TIMEOUT
	10/7 - passed all hand-crank and pneumatic testing at PCL; when board was taken back to hand cranks in run-in area, it now passed (10/7 3:04 PM)

Details of Testing Procedure

Note: During all thermal cycling, temperatures were raised or lowered as fast as the chamber would allow; ramp rates of >70C°were noted, and less than five minutes was required for the chamber to go from -100°C to +130°C. Also, all temperatures shown represent actual board temperature, not temperature in the chamber; temperature was measured with a thermocouple attached to the mounting hole nearest the engine connection.

Board AJ2X, continuous print testing (patterns from PC), no vibration

Board was cooled to -20°C and held there for ~5 minutes, then successively cooled in steps of -10C° at a time for 5 minutes each until the limit of the chamber was reached (at around -105°C). It was then cycled through the following profile: -100°C⇒+80°C⇒-100°C⇒+90°C⇒-100°C⇒+100°C; each temperature was held for about 5 minutes each before moving on. From +100°C, the temperature was increased 10C° at a time until a 79 SERVICE error was detected on the printer at +128°C, which stopped all printing. After power cycling the printer, a 63 NEEDS SERVICE error was reported, and printing was still stopped. The temperature was ramped down to -100°C and back up to +128°C when the same error occurred again. From here, temperature was gradually decreased and the printer was power cycled to attempt to restore function; the printer first responded with 8 black characters, then with a 05 SELF TEST message for about 0.5 seconds before switching to 63 NEEDS SERVICE error, then with 05 SELF TEST for about 2 seconds before switching to 63 NEEDS SERVICE, before finally at about 117°C, the printer was able to pass its own self test and continue printing. Temperature was then increased again; the same failure occurred at +128°C, and function was restored at +123° C this time.

Board AJ2X, DFTS tests, no vibration

The SIMM cards for the BEST-style testing and the related wiring were inserted, and temperature was increased from room temperature. At +127°C, a repeatable OBRAM_7 failure occurred at address 010F327C; the test was able to pass when temperature was reduced to +126°C.

Board 7AK7, continuous print testing (patterns from PC), no vibration

Temperature was reduced directly to -70°C and cooled in -10C° increments down to about -102°C, with five minutes dwell time allowed at each level. The board was then thermal cycled ⇒-100°C⇒+90°C⇒-100°C⇒

+100°C, then temperature was increased until the 79 SERVICE error occurred at +128°C. The same error messages as described with board AJ2X then occurred as temperature was gradually reduced to restore function; at +123°C, operation was restored. The board was cooled to room temperature to install the DFTS test wiring and SIMMs.

Board 7AK7, DFTS tests, no vibration
Temperature was increased until a NVRAM1_L error was detected at +128°C; the test continually failed as temperature was gradually reduced until it was able to pass again at +121°C. The board was then cycled down to -100°C and back to room temperature with no problems; all tests continued functioning.

Board 7AK7, DFTS tests, vibration
All tests were running; with 5G (at 6 d.o.f.) of vibration, board was cycled from +25°C⇒-50°C⇒+100°C⇒+110 °C with about 5 minutes dwell at each level. As temperature was increased from there, the NVRAM1_L repeatedly occurred again, this time at +116°C. Temperature was ramped down to -50°C, and all tests resumed functioning. Vibration was increased to 15G, and temperature was raised -50°C⇒+100°C⇒+110°C; as temperature was increased to +120°C again, the NVRAM1_L error recurred repeatedly at +116-117°C. Temperature was ramped down to -50°C, and all tests resumed functioning. Vibration was increased to 40G, and one cycle of all the tests worked fine. After that, though, RFRSH1_1 and SIMRAM_1 errors appeared and were consistent, even as vibration was gradually reduced back down to zero. Temperature was raised to room temperature, and all errors disappeared. Temperature and vibration were then increased to +106°C and 30G, and the RFRSH1_1 and SIMRAM_1 errors reappeared and were consistent as the board was cooled back down to room temperature.

Board 7AK7, continuous print testing (internal continuous self-print test), vibration
At 30G vibration, the board was taken to +100°C⇒+110°C⇒+120°C⇒+130°C; the only problems noticed were that the bottom line of the border box would be missing--this happened more consistently as temperature was increased. At +136°C, a 79 SERVICE (01bf) error occurred, and the board had to be cooled to +122°C before function was restored. (The same pattern of black squares, 05 SELF TEST, and 63 NEEDS SERVICE errors was detected as with the first testing on board AJ2X; see previous page for details.)

Board AJ13, continuous print testing (internal continuous self-print test), no vibration
Temperature was taken from room temp. directly to +130°C. At +131°C, a "79 SERVICE (3a10e)" error occurred. The printer was power cycled and temperature was gradually reduced until function was restored at +116°C. Temperature was taken back to room temp.

Board AJ13, continuous print testing (internal continuous self-print test), vibration
At 10G of vibration, temperature was increased. After about 2 minutes at +129°C, a "79 SERVICE (01bf)" error occurred. Temperature had to be reduced to +118°C before function was restored. Temperature was lowered to +25°C and vibration was increased to 20G. As temperature was again increased up to +130°C, the bottom lines on the border of the print page were sometimes entirely missing or partially missing. Vibration was increased to 30G, and after about 2 minutes the "79 SERVICE (01a1" error reappeared at +128°C. Temperature was reduced to +122°C before function was restored. Temp. was reduced to room temp., vibration was increased to 40G, and temp. was again increased. The "79 SERVICE (01bf)" occurred at +129°C, and function was restored at +123° C.

Board AJ13, continuous print testing (internal continuous self-print test), vibration, thermal cycling
At 40G vibration, temperature was cycled +130°C⇒-80°C⇒+130°C four times, with about five minutes dwell time at each extreme. In each cycle, the "79 SERVICE" error appeared in the range of +125°C to +130°C, and

everything else functioned OK except for occasional bottom lines missing off the print border. After the four cycles, temperature was reduced to room temp. to change boards.

Board 6QCH, continuous print testing (internal continuous self-print test), vibration, thermal cycling
At 40G vibration, temperature was cycled +130°C⇒-80°C⇒+130°C four times, with about five minutes dwell time at each extreme. No front panel errors occurred during the first two cycles, but about 1/2" of printing was "smeared" during each drop in temperature, at around +90°C to +100°C. On the third and fourth cycles, the "79 SERVICE" errors occurred again, and a couple data transmission errors occurred again as well, resulting in "smears". During these last cycles, printing could not be resumed until temperature was reduced to around +65°C or +75°C (but temperature was not gradually lowered here as it was in earlier tests to restore printing).

Board 6QCH, DFTS tests, vibration
At 40G of vibration, temperature was increased to +130°C. A SIMRAM_4 error occurred at +75°C, and OBRAM_1, SIMRAM_1, and REFRSH_1 errors occurred repeatedly at around +130°C. Temperature was gradually reduced; the SIMRAM_1 errors disappeared at +120°C, the REFRSH_1 and OBRAM_1 errors disappeared at +116°C, but SIMRAM_2 and SIMRAM_6 errors began recurring at +116°C as well. At +110° C, the REFRSH_1 errors stopped, but then RFRSH1_1 errors began occurring. The RFRSH1_1 errors and SIMRAM_1 errors were then repeatable all the way back to room temperature and with no vibration, but when the board was testing in a functional tester the next day, no problems were detected.

Appendix D - Analysis of Precipitation and Detection Methods for Board Defects

	Precipitation method	Detection method	Normally detected at Canon printer assy?	Normally detected at H-P printer assembly?	Prevention method (?)
Mechanical fail modes					
missing nut	vib (& temp extreme)	visual inspection	yes	if connector were used	shipping redesign/vision system
loose/missing bail clip	vib (& temp extreme)	visual inspection	maybe	less likely	shipping redesign/vision system
wrong nut	vib (& temp extreme)	visual inspection	yes	yes	vision system
misinserted (180 deg. out) connector	visual inspection	visual inspection	yes	yes (serial conn. can be flipped)	keyed connector/vision system
wrong connector	vib (if size difference), powering (to force it to misfunction)	connector test	yes	N/A for models assembled at H-P sites	vision system/add'l functional tests
cracked shroud on connector	vib, temperature cycling to accelerate cracking	visual inspection/ connector test	not always	not always	shipping redesign/vision system
insufficient solder	vib w/temp ramping, humidity	functional test	not always	not always	solder process monitoring
damaged housing (on 8,10-pin or port connector)	vib in shipping simulation	vision system/visual inspection	not always	not always	shipping redesign/vision system
damaged nut (on Centronics connector)	vib in shipping simulation	vision system/visual inspection	yes (on Payette)	not always	shipping redesign, incoming inspection
missing connector	vib or visual	vision system/visual inspection	yes (more so than at integ.)	not always	vision system/in-circuit test check
damaged PCB	vib (in shipping simulation)	fcn. test/in-circuit test	more so than at integ.	not always	shipping/depaneler redesign; incoming inspection
bent contacts in connector	vib, temperature cycling	vision system/visual inspection	not on SIMMs or I/O; those are not checked	only on mating ones (not SIMMs)	more robust contact/box redesign?
Electrical fail modes					

Appendix D - Analysis of Precipitation and Detection Methods for Board Defects

oscillator	power cycling thermal cycling	functional test/in-circuit test	more than at integ.	less (would get solid fails, but less intermittents)	???
NTF	vib/thermal cycling	functional test/in-circuit test			???
DRAM	thermal cycling, moisture, power	functional test (during or after stress testing)	more than at integ.	less (would get solid fails, but less intermittents)	screen/sample (if sporadic) test or process redesign (if class failure)
MILFORD	thermal cycling, moisture, power	functional test (during or after stress testing)	more than at integ.	less (would get solid fails, but less intermittents)	screen/sample (if sporadic) test or process redesign (if class failure)
buffer	thermal cycling, moisture, power	functional test (during or after stress testing)	more than at integ.	less (would get solid fails, but less intermittents)	screen/sample (if sporadic) test or process redesign (if class failure)
EEPROM	thermal cycling, moisture, power	functional test (during or after stress testing)	more than at integ.	less (would get solid fails, but less intermittents)	screen/sample (if sporadic) test or process redesign (if class failure)
micro-wire-short	vib/thermal cycling	functional test (during or after stress testing)	more than at integ.	less (would get solid fails, but less intermittents)	???
excess solder	vib/thermal/power cycling, humidity	functional test (during or after stress testing)			vision system/process controls
ROM	thermal cycling, moisture				screen/sample (if sporadic) test or process redesign (if class failure)
Notes:					
"vib" = vibration testing					
"How to prevent" measures assume that product redesign is not an option.					