Redefining Manufacturing Quality Control in the Electronics Industry

by

Maureen Fresquez Simington

B.S., Mechanical Engineering, New Mexico State University, 1992

Submitted to the Sloan School of Management and the
Department of Materials Science & Engineering
in Partial Fulfillment of the Requirements for the Degrees of

Master of Science in Management
and
Master of Science in Materials Science & Engineering

in Conjunction with the
Leaders For Manufacturing Program

at the
Massachusetts Institute of Technology
June 2000

©2000 Massachusetts Institute of Technology
All Rights Reserved

Signature of Author ______________

MIT Sloan School of Management
Department of Materials Science & Engineering
May 5, 2000

Certified by ____

Thomas W. Eagar
Professor of Materials Engineering
Thesis Supervisor

Certified by ____________________________

Roy E. Welsch
Professor of Statistics & Management Science
Thesis Supervisor

Accepted by ___________________

Carl V. Thompson
Stavros Salapatas Professor of Materials Science & Engineering
Chair, Departmental Committee on Graduate Students

Accepted by _______________________

Margaret Andrews
Executive Director of the MBA Program
MIT Sloan School of Management

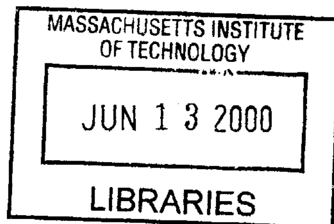
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
JUN 1 3 2000
LIBRARIES

ENG

Redefining Manufacturing Quality Control in the Electronics Industry

by

Maureen Fresquez Simington

Submitted to the Sloan School of Management and the
Department of Materials Science & Engineering
on May 5, 2000 in Partial Fulfillment of the
Requirements for the Degrees of

Master of Science in Management
and
Master of Science in Materials Science & Engineering

Abstract

The most time consuming and capital intensive portion in the assembly of power electronic devices is the test system. A comprehensive test system including functional and stress screening technologies can significantly increase assembly times and more than double the capital investment required in a new assembly line. The primary purpose of the test system is to screen components for early life failures and to verify proper assembly. Determination of key performance characteristics and the resultant test system are developed during the product design phase and are seldom revised after the product has been released to manufacturing. This thesis explores best practices in testing methods and develops new methods to analyze test system performance. Both efforts were conducted in an effort to optimize existing test regimes.

Upon completion of the above analyses the existing test sequence was reduced by 50%. This was primarily due to a discovery in the Burn In test cycle which indicated that failures correlated strongly with the on/off cycles inherent in the test sequence. A new test cycle was proposed to accommodate this finding and test results verified the initial hypothesis. Additionally, the summary of best practices identified new forms of product testing including Highly Accelerated Stress Testing (HAST), moving additional product testing into the development phase consequently reducing testing requirements during assembly.

Thesis Advisors:

Thomas W. Eagar
Professor of Materials Engineering

Roy E. Welsch
Professor of Statistics & Management Science

Table of Contents

List of Figures....6
List of Tables....8
List of Equations....8
Acknowledgements....9
Chapter 1 Introduction....11
1.1 Project Motivation....11
1.2 Project Definition and Resulting Priorities....12
1.3 Approach & Methodology....13
1.4 Project Goals, Measurements, and Deliverables....14
Chapter 2 Competitive Benchmarking....15
2.1 Industry Overview & Market Share....15
2.2 Overview of Subject Companies....16
2.3 Testing at Subject Companies....17
2.4 Summary of Benchmarking Study....26
Chapter 3 Product Line....27
3.1 Product Overview....27
3.2 Product Flow....27
3.3 On-line Production Test System and Database....29
Chapter 4 Hi Pot Performance Analysis....35
4.1. Failure Analysis....35
4.2. First Pass Yield Analysis....37
4.3 Pareto Analysis....39
4.4 Summary of Results....41
Chapter 5. System Tests 1 and 2 Performance Analysis....42
5.1. First Pass Yield Analysis....42
5.2. Pareto Analysis....44
5.3. Summary of Results....48
Chapter 6 Burn In Analysis....50
6.1 Burn In Testing Yields....50
6.2 Failures vs. Time Spent in Sauna....52
6.3 Fault Type Characterization....62
6.4 Test Reduction Hypothesis....65
Chapter 7 Product Performance Analysis....68
7.1. Initial Findings....68
7.2. Performance Analysis....69
7.3. Additional Findings....76
7.4. Recommendations for Future Analysis....77
Chapter 8: Solder Study....79
8.1 Solder Quality....79
8.2 The Soldering System....82
8.3 Soldering – The Existing System....84
8.4 Solder Tip Preventative Maintenance....85
8.5 Summary....92
Chapter 9 Results and Conclusions....95
9.1 Benchmarking Study Conclusions and Summary....95
9.2 Burn In Testing Reduction Results and Conclusions....97
9.3 Concluding Remarks....102
Bibliography....103
End Notes....104

List of Figures

Figure 1 - Market Share of Leading Micro Product Players 15
Figure 2 – Market Share of Leading Macro Segment Players 15
Figure 3 - Manufacturing Product Flow 27
Figure 4 – Burn In Test Cycle 33
Figure 5 - Failure Analysis 36
Figure 6 - Failure Rates vs. Units Produced 37
Figure 7 - Hi Pot FPY 38
Figure 8 - Hi Pot Failure Pareto 40
Figure 9 - Hi Pot Failure Pareto by Test Phase 40
Figure 10 - System Test 1 FPY 43
Figure 11 - System Test 2 FPY 44
Figure 12 - System Test 1 Failure Pareto by Test Phase 45
Figure 13 - System Test 2 Failure Pareto by Test Phase 45
Figure 14 - T1 vs. T2 by Test Phase 46
Figure 15 - System Test 1 Phase B Failure Pareto 47
Figure 16 - System Test 2 Phase B Failure Pareto 47
Figure 17 – FPY Device Size 1 51
Figure 18 – FPY Device Size 2 51
Figure 19 – FPY Device Size 3 51
Figure 20 – FPY Device Size 4 51
Figure 21 – Device Burn In Failures (no bins) 52
Figure 22 – Device Burn In Failures (30 second bins) 53
Figure 23 – Device Burn In Failures (60 second bins) 53
Figure 24 – Device Failures vs. Commanded Speed 54
Figure 25 – Cumulative Failure Rate 56
Figure 26 – F24 Failures vs. Command Speed 58
Figure 27 – F24 as a Percentage of Overall Failures 59
Figure 28 – Failure Curve Without F24 59
Figure 29 – Large Device Internal Temperature vs. Commanded Speed 61
Figure 30 – Control Board Temperature vs. Test Time 62
Figure 31 – Fault Code Grouping Pareto 64
Figure 32 – Projected Cumulative Failures with Spikes Shifted 66
Figure 33 – Comparison Between Existing and Shifted Cumulative Curves 67
Figure 34 - Fault Code as a Function of Frame Size 69
Figure 35 - Test Results: Units that passed all quality tests 70
Figure 36 - Units with 100% Cumulative FPY 71
Figure 37 - Test Results: Units that Failed a Production Test 72
Figure 38 - Units that Failed Production Tests 72
Figure 39 - Failures by Week of Manufacture 73
Figure 40 - Field Failure Percent by Week of Manufacture 74
Figure 41 - Faults by Week of Manufacture 74
Figure 42 - Failure Percent by Week of Manufacture 75
Figure 43 - Soldering Failures by Week 80
Figure 44 - Photo Identifying Solder Joints 81
Figure 45 - Operator Inspection Data 81
Figure 46 - Solder Fixture 82
Figure 47 - Solder Tip Data 86
Figure 48 - IR Camera Output 87
Figure 49 - Tip Temperature Over Time 88

Figure 50 - Initial Sign of Heat........90
Figure 51 - Thermal Signature of Tip at Steady State........90
Figure 52 - Initial Sign of Heat........91
Figure 53 - Stable Temperature........91
Figure 54 - Original 2 Hour Burn In Cycle........99
Figure 55 - New 1.5 Hour Burn In Test Cycle........99
Figure 56 - Failure Rate for New Burn In Sequence........100
Figure 57 - New Burn In Test Cumulative Failures........101

List of Tables

Table 1 – Fault Codes 63
Table 2 - Aggregate Failure Types 76
Table 3 - Benchmarking Results 97

List of Equations

Equation 1 – Statistical Capability Index 18
Equation 2 - First Pass Yield 38

Acknowledgements

First I would like to thank my husband Ken, for without him none of this would have been possible. His love and encouragement sustained me throughout the last two years. Ken has been a constant source of support for me. His patience with my busy schedule and my mood swings was angelic. Thank you so much Ken! I would also like to thank my family for their love and prayers.

Many thanks to my host company. I learned a great deal from the entire staff but especially from Mr. Harri Heimonen, my internship supervisor. I would also like to thank all the wonderful people who helped me survive the 6-month internship. Special thanks to my very good friends: Maria Lindroos, Meri Parm, Mika Winqvist, and Sami Wainio.

Next I would like to acknowledge Roger Maddocks and Charlie Brown from the Eastman Kodak Company for the corporate sponsorship I received during my time in the LFM program.

I would also like to acknowledge the overlap in thesis content between this thesis and that of my internship colleague, Erik Nelson. The collaborative effort throughout the course of the internship was a primary factor in the overall success of the project. Additionally, without the wisdom of Professor Thomas W. Eagar and Mr. Roy Welsch my academic advisors, this thesis would never have come together. Their guidance and advice throughout the internship and thesis authorship was priceless. Thank you.

Finally, much gratitude to Mr. Don Rosenfield and the entire LFM staff for their support over the last 2 years. LFM is an exceptional program and I feel honored to be an official alumnus of the program, MIT and Sloan.

Kokoo kokoon koko kokko	Gather the bonfire together!
Koko kokkoko?	You mean the whole bonfire?
Koko kokko!	Yes, the whole bonfire!

Chapter 1 Introduction

1.1 Project Motivation

The motivation for this project came from a large industrial concern operating in the power electronics industry. This company manufactures many products both large and small and has a major presence in both the domestic and international markets for its diverse product portfolio. Although this company has many products, which it supplies to many industries, the company was primarily interested in optimizing the manufacturing of a specific product. This product is an electronics product within the larger power electronics product family and is utilized for a very wide array of uses. For confidentiality reasons the company which is the subject of this study will be forthwith referred to as "The Company". In addition, to mask the identity of The Company the products, which are the focus here, will be simply referred to as "The Product".

The space which The Product operates in is currently rather dynamic by the staid metrics of traditional industrial concerns. "The market for "The Product" experienced an annual growth rate of just over 7% between 1995 and the year 2000[1]". Dramatic economic growth in developing countries such as China, India, and Southeast Asia is providing enormous potential for the product in these regions. Furthermore industrialized nations such as North America and Japan continue to purchase this product, to improve manufacturing quality and increase productivity.

Fierce competition and declining prices of power electronics are driving the average selling prices of products such as these down. Additionally, technology is moving so fast today that product life cycle is rapidly declining. Once lasting 10 to 15 years, product life cycles have been driven down to 3 to 5 years for micro (0-4 kW) and small (5-40 kW) products. Accelerating competitive pressures are driving suppliers to find the most cost-effective production methods.

Recent advances in materials and electronic component technology accompanied by price pressure, has made it necessary for most suppliers to out-source the fabrication of virtually all product components. Specifically, components such as printed circuit boards and power semi-conductors have become commodity products and have been outsourced as a competitive advantage. Outsourcing has reduced internal overhead costs and provides direct access to the latest technological developments in these areas. Today the core competencies of product suppliers lie in the realm of product design & development and assembly.

Perhaps the most time consuming and capital intensive portion of assembly of this product is the test system. Although the level and thoroughness of these systems varies by manufacturer, a comprehensive test system including functional and stress screening technologies can increase assembly times by up to 400% and can more than double the capital investment required in a new assembly line. Comprehensive test systems not only increase capital investment and assembly costs, but in many cases add an additional level of complexity to the process, ultimately increasing ramp-up times and reducing flexibility as volumes fluctuate.

The primary purpose of the test system is to screen components for early life failures and to verify proper assembly, ultimately resulting in increased outgoing quality levels and reduced failures during customer operation. However determination of key performance characteristics and the resultant test system are typically developed during the design phase of the product and are seldom revised after the design has been released to manufacturing. The Company believes there is value in optimizing test methods, without negatively impacting quality. This is the impetus behind this project and thesis.

1.2 Project Definition and Resulting Priorities

The original scope of the project was broadly defined as 'an effort to develop a new philosophy and new methods in production testing'. Early in the venture The Company's project definition was summarized into two concise statements:

> **Priority 1**: Benchmark industry best practices amongst competition and equivalent companies to determine optimal quality control procedures with minimal cost impact

> **Priority 2**: Determine optimal test methods to achieve desired quality goals and decrease overall test times by 50%

Four product lines were identified that accurately represented the manufacturer's business. This list of four was further focused down to one product line that had a significant amount of data available for analysis. This data was collected by an on-line production test system, stored in a database and was thought to be the most accurate data available. Additional information and description of the product line and test system will be provided in Chapter 3.

The development of two priorities and the identification of a key product line became the basis to approach development of new test philosophies and test methods over the course of the project.

1.3 Approach & Methodology

The two priorities, although based loosely on each other, were approached simultaneously. Each had an individual approach and required different resources, as outlined below.

1.3.1 Priority 1

Eight companies were identified for the benchmarking study. The companies included some competing directly in this industry and others in the computer, printer and copier industry. Additional industries were chosen, as they were expected to share like electronics components and have similar assembly, quality and test system issues. The benchmarking study was also used as a vehicle to leverage Massachusetts Institute of Technology (MIT) Leaders For Manufacturing (LFM) program relationships in an effort to gain access to similar manufacturing processes and to promote knowledge transfer.

Each company included in the benchmarking study was interviewed and questioned from which a correlation matrix was developed. The matrix became the basis for developing new test system methodologies and for making recommendations to the product design and development process.

1.3.2 Priority 2

An extensive analysis of data collected by The Company's on-line production test system was performed. The on-line test system records and stores test results from each of the four System Tests. The four System Tests consist of Hi Pot, System Test 1 (T1), Burn In and System Test 2 (T2). A complete description of each test can be found in Chapter 3. When possible the data was analyzed in product sub-groupings according to various product characteristics, including device size and voltage rating.

Using the production data, a failure analysis of the number of units that failed per units tested was completed. The Pareto analysis of test failures was then compiled and compared across products. Next, a first pass yield definition was agreed upon and

calculated for each System Test. Finally, field failure data was analyzed and compared with the findings from the on-line test system performance analysis.

The industry benchmarking findings and test system analysis results became the basis for recommendations to the product development process, redefining the production test methods, and the justification to reduce test times. This thesis explores the above analysis and findings in depth.

1.4 Project Goals, Measurements, and Deliverables

The project goals were again based on priority. From the benchmarking study, priority 1, a summary of industry best practices in quality control and test procedures was compiled. The summary compared test methods on a product rating, hp or kW, basis. This analysis and summary can be found in Chapter 2.

By establishing new test methods (priority 2) test times were to be reduced by 50%, without compromising current field quality levels. The baseline used for test time reduction was established in the first week of the on-site investigation. Recommended changes to the current product development process and modifications to the current test methods were presented upon the completion of the on-site investigation. Recommendations to improve the current test methods are discussed in Chapters 4 through 6.

Chapter 2 Competitive Benchmarking

2.1 Industry Overview & Market Share

These devices are typically categorized into five output capacity (kW rating) segments: micro (1-4 kW), small (5-40 kW), midrange (41-200 kW), large (201-600 kW) and mega (>600 kW). Growth in the industry will primarily come from the micro and mega segments. Companies such as The Company and Its other large competitors are considered the leaders in The Product's industry, accounting for the majority of worldwide revenues. In 1995, industry shipments of the product and associated services totaled $4.13B. Shown below are the respective market shares in the Micro (Figure 1) and Mega (Figure 2) segments:

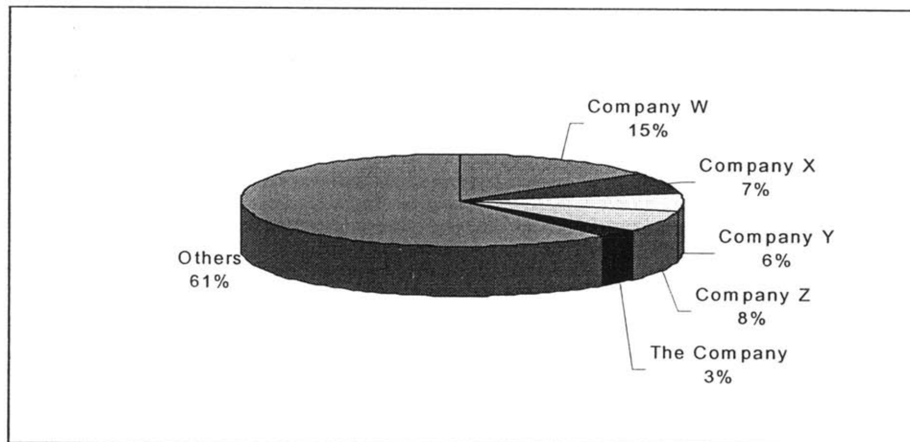

Figure 1 - Market Share of Leading Micro Product Players

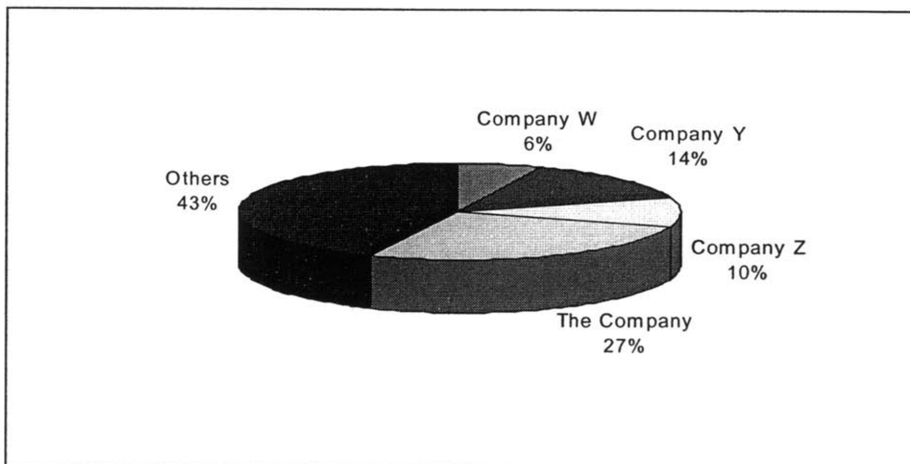

Figure 2 – Market Share of Leading Macro Segment Players

2.2 Overview of Subject Companies

The subject companies for this benchmarking study come from a wide variety of industries. The purpose was to gain insight into the testing systems used in the greater electronics manufacturing industry. Although the environmental conditions under which The Products operate may present both thermal and particulate-born challenges (i.e. dirt, dust, chips, and other contaminants) to the product design and operation, the underlying electronics should require similar testing to ensure product quality and operational excellence. Therefore, the scope of companies and industries examined were not limited only to The Product's industry or even to the power electronics industry, although, both are present in the subject population. The following presents an overview of the companies, their products, and the environmental conditions under which they may operate. This is done so as to provide a framework from which to base comparisons of testing systems to those used at The Company.

Company A is a producer of high-speed digital imaging and printing systems. Their products consist of very large units capable of handling global network printing responsibilities. The copiers are highly customized and produced on a make-to-order basis. The majority of their products operate at 220 VAC with both 50 and 60 Hz options. Under these power ratings internal thermal conditions are high, especially on a component by component basis, however, generally the units are operated almost exclusively in environmentally controlled rooms. The conditions here are generally less than 80° Fahrenheit and humidity is controlled.

Company B is a global player in the personal computer and workstation industry. Company B produces a wide range of both laptops and desktop units covering entry level units up to networked workstations. The power ratings on their products range from 110 VAC on the smaller units up to 220 VAC on their larger workstations. The division that was studied are responsible for building desktop units only, both for the PC and workstation markets. This division builds both standardized and customized products supplying both make-to-stock and make-to-order demand scenarios. Internal thermal conditions in these products is rather benign when compared to the industrial environment, however, component specific heating is a problem.

Company C manufactures high-end electrical devices mainly for the North American market. The company supplies these devices in the 0.4 to 750 kW range with power ratings of 230, 460, and 575 VAC. The larger sizes are very customized and built-to-

order, whereas the smaller devices have little to moderate customization and are mostly built-to-stock. These devices are used almost entirely in industrial environments where conditions are not controlled. The thermal conditions can be challenging as high temperatures are generally present. Additionally, dirt, dust, and various other contaminants common to industrial settings are present and may work into the devices.

Company D is a global producer of large and small printers and test and measurement equipment. The products have voltage ratings of 110 VAC and more. These broad ranges of products operate in a wide variety of environments. The office products generally operate in temperature and humidity controlled environments whereas the test and measurement equipment may be exposed to more deleterious environments such as those generally found in the industrial setting.

2.3 Testing at Subject Companies

2.3.1 Company A

Manufacturing at the division of Company A (note: only one division was investigated) consists mostly of final assembly. It is here that components and sub-assemblies from suppliers are assembled and tested prior to delivery to the customer. The cycle time for the product is approximately 33 hours; approximately 11 hours of which is attributed to assembly and approximately 22 hours of which is attributed to testing.

The test sequence consists of four steps. The first step is the Hi Pot test. This test is a safety check of the unit's circuitry. Here, as is typical within the electronics industry, a large potential is applied to the products main circuitry to ensure load-carrying capability. Next, the unit is subjected to its first power up. Here, generally, is where the first failures are encountered.

The automation and conveyance system responsible for controlling the printer's 'paper path' are both delicate and significant to the printer's proper operation. In addition, each component's harmonious operation with its neighbors is crucial to the overall performance. The next step, therefore, is to test this system by performing a continuous 'paper run'. The length of this test is approximately 25 minutes and can best understood as a functional test of the entire paper path. It is, however, the printer's first

period of extended operation and, therefore, the first test of component performance and reliability.

After assembling the outer cabinet a final quality check is performed on the printer prior to shipment. Some units may be exposed to an additional quality test at this point. This test consists of an extended paper run test and is performed on a sampling basis, only. This test is a quality check of both the manufacturing and design systems. It is here that small defects may be exacerbated into larger problems resulting in product degradation and, quite possibly, unit failure.

Although the extended paper run test might be perceived as a test performed specifically to highlight and expose 'early life' failures it does not necessarily perform that function. Notably, in fact, company A does not perform any tests to identify early life failures. This is considered to be the domain and responsibility of the component and sub-assembly suppliers. The company's perception of itself is as an assembly operation only and the extended paper run test is performed to ensure systemic quality is maintained and product wide functionality optimized.

Due to the fact that the unit was recently sold, field quality data was unavailable. It is just beginning to become available internally. Therefore, the data currently is only beginning to be evaluated for product performance and is not a good measure of overall test system effectiveness.

In addition to the quality processes used in printer assembly, company A also focuses on the quality of incoming components. During the product design phase components critical to overall product quality are identified and quality level specifications are established. The desired level for the supplier's process capability is also established. This number is typically set at a Capability Index of $C_{pk}=1.5$. The Capability Index is a statistical representation for process control widely used in industry. C_{pk} is defined by the following equation:

$$C_{pk} = \min\left[\frac{USL-\overline{X}}{3\sigma}, \frac{\overline{X}-LSL}{3\sigma}\right]$$

Equation 1 – Statistical Capability Index[2]

where

USL = Upper Specification Limit

LSL = Lower Specification Limit

σ = Process Standard Deviation

$\overline{X}$ = Process Mean

Company A then meets with suppliers, prior to establishing contracts, to verify the suppliers' manufacturing process can meet the required specifications. The Supplier Quality Assurance Group is assigned this job. The Group is a team of technical professionals involved in the product design phase as well as the manufacturing process. In notable situations where the company is concerned, the team has even gone into the supplier's factory and worked with the supplier to reduce process variability and meet required quality levels. Hands-on grooming, such as this, can prove immensely rewarding in understanding the supplier's underlying process capabilities. The company is, subsequently at a distinct advantage, such that the impact of proposed future changes and suggestions on product functionality and quality are better understood.

2.3.2 Company B

The testing at company B consists of three very different approaches which when combined is designed to minimize production testing. These three testing procedures are known as HAST, HASS, and run-in. HAST refers to Highly Accelerated Stress Testing. HASS refers to Highly Accelerated Stress Screening. Run-in simply refers to an extended period of operation.

HAST testing at company B is performed during the product design phase. In fact all HAST testing should be performed prior to product release to manufacturing. The idea is to test the 'goodness' of the design as one company representative noted. When the initial product design reaches a certain maturity in which component, sub-assembly, and/or full-assembly designs have been agreed upon the unit is tested to destruction under a wide array of environmental conditions. The designs are generally subjected to vibration and rapidly changing thermal environments to incur increasingly stressful conditions. When the unit fails under these conditions a quality examination loop is used to determine the failure mechanism. This information is then fed back into the design loop and the product is redesigned to correct the offending condition. After several

iterations the most likely failures are identified and corrected, thereby increasing the overall robustness of design and reducing the incidence of failure under normal or even extreme environments. Company B feels that they are testing not only component quality and design, but also assembly errors and manufacturing induced failures.

The next level of product testing performed at company B is HASS testing. HASS testing, due to the expense, is not performed on all products. Products, which receive HASS testing, are chosen based upon their cost and complexity. Only the more complex and/or costly products are chosen, as the cost for this level of testing may be high.

HASS testing backs off from the extremes of HAST testing. The objective of the one hour long test is not to take the product to destruction, but to test it to the outer operating design limits or specifications. The purpose is to highlight areas in the in-line production process and design which may be marginal and, therefore, not be able to stand up to the rigors of the outer operating regimes. As in HAST testing, the products are subjected to vibration and rapidly changing thermal environments to induce the greatest environmental stresses. Again, however, these stresses should only be at the design limits. As failures occur, a profile of failure mechanisms can be obtained. In addition, failure rates begin to be understood.

Unlike HAST testing, however, HASS testing is performed mostly during the manufacturing ramp up stages. At the beginning of product manufacture the sampling rate is 100%. Although timely, the low volumes allow for maximum testing. Quickly the sampling rate is dropped as failure profiles are developed and rates climb. Company B ultimately drops HASS testing as full volumes are generated and data acquisition is complete.

The final level of testing utilized by company B is the run-in test. This test verifies the entire manufacturing process is in control. The length of the run-in test is chosen at 96 hours. The units are allowed to operate in a normal environment and failures are monitored. The company uses the information from this test to develop a Mean Time To Failure (MTTF) profile. As the MTTF profile is better understood, the optimal point for run-in is identified and the test length can be modified to account for the changes. A sampling rate is used for the run-in test, however, this rate was not identified.

Company B has also established quality programs with their suppliers. Critical components, specifically PCB's, have additional testing requirements placed upon them. These tests are performed at the supplier. Specifically, boards are subjected to stress screening at elevated temperatures and with high and low voltages applied. The boards, along with other delicate components, are also exposed to a series of shipping tests, including drop tests to verify packaging requirements. Contracts with suppliers include specifications as to the type and degree of testing, as well as minimum acceptable quality levels. Company B also has a group of component quality engineers, who sample poor components and provide feedback to the supplier regarding quality levels. In some cases, they might also utilize their unique viewpoint as a customer to work with the supplier to recommend improvements to the manufacturing processes.

2.3.3 Company C

To ensure testing success company C maintains a rigorous quality program. This program begins with an incoming inspection on new and poor quality level components. Such inspections are carried out to verify supplier compliance with existing design and quality goals. All components are subjected to inspection. In addition, the manufacturing process is audited to ensure that it is capable. A random sampling of the process steps is conducted to ensure adequate coverage and in-process quality.

Interestingly, to maintain quality of the most delicate components, company C actually builds and assembles the PCB's in-house. The company has both through-hole and surface mount technology manufacturing capability. Leveraging a company wide competency in the manufacture, assembly, and test of PCB's allows the company to ensure a steady flow of quality boards. The company maintains a 'Center of Excellence' approach in the design and build of its boards. After board design the company will maintain board production in house through prototype and pilot stages. Only after the process has been determined to be in control and capable is the board manufacturing subbed out. Ultimately, they sub-contract up to 50% of the long term production, but always keep the remainder in-house to ensure adequate part flow. An additional benefit of maintaining up to half of the board production is that it gives the company a sense of the sub-contractors capabilities. By doing this the company can always discern between true production problems at the sub-contractor from others. Company C, therefore, effectively has a tight leash on its most delicate component's supply base.

As the PCB's are seen to be the most delicate part in company C's product with respect to reliability and manufacturing process the control of this step is seen as crucial to their product's manufacturing success. Using several SPC techniques as outlined within ISO 9000 guidelines the PCB manufacturing process is controlled. All product and process discrepancies are documented. Each discrepancy is then corrected and the corrective action is traced back to its source to ensure no further occurrences. This final step in preventative process 'maintenance' is key to long term process maturity.

The company has several quality improvement programs that it institutes as its primary continuous improvement tool. Utilizing field and manufacturing quality data the company has developed a fact based prioritization program to focus on the most dominant issues affecting its customer base and its manufacturing community. Issues in the manufacturing phase can be anything affecting the design, component, or assembly of the product. Focusing only on major product lines, the failure mechanisms are ranked and the top five are identified for investigation and solution recommendations. The top five problems can be identified via feedback from the field and/or issues affecting plant yields.

Using product return information, manufacturing data, service call information, and repair center data, the process is to sort, categorize, and pareto the data. This will allow the identification of the top five product issues. Next, ownership is assigned. Ultimately, corrective action is determined and implemented. The process output can be changes to product design, manufacturing process, testing procedures, or simply documentation.

The company employs a team-based solution to this process. The team is led by a quality group, but is cross-functional in its nature. The rest of team is picked from engineering, engineering quality, component engineering, quality/test engineering, manufacturing engineering, repair and technical support, purchasing, customer quality support, and marketing. Such diverse team membership ensures a wide scope to the solution and to the corrective action.

Quality programs at company C do not stop at the time of product delivery. The company has set up a Technical Rapid Response Team. The purpose of the team is to provide 24-hour response for all direct out-of-the-box failures. In such an occurrence a new product is immediately routed to the customer. The failed unit is, meanwhile, direct shipped to the company's engineering department for failure analysis. This process not only provides immediate customer response and, hopefully, satisfaction in the face of a

rather difficult scenario, but also provides the company with direct access to failure data. In light of the typical problems surrounding holes in field data this advantage can be notable.

Product testing at company C corresponds strongly with the life cycle requirements of the product. During the product design phase the company employs Highly Accelerated Life Testing (HALT) very similarly to the HAST testing employed by company B. During the pilot and ramp-up stages a typical Burn In testing process is used. When the product is at full manufacturing run, an audit Burn In testing strategy is used. Finally, to support the long term product life, reliability testing is conducted.

During the design phase HALT testing is used to validate and demonstrate the design limits and tolerances of the product. Here the upper and lower destructive limits of the product are explored and documented. Through a series of repetitive applications of thermal, electrical, and/or vibrational stresses the lowest common denominator of failure mechanisms can be quickly identified. Additionally, their solutions can be verified via the same tests ensuring a long-term fix. It is important here to note that the stresses that are applied need to be applied not only with a large enough magnitude to induce failure, but also at a high rate of change. For instance, a thermal stress may be applied, however, a constant thermal stress is of less value than a changing thermal gradient on the order of 10-20 °C per minute. The applied stresses in these cases are far more severe. Correctly applied, HALT testing during the design phase should capture design failures and component defects.

Burn in testing is used during the pilot stage of the product life cycle. Using the operational limits as boundaries for load application, the products are subjected to thermal loads and power cycles. This testing should discover process, supplier, and design defects. Additionally, the Burn In process should be a validation of the proper use of Design For Manufacturing techniques, the manufacturing process, and supplier quality. During the pilot stage a 100% Burn In testing rate is conducted on an identified pilot lot. The testing is conducted at a fixed elevated temperature. This temperature is chosen to be the product rating temperature.

The length of the Burn In time is quite long at the beginning of the pilot stage and is shortened throughout the stage. Initially, a Burn In time of 96 hours is chosen. The units, therefore, are run using inductive loading only (no complementary device loading) for a full 96 straight hours in a chamber with an ambient temperature equivalent to the

product's design operational limit. Eventually as product design and operational confidence is achieved the Burn In time is reduced... first to 48 hours, then 24, and finally 12 hours. Each pilot lot receives 100% Burn In coverage during these stages.

After the pilot and ramp up stages Company C employs an audit Burn In or final product Run In test. The purpose of this type of testing is to assure process control and supplier quality; effectively to maintain overall quality assurance while minimizing field quality problems. Run in testing differs from Burn In testing in the length of time tested and the applied ambient thermal loads. In the case of the Run In tests 100% of the units produced experience full load dynamic testing. Depending upon the product size test lengths vary from 12 minutes to two hours and are based upon ensuring that, at minimum, three product time constants are achieved. Additionally, the current load is cycled approximately every two minutes from 50% to 100% and then to 150% to induce the dynamic loads. The loads are simulated using matching motor dynamometers. Finally, in lieu of operating the Run In tests at an elevated ambient temperature the heat generated from the electrical loads is deemed to be sufficient to simulate thermal loading.

The long-term reliability of the product is also tested. Here units are randomly sampled from the line. Long-term operational tests are conducted at the rated temperatures, loads, and power cycles to help understand duration capabilities. Such testing tests the long-term reliability of the design, process, and components.

Lastly, Company C, in addition to its product life based testing, performs extensive PCB testing. In circuit testing and parametric functional testing are conducted in an attempt to achieve high rates of out-of-the-board shop quality levels. Again, PCB's are considered to be a company key technology, or core competence, and this testing attempts to maintain high levels of quality in what they consider to be the most delicate component. Additionally, the vendors that Company C ultimately chooses perform board level testing in the form of "bed of nails" conductance tests[3]. Coverage, however, was estimated to be only 70% at best. Interestingly, Company C used to perform Burn In testing on the PCB's but no longer does so. Unfortunately, the reason for the cessation of testing was not determined.

2.3.4 Company D

Due to the diversity of its product testing at Company D is, by necessity, similarly quite diverse. Due to this diversity, specific testing procedures were not explored. Instead, investigation focused more on general testing trends at the company. Specifically, the use of HAST testing, in lieu of the traditional Burn In testing, was explored.

At Company D the use of Burn In testing was slowly phased out over many years in favor of HAST. Although the company had a long history of performing Burn In testing to identify infant mortality failures there is almost no Burn In conducted anymore. Company representatives did, however, note that pockets of resistance to new procedures did remain as members in these areas were not yet convinced of reported testing benefits provided by HAST.

To understand the benefits of Burn In Company D representatives required data. To convince themselves, tests were performed in which even serial numbers on a given product line were shipped without Burn In and odd serial numbers were shipped after the normal Burn In procedure. After waiting upwards of 1-2 years for the necessary field quality data to be given adequate time to be reported, the experimenters determined that there was no statistical evidence to buttress the additional testing. Specifically, they could not show with statistical clarity that field quality improved (i.e. the rate of early failures was reduced) by performing a traditional Burn In test. To ensure they had enough data to capture the necessary failure instances the even/odd experiment was performed for an entire year's production run.

The results from this study have apparently changed the way production testing is performed at the company. Incorporating HAST into the normal design phase is much more commonplace. To enhance these procedures the process includes the use of thermal stressing with a very high rate of thermal change ($\Delta T > 10\,^{\circ}C/minute$). In fact, the company believes subjecting the product to a constant temperature is a very poor application of stress and is not productive in ferreting out failures.

In addition, the company is exploring the use of vibrational stressing to investigate how applied vibration can affect the product life and failure rates. The company claims to have shown that subjecting a product to five minutes on a vibrating table induced the same failure rates as a week of stress testing. Unfortunately, these claims are unsubstantiated otherwise. In vibrational testing it was pointed out that pseudo random

vibrations should be applied to ensure no directional bias is applied. Such a bias might induce rapid failure rates in favored directions and ignore failure rates in non-favored directions. To support such pseudo random vibrational testing the company uses a 6-axis vibration table. Internally, the company refers to the application of thermal cycling and random vibrational testing as STRIFE testing (STress and LIFE testing).

To implement the HAST testing the company uses a sampling strategy to ensure consistent testing application during the production run. The reported sampling rate is approximately 1%. It was not discovered how this rate was determined. To ensure useful results from the HAST testing the company tracks down the failure's root causes and corrects them. This is the most important role of the testing procedure.

2.4 Summary of Benchmarking Study

This benchmarking study provided insight into some of the testing procedures currently being employed within the electronics manufacturing industry. Apparently procedures such as HAST testing are being utilized in an effort to provide more robust designs to the manufacturing floor so that in-line testing such as traditional Burn In may be reduced. This philosophy, if employed correctly appears to provide a more robust product design greatly reducing cycle time and subsequently product cost.

Other results from the study showed that there is a wide tendency within the industry to work much closer with the supplier base. Several of the companies investigated were conducting extensive supplier relations. To the extent that the suppliers were comfortable allowing these companies onto their manufacturing lines and even their product design groups, these companies were quite successful in building mutually beneficial relationships. The relationships served both sides by improving incoming component quality to the company's manufacturing lines and ensuring future business for the suppliers who showed a willingness to grow and adapt to process improvements. Such improvements with component quality are sure to improve yields for in-line testing possibly allowing the companies to reduce their dependence upon such testing

Additionally, by evaluating the capabilities of their supplier's processes the companies were able to successfully manage future product and process changes. Specifically, they were able to reduce incoming product variability reducing the need for more extensive in-line testing. Effectively, such testing was successfully removed from the company's manufacturing and pushed upstream, where appropriate, to the supplier.

Chapter 3 Product Line

3.1 Product Overview

The product line consists of four frame sizes. The term frame size refers to the physical size or footprint of the unit. A further sub-set of product differentiation within each frame size is the type series. The rated output power (kVA), voltage rating (Vac) and enclosure class designate a type series. For the purpose of this study, only units manufactured at the internship site were analyzed, yielding a final number of type series of 10.

3.2 Product Flow

The product flow is shown in Figure 3 below:

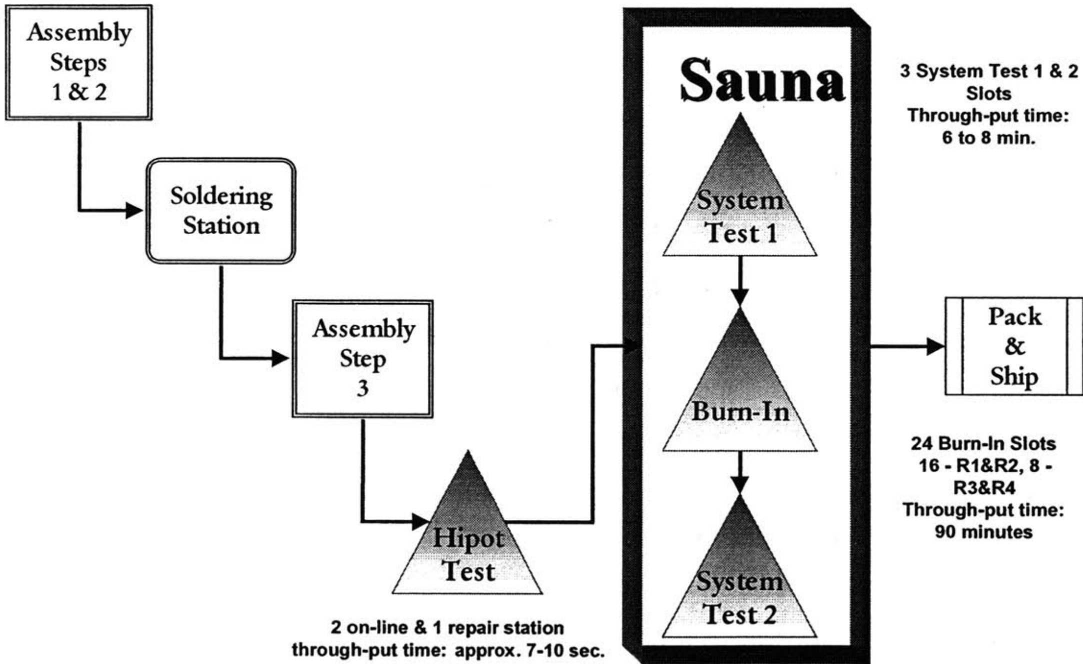

Figure 3 - Manufacturing Product Flow

All products are assembled on this production line, in various sizes and voltage ranges. Batches can run from as little as 1 unit to as many as 15. Assembly step 1 is a basic assembly step where the heat sink and various components such as fans, etc. are attached to the frame or platen. The platen is the board that supports the unit

throughout the production process, allowing it to be transferred by powered roller conveyors. The platen has a bar code that contains all vital product information and is used in all product-tracking software throughout the assembly process. Assembly stage 1 feeds assembly stage 2 via powered conveyors.

Assembly stage 2 consists of two identical and parallel assembly cells, with one operator in each cell. As in all cell designs, each operator has all the tools and components in the cell necessary to perform the assembly process. The operator assembles additional components onto the unit, performs a visual quality check, and releases the unit to another powered conveyor. Immediately after assembly stage 2 the powered conveyors merge, feeding a single soldering robot.

As the unit enters the soldering station, the robot visually aligns itself to the size of the platen/unit, reads the bar code and proceeds with the soldering process. Upon completion of the soldering sequence, the unit is transferred to a power conveyor which, using a queuing algorithm built into the system software, transfers the unit to one of two assembly stage 3 cells.

At assembly stage 3, the operator performs a visual inspection of the solder quality, attaches additional components to the unit and releases the device to the first of four consecutive test stages. Each test stage is highlighted in the diagram above for clarity.

There is one Hi Pot test station on each line and also one at the repair station. Upon passing the Hi Pot test, the device enters the 'sauna' for the remainder of the test sequence. The sauna is a thermally controlled environment, operating at ~45° Celsius. While in the sauna, the device is subjected to three additional tests: System Test 1, Burn In, and System Test 2. The device navigates within various testing slots with the aid of a manipulator arm. The diagram above indicates the number of test slots designated for each test and frame size. If the device fails any of the four tests, it is immediately rejected and sent to the repair station (located outside the sauna). Additional explanation of each test sequence is found in section 3.3.

After completing the test regime successfully, the device is transferred from the sauna to the pack station. The device is removed from the platen, the cover is placed on the unit and it is boxed. Product manuals and information sheets are placed in the box prior to the taping operation. The boxes are then transferred to the palletizing line for

stretch wrapping and shipping. Platens are subsequently returned to assembly stage 1 via a high-tech, over-head powered conveyor.

3.3 On-line Production Test System and Database

The product-line utilizes an advanced test system to ensure product quality. This test system was developed in partnership with a local software provider. Tests are performed after the final assembly stage of the unit to verify the unit is assembled and working correctly. The test system consists of four product tests: Hi Pot, System Tests 1 & 2 and Burn In. Each test has several potential outcomes. The two most common outcomes are pass and fail. If the device under test (DUT) receives a test outcome of 'pass' the unit proceeds to the next test or the final stage of production. If the DUT fails a System Test, the unit is automatically sent to the repair station where a trained operator inspects the unit. After the unit is repaired it returns to the production flow at the first System Test, Hi Pot. Each unit must receive 4 consecutive passes, one from each test, prior to shipping.

All test results are stored in an on-line production database in various data tables. A brief explanation of each System Test and potential outcomes is discussed below.

3.3.1 Hi Pot Test Description

The purpose of the Hi Pot or insulation test is to ensure that the construction of the converter has not been damaged in any way and that the device has the capability to withstand the demanded voltage. There are three Hi Pot test stations, one after each assembly stage 3 cell and one at the repair station. The duration of the Hi Pot test is 7 to 10 seconds, depending on the frame size being tested. The two possible results of the HIPOT test are pass and fail.

The test consists of 3 primary phases: insulation test 1, voltage withstand, and insulation test 2. A brief description of each phase follows:

1. **Insulation Test 1:** Measures the insulation resistance between the main circuit and the chassis. Measures insulation resistance. This test is considered a safety check

2. **Voltage Withstand:** Connects power between main circuit and chassis. Measures leakage current and verifies insulation is OK.

3. **Insulation Test 2:** Repeats insulation test 1. Verifies the voltage withstand test has not compromised the integrity of the unit.

If the unit passes the test, it is sent into the 'sauna' to continue the test sequence. If the device fails, the test system records the failed phase and the measured values of the test. The unit is subsequently released to the repair station. When the repairperson is confident the unit is operating correctly, she re-tests the unit at the Hi Pot station located at the repair station. If the unit passes, it is transferred to the sauna and proceeds to the next test stage. If the device fails, the repairperson continues to troubleshoot the problem until the device passes the Hi Pot test. Upon each trial, the repairperson records her efforts in the on-line database, including the action taken or component replaced.

3.3.2 System Test 1 & 2 Test Description

The purpose of System Test 1 & 2 is to ensure the components of the converter are not damaged or defective and that the converter operates as designed. The two System Tests are virtually the same; however, in System Test 2 the device has been exposed to Burn In and is operating at an elevated temperature (see the next section for a detailed description of Burn In). Each of the System Tests has four possible outcomes: pass, fail, marginal, and abort. The pass and fail qualifiers were discussed in the previous section. A marginal outcome indicates the unit has passed the test but the measured values are within 10% of test limits. Upon receiving a marginal qualifier, the unit proceeds to the next step in the production process without delay. Operators are unaware of a marginal test outcome. The abort qualifier indicates an operator aborted the test due to operator error or test system error.

The System Tests are anywhere from 6 to 8 minutes in length. The test time varies based on System Test 1 vs. System Test 2 and the frame size. The tests consist of the following 15 phases. Each description is numbered first with the phase number corresponding to System Test 1, followed by the phase number of System Test 2. An X indicates phases absent from the test pattern and exceptions are noted in the test description:

> **1,X. Verification test of the intermediate circuit:** Consists of a low voltage and high voltage level. Verifies there is no short circuit in the main circuit, DC capacitors are operating correctly, and that discharge resistors are good. This test is specific to System Test 1.
>
> **2,1. Connection of Supply Voltage:** Exposes the unit to normal operating condition and voltage. Measures the supply voltage, verifies capacitors are balanced and that electronics have supply voltage.

3,2. Voltage measurement and serial communication verification test: Measures that the unit sees the same voltage as measured in the test above and that the software is running.

4,X. Size Check Test: Reads the unit specification from the palette and verifies the correct printed circuit board has been used in the assembly process. This test is specific to System Test 1.

5,X. Software Version: Verifies the correct software or chip has been loaded on the printed circuit board. This test is specific to System Test 1.

7,4. Temperature Measurement Verification Test: Makes a temperature measurement below the IGBT, compares this value to the high and low limit. Checks for overheating.

6,3. Loading of Parameters: In System Test 1, loads test parameters for remainder of System Test 1 and 2. In System Test 2, load type specific parameters and set fault memories to zero.

8,5. Current Measurement Verification Test: Verifies the current the unit is experiencing to the known current placed on the unit.

9,6. Verification Test of the Ain1, Ain2, Digital Inputs, Aout, and Relay Outputs: Tests the functionality of the control board; checks board inputs and outputs.

10,7. Test of Basic Functions: Tests the unit under light load and nominal load conditions. Verifies complementary device runs in the right direction, charge relay is energized, line currents are symmetric, and the fan is rotating.

13,10. Earth Fault Trip Test: Verifies the complementary device is running and the unit self-protection (earth fault trip) function works correctly. This test is specific to System Test 2.

11,8. Overvoltage Trip: Verifies the complementary device is running and the unit self-protection (overvoltage trip) function works correctly. This test is specific to System Test 2.

12,9. Overcurrent Trip: Verifies (overcurrent trip) the unit self-protection function works.

14,X. Loading of Burn In parameters: Sets test parameters for the next sequence, Burn In. This test is specific to System Test 1.

X,11. Loading of factory default parameters: Sets factory default parameters in preparation for use at the customer site. This test is specific to System Test 2.

X,12. Set Other Parameters: Sets test date, sets previous faults to zero, and checks the unit ID code. This test is specific to System Test 2.

15,14. Discharge the intermediate circuit: Disconnect mains, power supply shuts down.

Upon completing System Test 1 successfully, the device is transferred to a Burn In slot, if the device fails it is transferred to the repair station. The repairperson follows a similar trouble shooting process as outlined above but also checks the online database

to learn what phase of the test the device failed. This can often lead the repairperson to a specific root cause, such as a defective board or loose connection. When the repairperson determines that the problem is resolved, the unit is once again subjected to the Hi Pot test. If the unit passes, it returns to the 'sauna' to resume the test sequence.

The above process is also true for System Test 2, however, if the unit receives a pass after the test, the unit is transferred to the packing and shipping station. If the device fails, it is similarly inspected and re-tested. Each device must receive four consecutive passes before reaching the final production stage, packing and shipping.

Similar to the Hi Pot test, when a unit fails either System Test 1 or 2, the test system records the failure including which phase of the test failed and the measured values for the phase.

3.3.3 Burn In Test Description

The purpose of the Burn In test, as described by The Company, is to identify those components that may fail during the early periods of unit life and thus induce so-called early-life failures. Specifically, this is attempted by running each unit at an elevated temperature so as to induce thermal degrading loads on the unit's circuitry and components. Since, the DUT remains in the sauna after System Test 1 this is relatively easy. Placed between the two System Tests, the Burn In test is conducted within the elevated temperature environment (~45°C) of the sauna.

While the test steps performed during the two System Tests are aimed at verifying the unit's functionality and stressing individual components, the purpose of the Burn In test is to operate the device for a set period of time and verify system robustness. While in the sauna the unit will operate with nominal current and through a pre-ordained loading sequence to simulate an actual working environment (see Figure 4). As such, the Burn In test period is significantly different from the other three tests conducted. The Burn In test consists of an automated series of changing complementary device speeds that the unit must command.

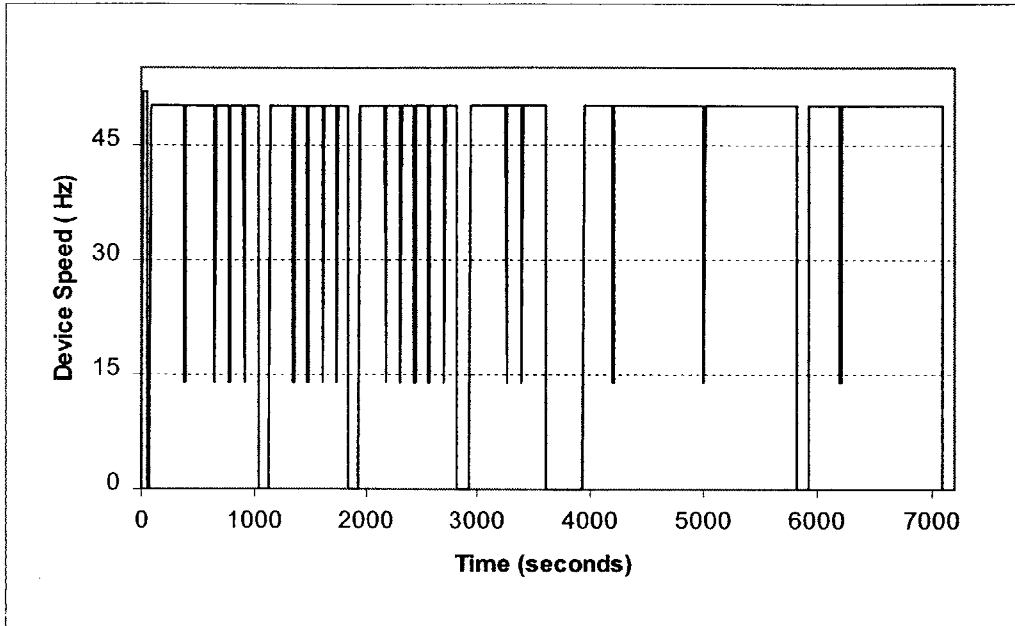

Figure 4 – Burn In Test Cycle

If, as in the two System Tests, the unit should fail it will be automatically removed from the sauna and transferred to the repair area. The test database automatically records the time to failure in seconds beginning with zero at test inception. After arriving at the repair station the repairperson will analyze the unit and attempt to diagnose the problem. If the diagnosis is successful the repair will be implemented, a repair code will be assigned based upon the nature of the failure, and the unit will be returned to the product flow. At the same time, the repairperson will manually enter the repair instructions and any other pertinent data into the database. Whether the unit returns to the product flow upstream of Hi Pot or System Test 1 depends upon the nature of the repair and whether it is safe to do so.

In the case where the diagnosis is unsuccessful in identifying the failure mode an "unknown condition" type code will be assigned to the event and the unit will begin the Burn In test again. Ultimately, the unit will pass all tests and then it will be directed to System Test 2 that was described above. To do this, the DUT will be automatically removed from the Burn In slot that it occupied during the Burn In test and transferred to the System Test slot. As identified in Figure 3 there are three slots dedicated to performing the System Tests while there are 24 slots (16 slots for the smaller frame sizes and 8 slots for bigger frame sizes) dedicated for the Burn In operation. The large difference is due to the large disparity in time required for the System Tests versus that

dedicated to the Burn In operation. As the figure shows, each System Test requires approximately six to eight minutes to complete. Similarly, the Burn In operation lasted 120 minutes at project inception.

The pass criterion for the Burn In test is simple, either it fails to operate at the prescribed speeds or not. In the case of failure, "FAILED" is recorded in the database for that sequence and in the acceptable case "PASS" is recorded.

Chapter 4 Hi Pot Performance Analysis

As indicated in previous chapters, each unit is subjected to the four-phase test regime prior to shipment. The 100% test philosophy has established world class levels of reliability and performance in customer applications.

The results from each test phase are stored real-time in the production database and organized in information tables. Using the data from multiple tables, a Microsoft Access query was developed to extract and organize the data in a user-friendly format. The resulting information was transferred to a Microsoft Excel spreadsheet for further manipulation.

As Hi Pot was the first test in the system that was analyzed, it served as the vehicle to become familiar with the databases and to understand what data was available and how it was organized. As the Hi Pot analysis progressed, the findings were used to establish the process with which to analyze the remaining test system. The discussion below describes the trial and error process that was used on the Hi Pot analysis, the identified test time reduction and the method that was established to analyze the remaining test system.

4.1. Failure Analysis

The first study performed on the Hi Pot test was a failure analysis. A failure analysis was considered to be a good indicator of failure trends in the test system. From this data, it was assumed that a relationship would be identified between various levels of production and test performance. The analysis was performed at the daily production level, calculating the number of units that failed the test and the total number of individual units tested. As explained in previous sections, units that fail a test are diagnosed, repaired and returned to the production line. Any unit that fails a test phase must return to the beginning of the test system and be re-tested at each test station. The process repeats until that unit receives concurrent passes at each test phase. The test system tracks the performance of each unit by assigning a sequence number each time a unit is tested at a particular station. Therefore, units may be tested multiple times at each test station and thus the importance of measuring the number of individual units tested as opposed to the total number of tests per day. A query was designed to extract this data. A second query was used to establish the total number of failed tests per day.

Utilizing the output from the two queries, the daily failure rate was calculated. The failure rate was defined as the number of failed units per number of units tested. The number of units tested and the failure rate were plotted versus day of production. The intent was to identify a trend between the total number of units manufactured and the number of failures. It was expected that a higher failure rate would be seen on days when production was low and lower failure rates seen on days when production was high. An example of this chart is shown in Figure 5 below:

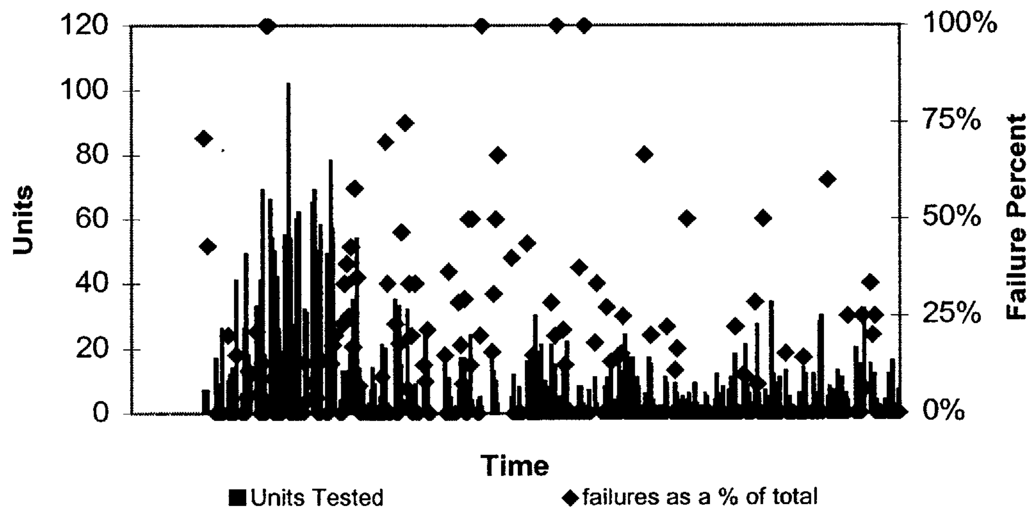

Figure 5 - Failure Analysis

No conclusive relationships were identified from the information. As the chart did not provide the insight intended, it was decided that a scatter plot might be a better format to investigate correlation between units produced and failure rate[4].

The scatter plot in Figure 6 was created using the same data as depicted in Figure 5. It is apparent that failure rates diminish as the number of units produced increases. However, the data populations also indicate an equal likelihood of low failure rates on high or low production days. After repeating this analysis on all available data at the type series and frame size level, no conclusive relationships were established. The original hypothesis to establish a correlation between production levels and test performance was questioned. In retrospect, the value derived in identifying a correlation would have been for the manufacturing planning process and not for improvements to

the test system. It was decided there was little value in pursuing this type of analysis for the remaining test system and the failure rate analysis was abandoned.

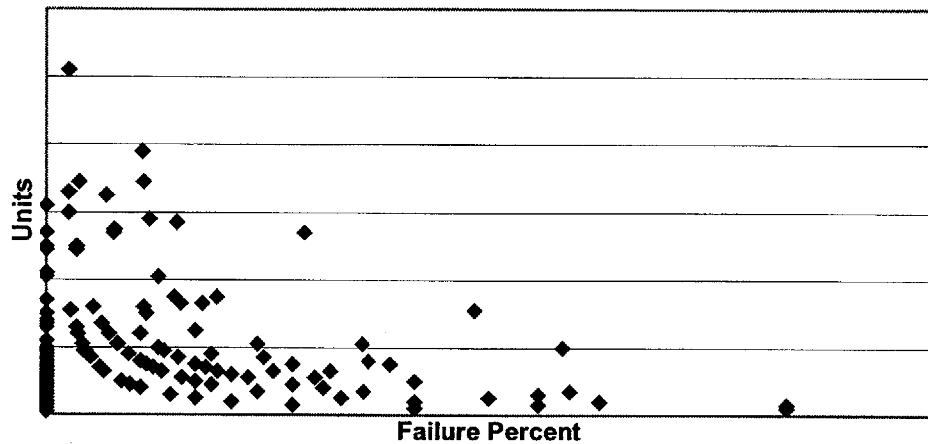

Figure 6 - Failure Rates vs. Units Produced

It was learned at this time that the production line under analysis was installed one year prior to this investigation, thus data in the earlier months of production could be skewed due to learning curve effects. After considering all product variations and corresponding time in production, it was determined that future analyses should include data from the present year only. The remaining discussion regarding data analysis will reflect this protocol.

4.2. First Pass Yield Analysis

Industry literature indicates when a manufacturing process achieves higher levels of process control and reliability, quality inspection can be reduced, resulting in significant cost savings to the company. First pass yield has been used as technique to justify reduced testing schemes in many industries. In this case, if first pass yield levels demonstrated consistent performance and approached six-sigma levels; perhaps the company could challenge the 100% test philosophy and justify a new methodology based on a significantly reduced sampling level. This redefinition would reduce the load on the existing testing station and simplify the test process. In the future, it would also reduce the level of capital expenditure required for installation of new production lines. It was agreed that this was a potentially beneficial analysis.

As mentioned in the previous section, the test system assigns a sequence number to a unit each time it is tested at a test station. Of specific interest in this case were test results with a sequence number '1' representing the first time each unit was tested at the Hi Pot test station. A query was designed to extract the total number of sequence number 1's per day and the outcome (pass or fail) of the test. A daily first pass yield was calculated and defined as the number of units that received a pass on sequence number 1 divided by the total number of sequence 1's. See Equation 2 below for the formula:

$$\textit{First Pass Yield (FPY)} = \frac{\textit{\# of Units with a pass on Sequence \# 1}}{\textit{Total number of Sequence \# 1's}}$$

Equation 2 - First Pass Yield

Figure 7 shows the FPY performance over time for all product variations:

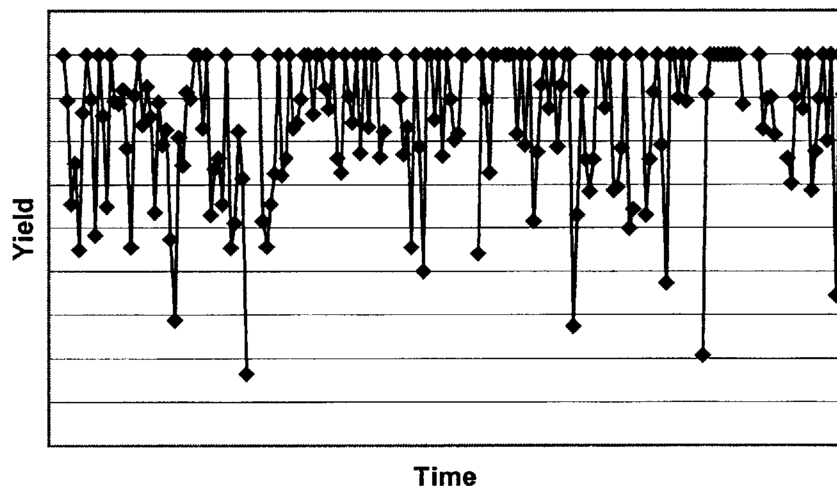

Figure 7 - Hi Pot FPY

This analysis was also conducted at the type series and frame series level. None of the product variations exhibited consistent first pass yield levels or approached six-sigma performance. Using these characteristics as the basis for test reduction, no recommendation could be made to modify the current 100% test philosophy.

Upon discussion of the results with the technical experts, it was discovered that there were two additional considerations that would prohibit future reductions in testing levels. The first consideration was that the Hi Pot test is a regulatory requirement for all products in this industry. Thus, any modifications to the Hi Pot test must continue to

meet the minimum requirements of the regulation. The second consideration was potential risk to downstream operations. It was widely believed that the risk of passing a defective unit to System Test 1, where a unit could potentially explode in the 'sauna', more than justified the test and resultant capital investment. Given these two issues, regulatory requirements and risk level, the technical team decided it was in the company's best interest to maintain the 100% test philosophy for the Hi Pot test, even as test performance consistently achieved six sigma levels of performance.

4.3 Pareto Analysis

The final analysis performed on the Hi Pot test was an analysis of failed units. There were two purposes to this analysis. The first was to characterize the modes of failure and correlate the data with field returns. If a relationship were identified between test system failures and field failures, it might be possible to use the on-line test system as a predictor of field performance. On the other hand, if there were no correlation, the test system may need to be re-designed to increase the 'robustness' of the test in an effort to reduce failures in the field. Chapter 7 has been devoted to this analysis.

The second purpose of the analysis was to Pareto the failures by test phase. As mentioned in Chapter 3, there were 3 primary phases in the Hi Pot test pattern. When a unit failed the test sequence, the test system recorded the phase of the test pattern that the unit failed. The intent here was to Pareto the modes of failure by test phase. This might better identify the relationship between failure mode and the phase in which the failure occurred.

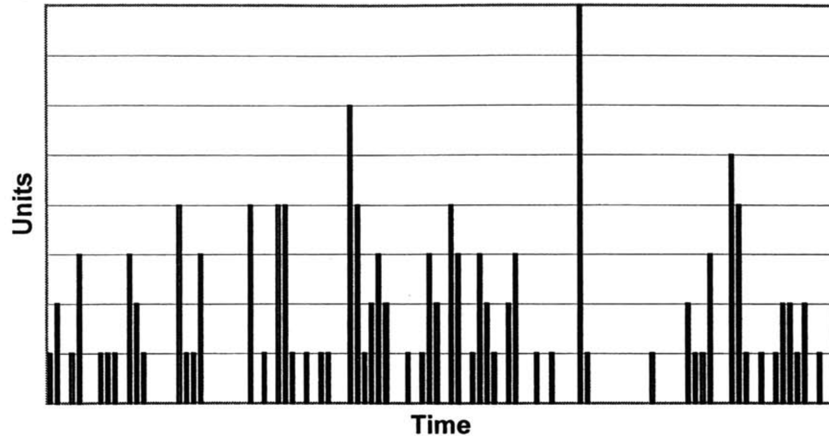

Figure 8 - Hi Pot Failure Pareto

A query was created to extract this data from the database. In this case all data was pertinent, including units that failed Hi Pot multiple times. Thus sequence number was ignored. Secondly, the query provided the 'test name' data, which made it possible to identify which phase the unit failed. Upon immediate inspection of the data, there were only two failures: A and B. The data was plotted and is shown in Figure 8. Here it is seen that for every "A" failure, there are two "B" failures.

Upon further exploration, the initial finding of the 2:1 relationship of failures between failure mode "A" and "B" was an outcome that was considered normal and not particularly insightful. The next step of the analysis was to sort failures by test phase.

Failure Mode: A(diamond) and B(diagonal)

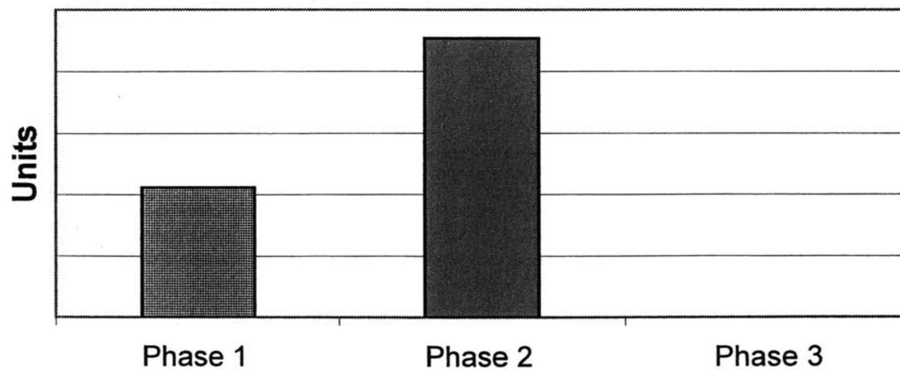

Figure 9 - Hi Pot Failure Pareto by Test Phase

In Figure 9 above, two pieces of information are revealed. First, it is apparent that the two failure modes correspond directly to a specific test phase, meaning that failure

mode A occurs only in phase 1 and failure mode B occurs only in phase 2. Secondly, failures occur in phases 1 and 2, but not phase 3. As previously described in Chapter 3, phase 3 verifies that phase 2 has not compromised the integrity of the unit and is identical test to phase 1. As there are no failures in phase 3, it is evident that phase 2 does not disturb the operation of the unit. A case could be made that phase 3 is redundant and can be removed from the test pattern. Removing the test phase reduces test time by 30% and simultaneously simplifies the test pattern, minimizes test complexity and ultimately reduces database storage requirements. This was thought to be a highly attractive justification for eliminating the test phase. However, this was ultimately rejected. The technical experts determined that reducing the Hi Pot test by a factor of seconds was not a value-added use of their software resources.

4.4 Summary of Results

Two data analysis techniques were identified in this analysis. The first pass yield technique was thought to be beneficial in challenging the 100% test philosophy. Although the Hi Pot test was not currently performing at six-sigma levels and other external factors prevented a reduction in testing levels, first pass yield was identified as a leading indicator for challenging the 100% test philosophy in the remaining test system. The second technique identified was the phase Pareto analysis. Utilizing this technique, the value of each test phase is evaluated. From this analysis, it is clear that this technique can aid in determining unnecessary or redundant test phases. Again, this technique can be applied throughout the test system and across product lines.

Although no significant benefits were achieved from the Hi Pot analysis, the technical team embraced the new data analysis techniques. More importantly, the analysis completed here served as the test subject for the project and the learning's were applied to the remaining test system analysis.

Chapter 5. System Tests 1 and 2 Performance Analysis

As discussed in Chapter 3, the purpose of the System Test is to ensure the unit operates as designed. The actual test pattern differences between the System Test 1 (T1) and System Test 2 (T2) are minimal. The primary difference being that T2 operates the device immediately after the Burn In phase, while the device is warm, in an attempt to screen thermally induced failures. Consequently, as the T1 analysis progressed, the similarity of the System Tests became an increasingly valuable factor. So much so, that the analysis evolved into a simultaneous study of T1 and T2 performance. Acknowledging the interdependence of the analysis demands that the subsequent discussion be addressed concurrently, rather than addressing each test independently.

In the initial analysis phase, it was learned that this test system was significantly unlike traditional testing schemes. The distinction was the addition of T2 to screen for thermally induced failures. Thus, it was important to validate the value-added by System Test 2. The value of T2 phases was challenged throughout the analysis and will be discussed in this section. Additionally, the discussion addresses new analysis techniques and opportunities for test time reduction.

5.1. First Pass Yield Analysis

The first portion of the System Test analysis proceeded much like the Hi Pot analysis. Each of the tests was subjected to a first pass yield analysis. The hypothesis was similar to that of Hi Pot as well, in that if the yield levels exhibited consistent performance and approached six-sigma levels, the 100% test philosophy would be challenged. An additional theory tested was that of the value-added or necessity of System Test 2. Given the System Tests were virtually the same, the added benefit to System Test 2 was the ability to test the unit at a higher operating temperature. It was thought that this test was more a test of product design and 'robustness' to temperature issue, as opposed to testing the unit's functionality. Should the findings from the analysis show that System Test 2 is providing no value, a recommendation to eliminate all or a portion of the test would be made.

Similar to the Hi Pot analysis two queries were created, one to provide test results on sequence number 1 and the other to provide the total number of units tested. As discussed in Chapter 3, section 3.3.2, there were four possible test outcomes for the System Tests. In this case, two outcomes required additional consideration: marginal and abort. When a unit receives a marginal test outcome, this is an indication that the unit is operating within the outer 10% of the product limits but the unit proceeds through the manufacturing operation without intervention. Given the lack of intervention, it was decided to treat the marginal test result as a pass outcome. The abort test result, on the other hand, was considered a failure. In most cases an abort outcome was due to operator error or test system failure. Accounting for the abort results as a failure was an attempt to gather information on operator variability and the impact a complex test system might be having on yields. The first pass yield data was calculated for System Test 1 and is represented in Figure 10:

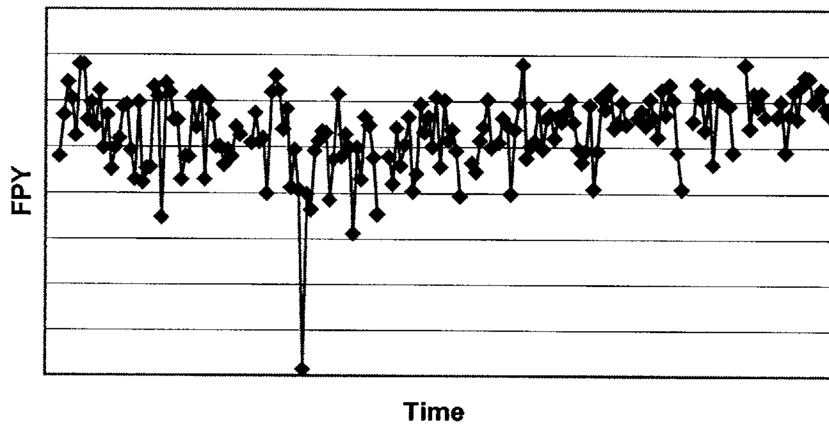

Figure 10 - System Test 1 FPY

It is apparent that although the process exhibits improvement over time, manufacturing was only now performing at original yield levels. Based on these performance levels, which were inconsistent and below six sigma levels, it was not possible to recommend a reduced testing strategy for System Test 1. The first pass yield data for System Test 2 was subsequently calculated. This data is represented in Figure 11

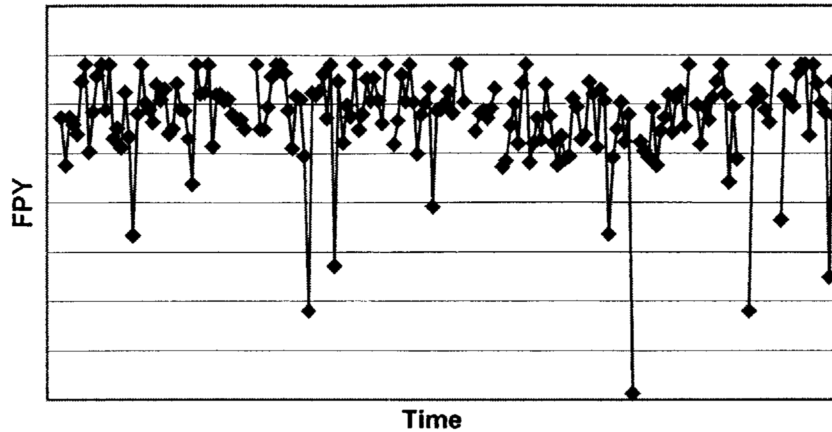

Figure 11 - System Test 2 FPY

For confidentiality purposes, the axes have been altered to mask performance levels; however, the performance levels for T2 are higher than for T1. This is to be expected, as System Test 1 should screen out most of the product or component failures. Again, the primary benefit between the System Tests is that the unit has been exposed to Burn In and is running at a higher temperature under T2. However, the test performance is clearly inconsistent and below six sigma levels, hence it is not a candidate for a reduced testing scheme either. Additionally, it is apparent that T2 continues to capture failures that T1 and Burn In have failed to eliminate. T2 was subsequently considered a value-adding test and would not be eliminated. However, the following phase Pareto analysis would be used to evaluate whether each phase in System Test 2 was necessary or if some could be eliminated for redundancy reasons.

5.2. Pareto Analysis

The next step in the analysis was to complete a phase Pareto of failures for T1 and T2, and compare the results. Given the similarity of the tests, the comparison of the failures by test phase should identify failure modes in phases of T2 that are specifically impacted by the thermal change in the unit. As in any production system, the ideal test system would identify failed components early in the process, before significant value was added to the product. In this case, it was desirable to identify as many failures as possible prior to the Burn In test phase. The challenge was to identify those phases in T2 that screened-out failures and determine whether they added value in T1. Theoretically, T2 should capture the same failures as T1; therefore the T1 test phase may be the more appropriate candidate for elimination. The primary purpose was to

eliminate any overlap between the test phases and retain only those phases that identified failed components.

Once again queries were developed to extract the required data from the database.

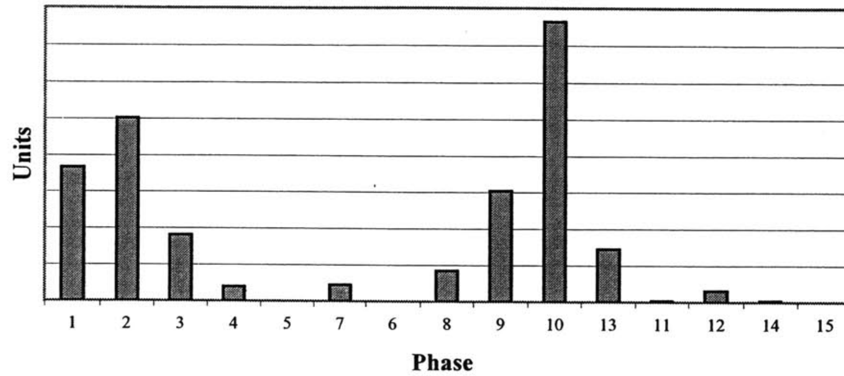

Figure 12 - System Test 1 Failure Pareto by Test Phase

The phase Pareto of failures for System Test 1 and 2 are shown in Figure 12 and Figure 13:

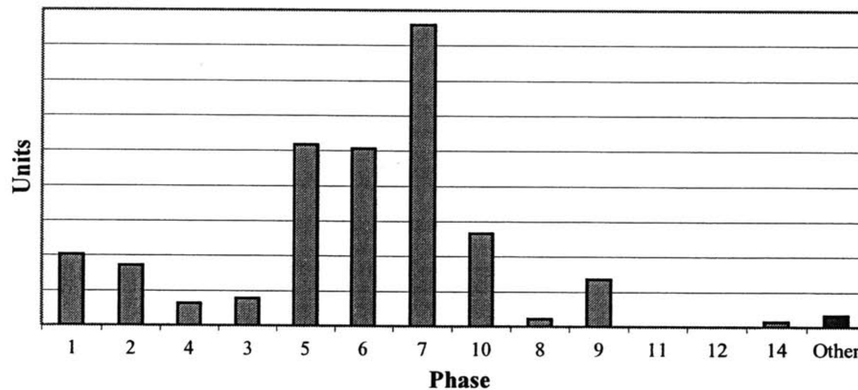

Figure 13 - System Test 2 Failure Pareto by Test Phase

The obvious disorder of the x-axis on both charts are attributed to three things, the first being the order in which the phases are run i.e. Phase 3 is not necessarily run prior to Phase 4. Second are discrepancies in designation between T1 and T2. Although the tests are virtually identical, phase 1 in T1 does not necessarily correspond to phase 1 in T2. Third, not all of the phases are present in both tests and the numbering system does not account for that variation. Additionally, the 'other' category is failures attributed to T2 but for which the phase is non-existent. Consequently, the order of the axis in the

charts above is the true representation of the order in which the test phases are run. A comprehensive description of each test phase was presented in Chapter 3.

Given the nomenclature discrepancies, a direct comparison of the two charts was difficult at best. Additional charts were created to facilitate the comparison by phase. In these charts, phases are referred to by test phase description. In an effort to maintain the discussion on the analysis method, as opposed to specific findings, one example analysis will be discussed from this point forward. Specific results of the analysis will be discussed in the final section of this chapter, 5.3. An example is shown in Figure 14[5]:

Figure 14 - T1 vs. T2 by Test Phase

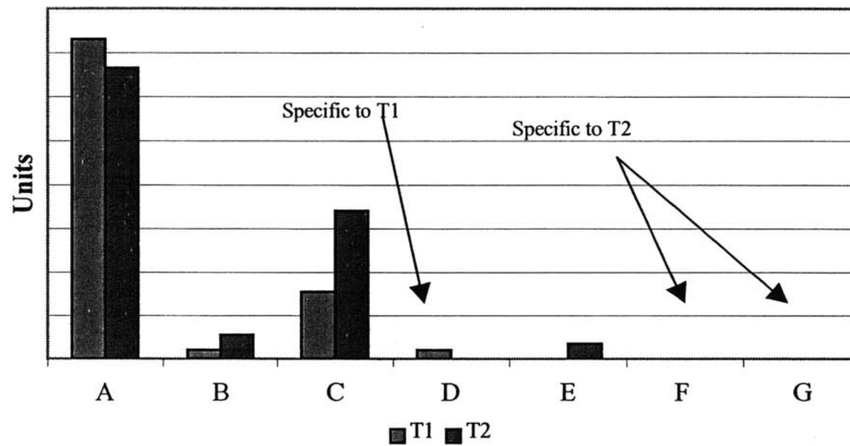

The chart highlights an important issue: phases that are exclusive to System Test 1 or System Test 2. Of the remaining phases, A and C each have a significant failure level in both T1 and T2. Clearly, the large number of failures in T2 indicates thermal degradation and validates the need for T2. However, given the equally large number of failures in T1 and manufacturing's desire to eliminate defects prior to the Burn In test, this phase clearly can not be eliminated from T1 either.

Phase B on the other hand, had a lower number of defects in T1 and could potentially be eliminated without great risk to manufacturing. Similarly phase E had no defects in T1 and should be considered for elimination. Accordingly, a second query was designed to capture an additional level of information for the two phases in question. Within each of these phases were sub-sequences where the unit was subjected to additional test patterns and performance measurements. Analyzing at the sub-phase level allowed for an additional Pareto of failures.

The Pareto by sub-phase was created for phase B and is shown in Figure 15 and Figure 16 below. Note: each pattern depicts a unique failure mode. A difference in failure modes is noticed upon comparing T1 and T2, specifically in sub-phases a, b, and d. In this case, the only sub-phase that exhibited similar failure modes and could be a candidate for elimination was 'g'. Given this was a small percentage of the total phase sequence, the technical team chose not to make the required changes to the test sequence.

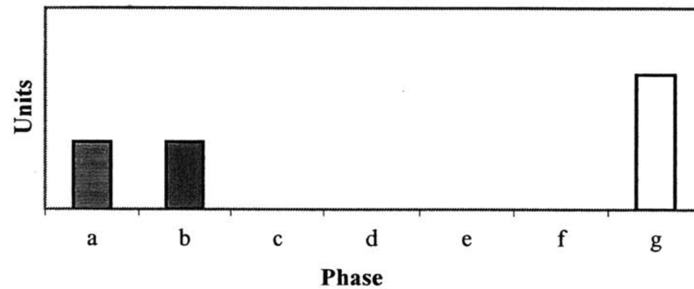

Figure 15 - System Test 1 Phase B Failure Pareto

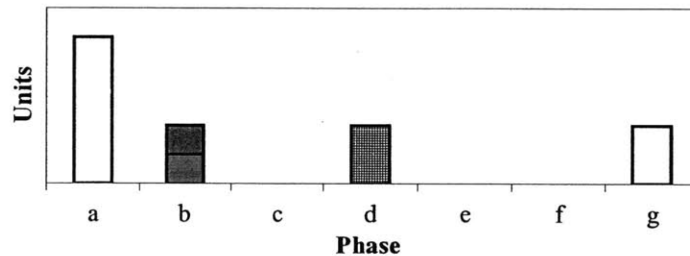

Figure 16 - System Test 2 Phase B Failure Pareto

As mentioned above, this analysis was completed for each phase of System Test 1 and System Test 2. Upon completion of the analysis, the technical experts were consulted. Each test phase and subsequent sub-phase Pareto was investigated for overlapping test patterns. Two significant issues were identified. The first was the fan failure mode. Fan failures represented a significant number of test failures. Given the simple design and mode of operation, the excessive failure rate was disturbing. Upon additional investigation, it was learned that the failures were primarily due to test system error and improper assembly.

The fan test system is relatively simple. An arm was lowered to the approximate height of the fan and an airflow measurement was taken. The results were compared to upper and lower specification limits. If the fan is not spinning or spinning in the wrong direction, the measurement would violate the specification limits and the unit would fail the test. Although the test is simple, it is not robust. Often as the arm lowered to fan height it would hit an obstruction, typically cabling, preventing the arm from completely reaching fan height. The airflow measurement is subsequently taken at the wrong height, causing a failure of the performance limit. Upon realizing the magnitude of test system errors, improvements to the fan test were discussed. Two ideas were seriously considered. The first was complete elimination of the test, citing that the unit would overheat in Burn In if the fan were not operating correctly. The second suggestion was to test the fan at assembly station 3, allowing the operator to replace the unit immediately if it was found to be defective. The first suggestion was considered the most fruitful and is to be piloted in production.

The remaining fan failures were due to improper installation. Upon further discussion with the technical group, it was revealed that there was a significant design for manufacturing/assembly issue. The unit had been designed in a way that allowed the fan to be installed with the fan spinning in the wrong direction. This issue would be addressed with the product design & development group to identify mistake-proof/fail-safe assembly procedures for current and future products.

The second issue identified was the order of test phases. The test system was designed with ramp-up phases to reduce stress on the unit. The need for this ramp-up was questioned. The team determined it could alter the order of the test phases to verify that the light load/ramp-up theory was indeed adding value. This benefit was a 30% reduction in test times.

5.3. Summary of Results

Several issues were raised in this analysis:

- Nomenclature discrepancies and inconsistent labeling practices between the System Tests made data analysis difficult. Addressing this issue would simplify analysis efforts in the future.

- Failures were attributed to test phases that were nonexistent. It was believed that there were 'bugs' in the test system software that needed to be addressed by the software provider.
- The miscellaneous failure mode accounted for a significant number of failures. The process to use the miscellaneous failure code should be constrained or completely eliminated.
- Extremely high levels of fan failures were attributed to System Test errors and product design and development.
- Validate the benefit of using the marginal qualifier in the test system. In an effort to simplify the test pattern, the qualifier could be eliminated from the software, as it is currently not used.
- The phase by phase Pareto analysis should become the method to analyze test system performance in the future. It was believed that component and product performance would continue to improve and that in the future, portions of T1 and T2 could be eliminated.
- Altering the order of test phases to verify the ramp-up theory was adding value. The analysis highlighted an opportunity for a 4-minute test time reduction, approximately 30%, of T1 and T2.

Chapter 6 Burn In Analysis

Although there are four different testing regimes within the manufacturing sequence and, presumably each one should add correspondingly towards the stated project goal of an overall test-time reduction of 50%, each test's contribution to this goal was biased by its underlying contribution to the pre-existing overall test-time. As Hi Pot accounted for only 7-8 seconds of the original 134 minutes of overall test-time, its contribution would be negligible. Additionally, System Tests 1 and 2 contributed approximately 10% to the original test time. Their contribution, although not negligible as Hi Pot's, would be severely limited by the absolute test times each test provided. From the outset, it was expected that Burn In would contribute the greatest test time reduction towards the goal as fully 90% of the original test-time was contributed by the Burn In phase. Based upon this, the testing procedures within the Burn In sequence were the focus of intense scrutiny and dissection.

6.1 Burn In Testing Yields

To begin the analysis of the Burn In testing procedures it was necessary to understand the quality and efficiency of the production system. To this end, the yields of the testing system were analyzed.[6] Although the specific data has been masked for proprietary reasons, the first pass yields of each of the four device sizes are represented in Figure 17 to Figure 20. The results discerned from examining this data were encouraging as they displayed a mature production system for a moderately complicated product. After reviewing the production rates, however, there remain a significant number of failure occurrences.

It is of additional interest to note that The Company is capable of ramping up production rates as yields quickly grow to normal rates. This is best characterized by the introduction of the larger frame sizes (see Figure 19 and Figure 20) where yields begin, not unexpectedly, markedly lower and over a period of weeks quickly attain typical production efficiencies.

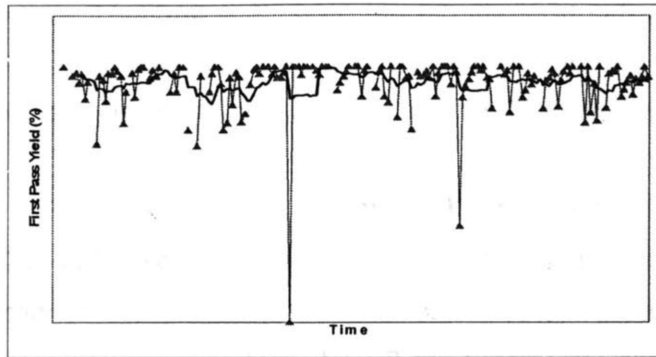

Figure 17 – FPY Device Size 1

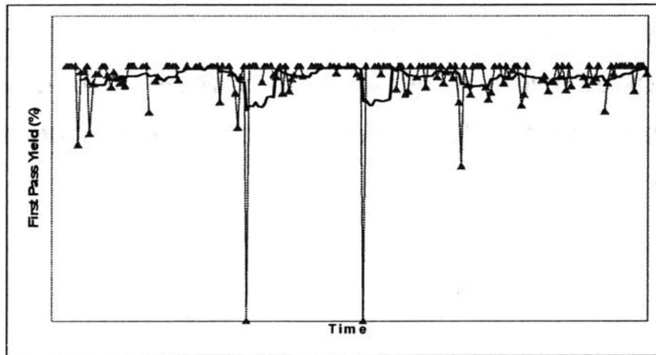

Figure 18 – FPY Device Size 2

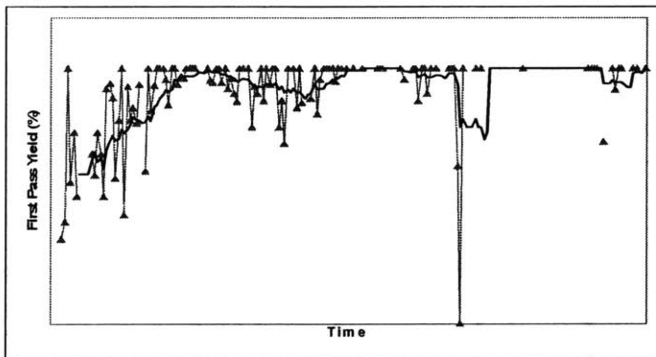

Figure 19 – FPY Device Size 3

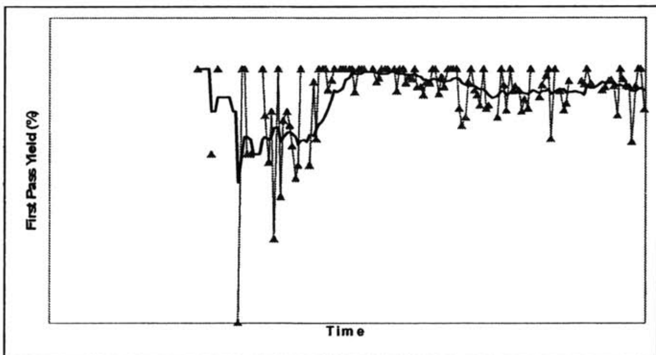

Figure 20 – FPY Device Size 4

Unfortunately, this analysis did not produce a particular frame size to be the main contributor to the overall test failure rates. The first pass yields are generally the same across frame sizes and no particular frame size is drastically worse than the others.

6.2 Failures vs. Time Spent in Sauna

At project inception the devices were subjected to a full 120 minutes of elevated temperature testing in the so-called "sauna". Based on the apparent random distribution of failures occurring throughout the Burn In test phase the next avenue of investigation was to understand the true time distribution rate of failures. As shown in Figure 4 the devices command the complementary devices to operate through a series of changing speeds. This speed cycle is designed to validate their operational abilities and production and component quality. Figure 21 shows the distribution of product failures as a function of time spent in the sauna.

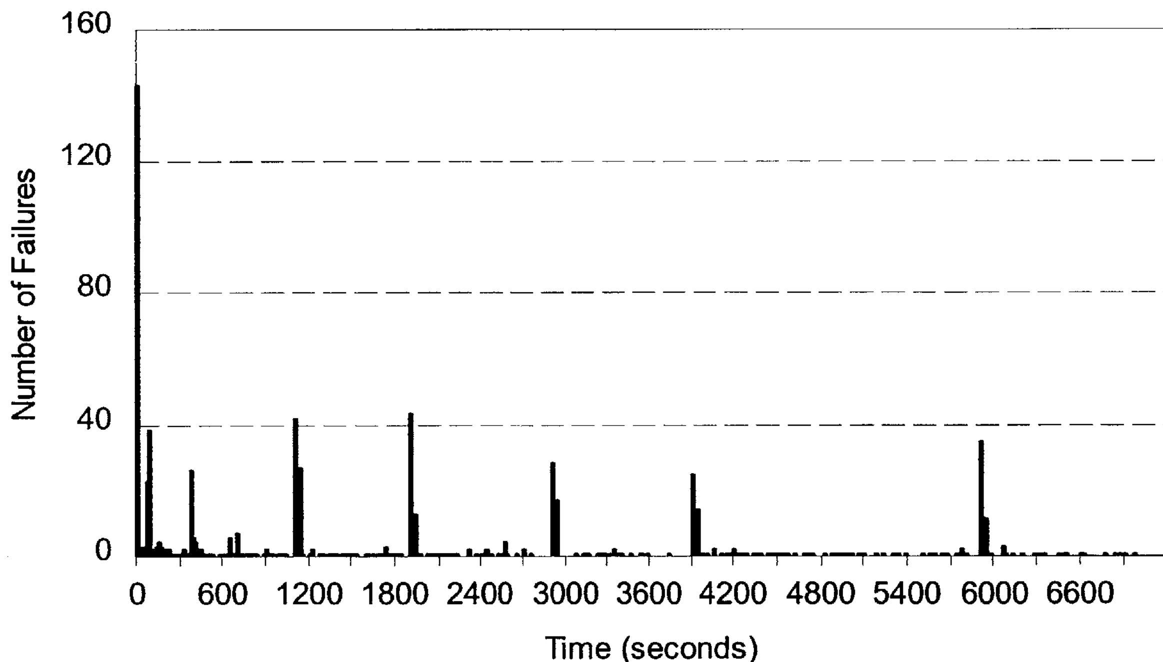

Figure 21 – Device Burn In Failures (no bins)

For ease of viewing and analysis the failures were also grouped into 30 and 60 second bins. This allowed further analysis of any time dependent macro effects to come to light. Both representations are shown in Figure 22 and Figure 23.

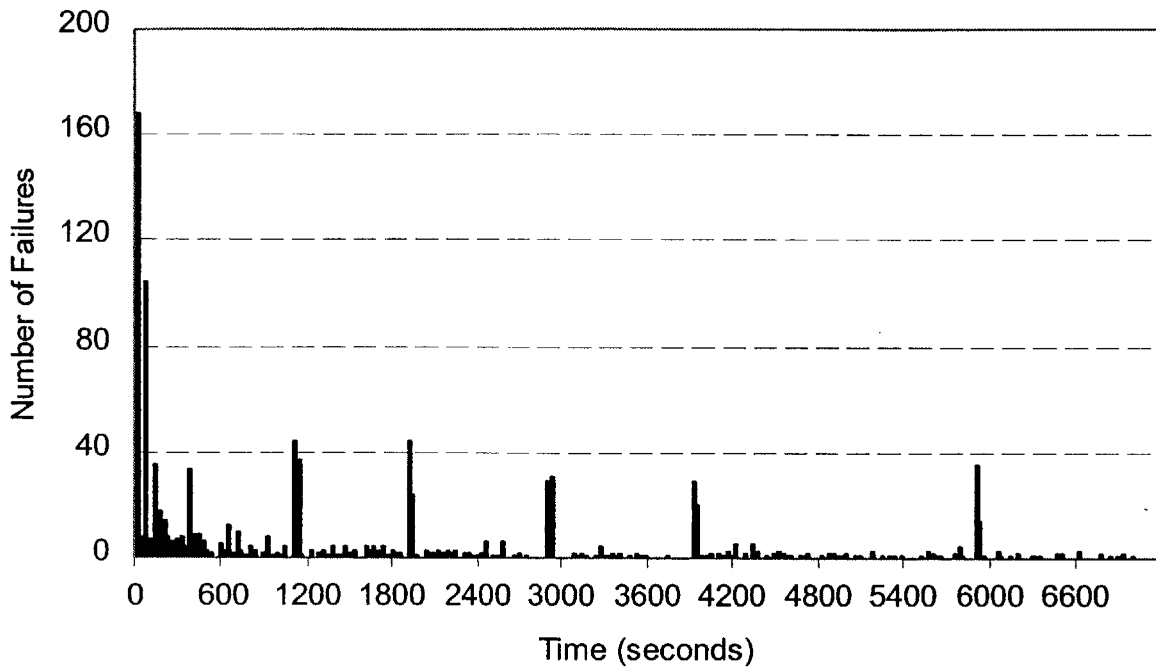

Figure 22 – Device Burn In Failures (30 second bins)

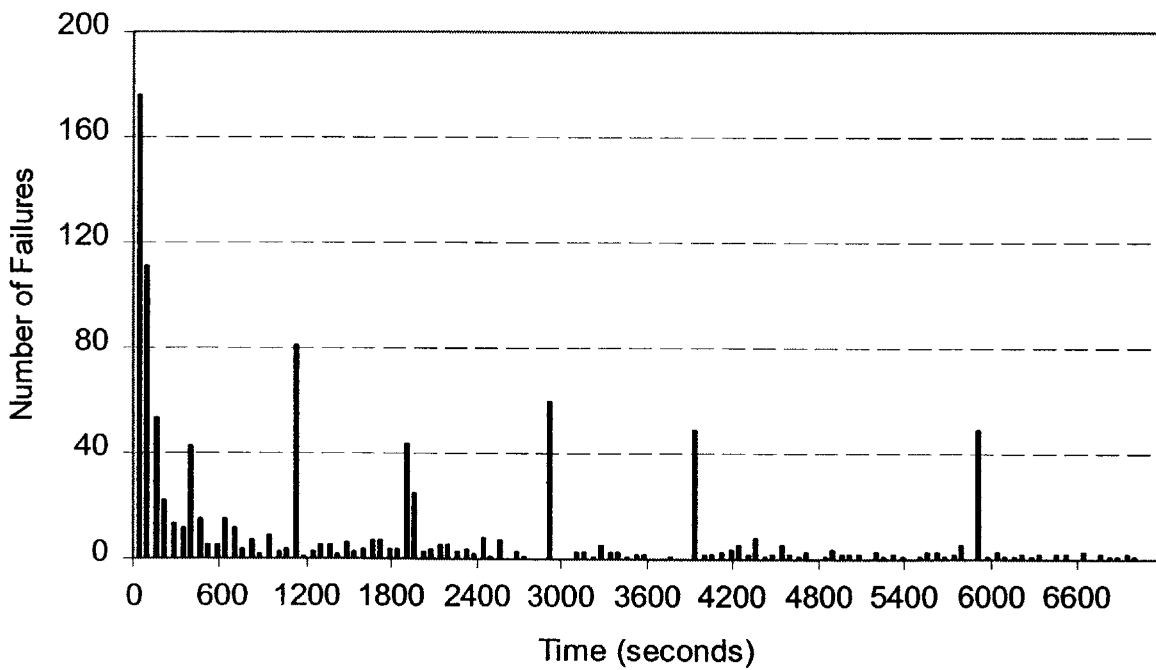

Figure 23 – Device Burn In Failures (60 second bins)

One thing that begins to become apparent from Figure 21 through Figure 23 is that the failure distributions consist of a series of large spikes surrounded by low level background 'noise'. These spikes are the first indication that unit failures occurring in the Burn In cycle are not randomly distributed, but in actuality have a time dependence. In addition, closer examination of Figure 22 and Figure 23 show that the spikes are actually two events occurring in close time succession.

6.2.1 Failure Rates vs. Commanded Speed

To better understand the correlation between these spikes and device operation the commanded speed cycle is overlaid upon the failure distribution (no bins). This is shown in Figure 24.

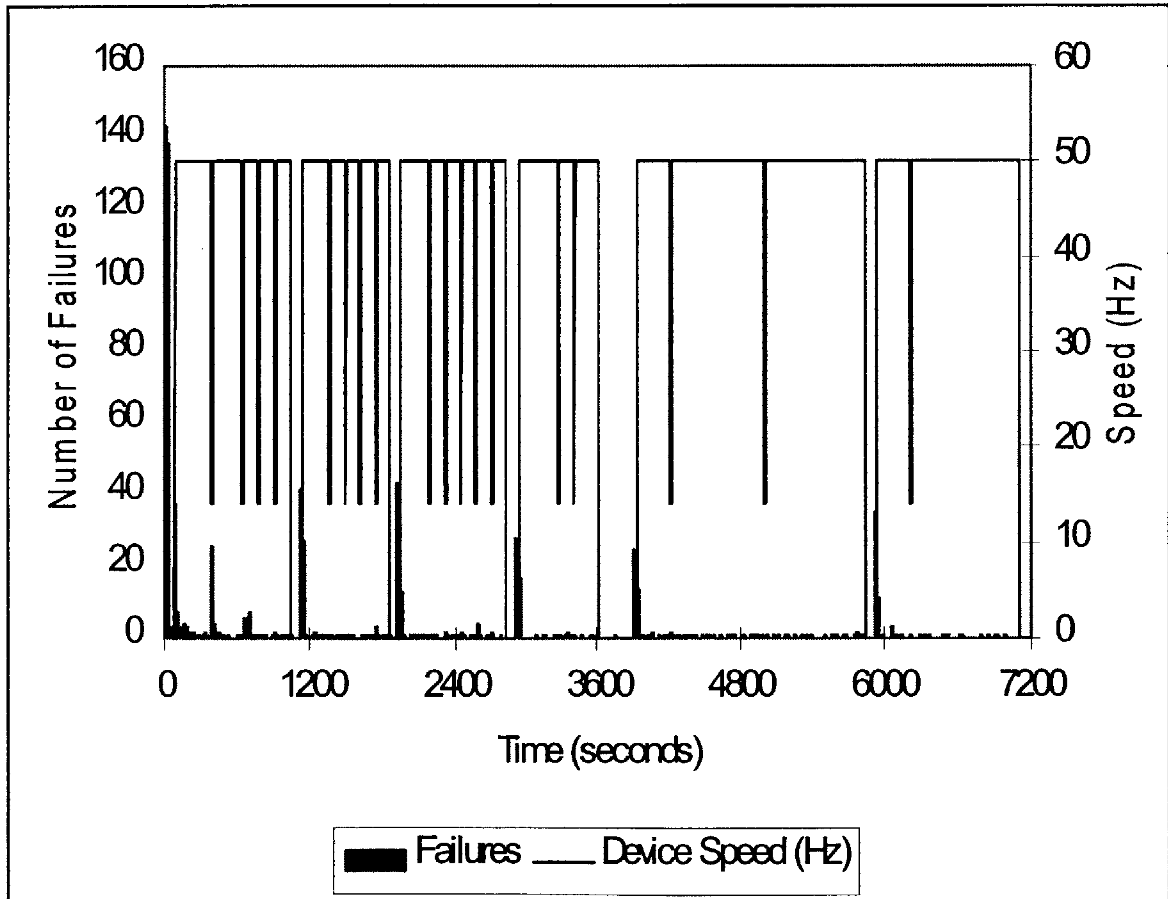

Figure 24 – Device Failures vs. Commanded Speed

A better understanding of the failure spikes begins to reveal itself from an examination of Figure 24. Each large spike occurs at a time when the operational cycle has commanded the complementary device to stop. Such events are intended to simulate a real world environment in which the complementary device may switch many times from full on to full off. The testing hypothesis here being that this is a much more life-like situation and, therefore, will better validate the quality and functionality of the device.

Apparently this is also a time that the tested devices have the most difficulty in adjusting to the stimulus that the systemic changes induce in them. This is not unexpected, however. It is anticipated that failures should occur as stimulus from external changes upset the internal systemic order which is achieved during long periods of steady state operation. Once this order is upset the system must adjust to a rapidly changing environment and settle into a new state.

Such a strong correlation exists between the 'complementary device off' state and the appearance of the failure spikes, that at each of the six commanded 'off' events the rate of failure is significantly higher than the existing background failure noise. Indeed, only one other spike exists outside of this set. This spike occurs at the first complementary device 'slow down' event at approximate time t=380 seconds. These complementary device slow downs are similar to the 'complementary device off' events in that they are conducted to simulate a real-world situation in which the complementary device speed may be externally instructed to change for some period of time. The complementary device speed is reduced to 15 Hz from 50 Hz during these events and then quickly returns to full speed. The device goes through 18 such slow downs throughout the Burn In cycle.

Since the complementary device 'slow down' events induce similar system changes as those in the full 'off' events, it is not unexpected that they would induce failures. However, it is interesting to note that only the first (and maybe the second) such 'slow down' actually does induce significant failures. Indeed, the remaining 'slow downs' do not appear to induce any failures discernible from the background failure noise. As the first slow down occurs only six minutes into full operation it is expected that early life failures associated with initial power up and operation may be the overriding failure mechanism and, therefore, not necessarily the same mechanism inducing the failures in the complementary device 'off' events. This hypothesis is supported by the existence of

only slightly more failures in the second 'slow down' event than discernible above the background failure rates than that which occurred in the first event. If infant mortality is the overriding failure mechanism for the slow down events it would be expected that the rate of failures would quickly tail off after the first few minutes of operation. This is exactly the phenomenon exhibited in the first two 'slow downs'. Of course, it is important to note that it is not exhibited for the complementary device 'off' events that may suggest the existence of other failure mechanisms.

6.2.2 Cumulative Failure Curve

Although the spikes are the most readily apparent artifacts on the failure curves, the presence of the background 'noise' and the rate at which failures are induced there are not represented by the analysis of these spikes. To understand the rate at which failures accumulate during the entire Burn In cycle cumulative failure curves were generated to graphically represent this. Figure 25 shows this effect.

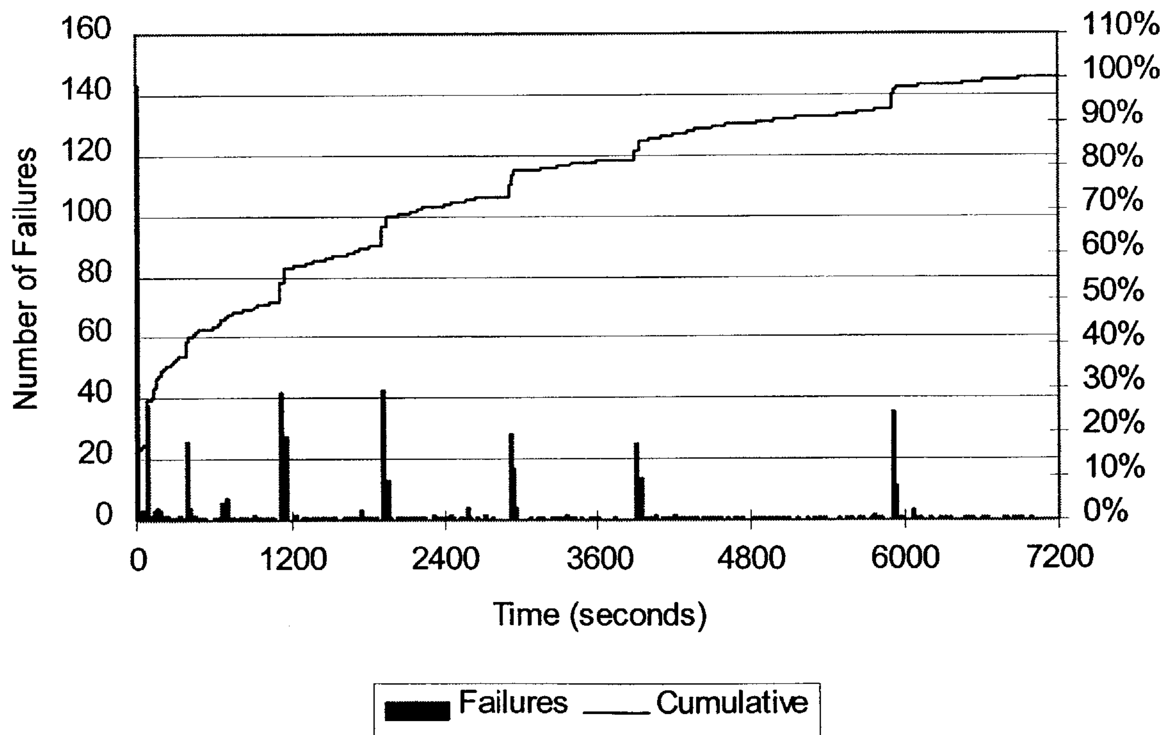

Figure 25 – Cumulative Failure Rate

The presence of the extremely high failure spike at the beginning of the test sequence dramatically affects the initial shape of the curve biasing it upwards and towards the right. Additionally, the high rate of failures in the first ten minutes of testing quickly push the failures to 40% of the entire two hour run time. Considering the goal of reducing the test time this is desirable however. Ideally, all of the failures would happen very quickly (assuming they will occur at all) at the beginning of the Burn In cycle, thereby, allowing the test time to be reduced to only that necessary to capture in-house all of the failures. So the rapid initial rise is very desirable. In fact, the goal should be to continually push this curve upwards and to the left so as to minimize Burn In time.

6.2.3 Analysis of Spike Components

Each of the spikes occurring after the first ten minutes of operation actually consists of a set of spikes. For ease of reference, these spikes are referred to as follows: Ia&b, IIa&b, IIIa&b, IVa&b, and Va&b. The Roman numeral refers to one of the five major spike events occurring at approximately t=1100, 1900, 2900, 3900, and 5900 seconds and the 'a' or 'b' refer to either the first or second spike component in each pair respectively.

Analysis of the online failure data shows that the 'a' component and 'b' components of each set are similarly spaced in time. They are separated by exactly 33 seconds leading to the assumption that a similar mechanism is driving the production of all spikes. After reviewing the dataset which controls the devices during the Burn In cycle it turns out that the 'a' component of each spike occurs at the end of the 'off' cycle when the device commands the complementary device to 'Power On'. After a period of time has elapsed for the unit to fully power up and prepare the system for full operation the full complementary device speed command is given. At each of the five spike sets this time between commands is exactly 33 seconds and precisely corresponds to the immediate onset of a failure event. It is interesting here that the 'Power Off' command is not the culprit inducing failures but actually the 'Power On' command and subsequent 'Full Speed' command are.

One cautionary note that was considered when analyzing the spike sets was that it was conceivable that the failure mechanism could be dependent upon the length of steady state operation between subsequent ON/OFF cycles. This would suggest a

result where the longer the device operated the larger the failure spikes afterwards. This condition was not realized though through a visual analysis of the failure spectrum. Here spike sets II, III, and IV are all separated by exactly 1000 seconds while spike set V is actually separated from III by 2000 seconds. V is similar in magnitude than I and II and only slightly larger than III suggesting that such a mechanism was not directly tied to the results.

6.2.4 Influence of Fault 24 on Failure Curves

After consultation with company engineers knowledgeable about the device's Burn In cycle and subsequent failure characteristics a particular failure mechanism explaining the curious spiking conditions was explored. Apparently a known failure mode referred to herein as Fault 24 (F24) might explain the predominant spiking effect. It was suspected that F24 failures were the result of Electromagnetic Interference (EMI) occurring internally to the device. The specific circuitry and current loads induced at power up were suspected of exacerbating the condition, thereby, resulting in a large component of the spike failures being the result of the F24 mechanism. In addition, as the F24 condition was largely understood, company engineers felt confident that a solution was at hand and, therefore, a solution to the known spiking was also at hand.

It was not known, however, whether the spiking condition was completely explained by the F24 mechanism. To better understand this the incidence of the faults were mapped and the results are shown in Figure 26.

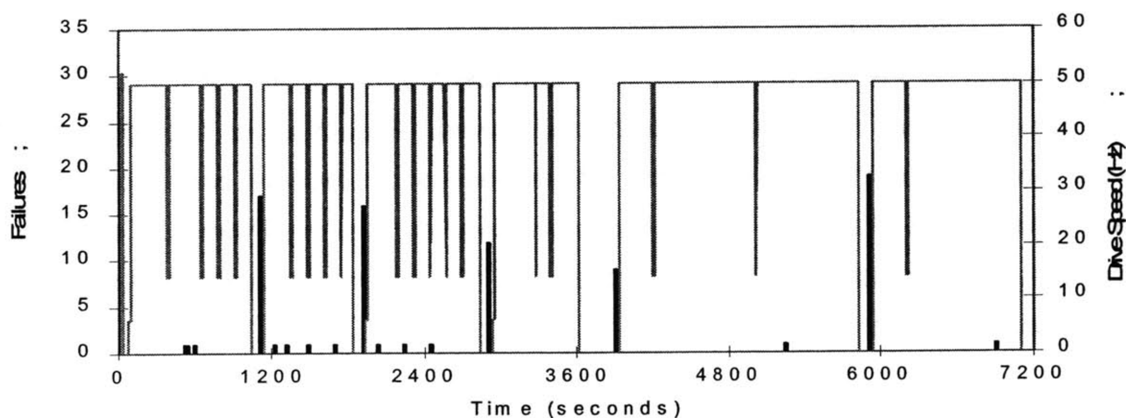

Figure 26 – F24 Failures vs. Command Speed

Figure 27 – F24 as a Percentage of Overall Failures

Figure 26 shows that F24 failures do indeed exist within the spikes, but curiously they only account for one of each spike pair (a). Additionally, Figure 27 shows the relative percentage of each singular spike accounted for by the F24 events. From this figure it can be seen that, in general, only about 40% of the a spikes are accounted for by F24, with Va being closer to 55%.

Figure 28 – Failure Curve Without F24

The removal of the F24 spikes is shown in Figure 28. Since, F24 was not the absolute culprit there remained other mechanisms inducing failures during the power up instances.

6.2.5 Temperature Effects

An overriding concern of the company, and the principle reason that the Burn In testing is performed in the 'sauna', is the assumed, degrading effects of an elevated thermal environment. Indeed in many instances it is believed that the industrial customers of the company will need the devices to operate at length in hostile environments. The environment is often challenged by high thermal conditions and also (although not specifically tested in production) dirty and foul as many industrial environments can be.

If the presupposition that the thermal environment is truly degrading to the device's performance is correct, than any reduction of overall Burn In time should only be carried out after verifying that the unit being tested has thermally soaked to a steady state thermal condition. If the test were reduced to the point that the unit was still heating up then additional thermally induced failures may not be captured by the Burn In test, thereby, releasing them into the field and enduring the pain of allowing the customer to detect them. As this is the whole point of conducting the Burn In test in the first place this does not appear to be an acceptable conclusion and, therefore, a good understanding of the thermal characteristics of the device while it is subjected to the sauna environment must be conducted.

Figure 29 shows thermal data obtained previously from within the company for the device under consideration. Here the largest frame size is assumed to be the worst case scenario as it is the most powerful device currently manufactured and, therefore, has the greatest internal heat generation capabilities.

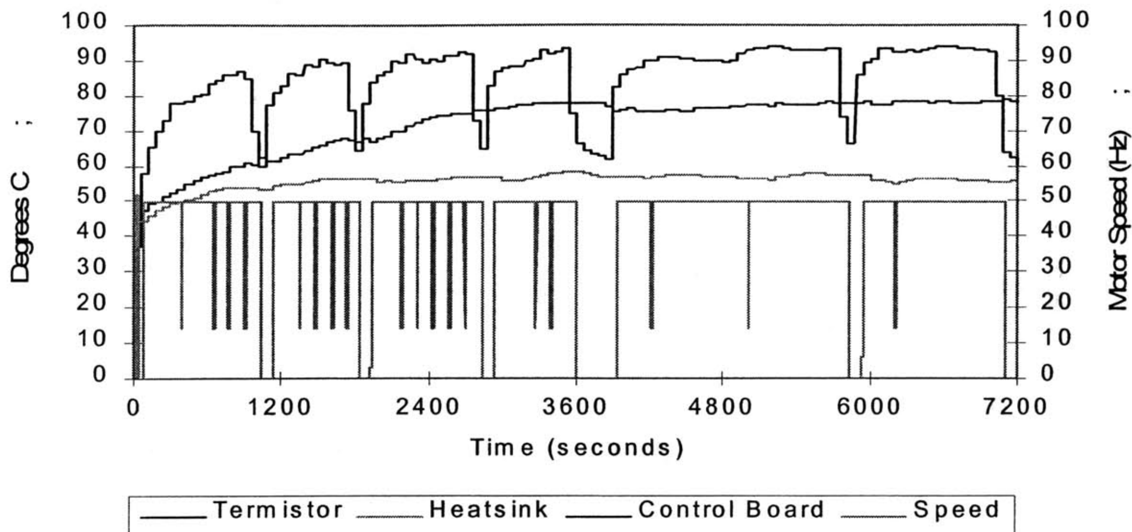

Figure 29 – Large Device Internal Temperature vs. Commanded Speed

Figure 29 shows the thermal characteristics of three representative locations internal to the device. The termistor component has a very fast thermal constant, which mirrors the power ON/OFF characteristics of the device. While the termistor sees the highest relative thermal environment it also has a very fast thermal time constant (approximately 900 seconds) and will be unlikely to suffer from a shorter thermal soaking. Likewise, the heatsink has a very fast thermal heat constant (approximately 600 seconds) and should be unaffected. Indeed, by its very nature, the heatsink should be unaffected and it nearly remains at the ambient temperature of the sauna anyway. Lastly, the temperature at the critical and sensitive control board appears to have the greatest thermal time constant (approximately 3000 seconds) and, therefore, is the most susceptible to a test time reduction.

Figure 30 highlights this relationship.

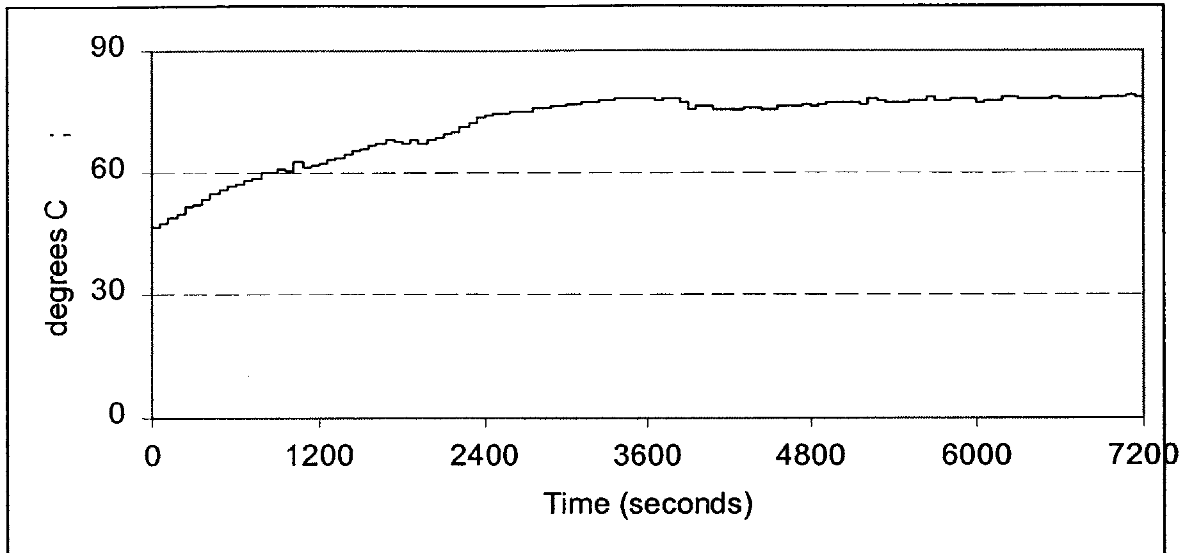

Figure 30 – Control Board Temperature vs. Test Time

From Figure 30 it is apparent that a test time reduction of 50% from the original time of two hours is feasible from the thermal standpoint. The device achieves a steady thermal state after about 50 minutes of operation; well under the goal of reducing the Burn In time to 60 minutes.

6.3 Fault Type Characterization

If the units fail during the Burn In test they are automatically removed from the chamber and routed to the repair area where they await failure analysis and further routing instructions from a trained repairperson. Failure analysis consists of a quick identification of the underlying problem and the assignment of a failure code, which gets manually entered into the test database. These codes are predefined and meant to capture the underlying failure mechanisms for increased failure analysis. The codes are grouped into seven categories as Table 1 shows (see appendices for entire listing).

Fault Grouping	Base Fault Code
Component Faults	1xxxx
Assembly Faults	2xxxx
Planning Faults	3xxxx
Soldering Robot Faults	4xxxx
Test System Faults	5xxxx
Conveyance System Faults	6xxxx
No Fault	7xxxx

Table 1 – Fault Codes

The notable groupings in Table 1 the Component faults (code 1xxxx), Assembly faults (code 2xxxx), Test System Faults (code 5xxxx), and No Faults (code 7xxxx). When failures can be attributed to the components or product assembly practices the test failure event is given the most appropriate code within each of the respective failure groupings. For instance, if a device fails while undergoing the Burn In test and the failure is attributed to the unit's control board being electrically defective it might receive a fault code of 10126. The failure here is attributed to one of the components, thereby receiving code family 1xxxx and the control board being electrically defective designating the remaining four digits (x0126). A similar family of codes exists for the other designations.

The interesting fault code families are the 5xxxx and 7xxxx series. These families represent failures directly attributable to the Test System and 'unknown' failures respectively. While it is important to capture failures in the test system so that system improvements might be considered and employed, and even on a very basic level of internal test system monitoring, it is desired that such failures would be minimal at best. In addition, failures which are not attributable to any known failure and are, therefore, 'mystery' failures are disconcerting and should be avoided. Interestingly, investigations into the latter failure events generally did not lead to further fault investigation, but to simply running the unit back through the test sequence to determine whether it was a 'nuisance' error or not.

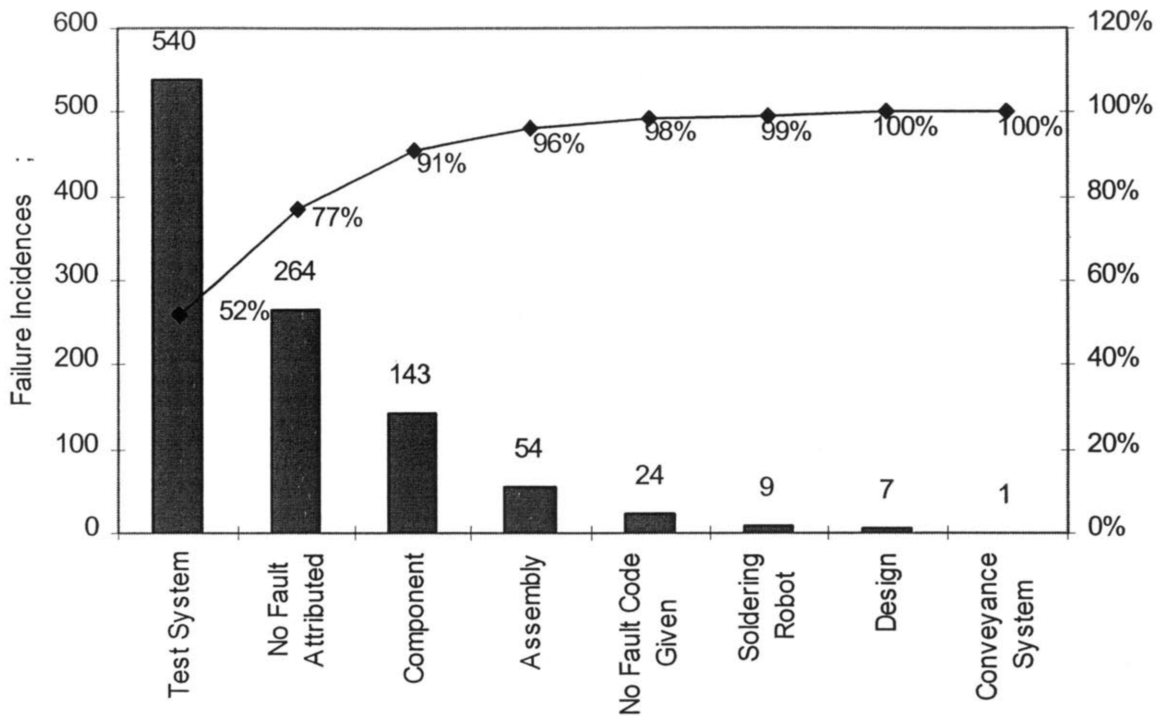

Figure 31 – Fault Code Grouping Pareto

Figure 31 shows the relative failure modes by fault code family. The most striking feature highlighted here is that failures attributed to the Burn In test system account for 52% of the entire failures recorded in the investigation period. This is unfortunate as the very system assigned with evaluating the product's quality may be masking the underlying problems through its domination of the failure event population.

Another striking aspect of Figure 31 is the secondary influence of the "No Fault" failures. No Faults account for 25% of the recorded failures. This is similarly disturbing in that a quarter of all failure events are not understood. There is little or no information as to whether they were true failures, system anomalies, or other events. Of course, the failures are, in one way or another, 'solved', as they must eventually pass the entire testing scheme. This, however, may take one or more testing iterations increasing critical throughput times and continually detracting from resources, as they require additional time and effort. The presence of these failures in such large numbers may be

even more insidious to the organization than the Test System Failures due to their unknown status and lack of corrective systemic action.

It is not until the third and fourth columns in the Pareto chart (Figure31) are reached that true device failures are identified. The Component and Assembly Failures, here, account for 14% and 5% respectively of the overall failures. These failures constitute the traditional purpose for Burn In testing, but amount to less than one out of five occurrences. The existence of component and assembly failures within the 'No Fault' groupings cannot be ruled out, however, this valuable artifact is lost within the unknown status of the group.

6.4 Test Reduction Hypothesis

The result of all of this analysis was to understand the underlying fault mechanisms and failure rates. In summary, the first pass yields of the Burn In test phase, while good, left some room for improvement. At the very least, if the Burn In time was to be reduced a full 50% to meet initial project goals, the yields should not be allowed to deteriorate. Hopefully, a new test regimen could be ascertained to maintain yields and quality with the benefit of a shorter test time.

In addition, the failure rates show that the overriding failure aspects are controlled by the device's On/Off command structure. These failures are not entirely explained by the F24 events and are, therefore, still not entirely understood. The spikes which constitute this structure significantly stand out from the background failure noise rate. This background noise rate tapers off slowly throughout the test time but does not appear to go to a rate of zero. Likewise, the initial 10-15 minutes of test time appear to be dominated by early life failure (infant mortality) syndrome causing the cumulative failure curve (Figure 25) to initially grow very rapidly and then grow slower, approaching a near constant growth rate.

If the goal is to reduce the overall test time to one hour from its initial length of time of two hours then this new time must be greater than the longest thermal time constant. Figure 31 ensures that this constraint will be maintained.

In light of the knowledge that a strong failure inducing mechanism was the On/Off structure of the test it was suspected that a new test which entailed 'shifting' these

failures earlier in the test would allow their capture at the same time as the test length was reduced. The goal in reducing Burn In testing is to induce as many failures as possible as early as possible, thereby ensuring the shortest test with the highest failure capture rates. This maximizes throughput while minimizing field failure problems. If all of the failures induced by the device's On/Off mechanism could be induced in exactly half the time it currently took the expected failure rates would be as shown in Figure 32.

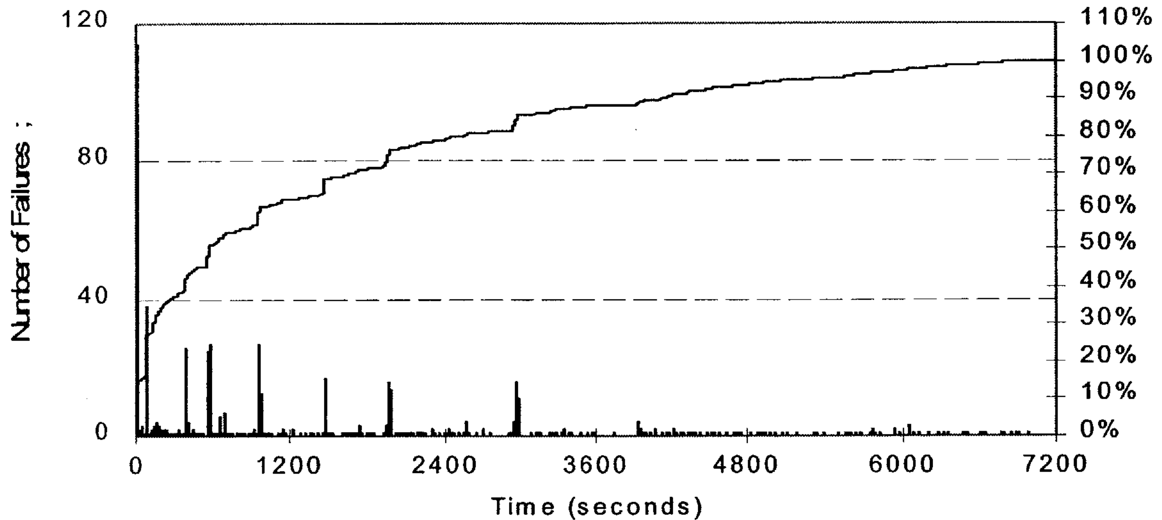

Figure 32 – Projected Cumulative Failures with Spikes Shifted

Figure 32 shows the failure rates expected if spikes I-V were all shifted 50% to the left. In addition, the cumulative failure curve is shown here to give an idea for how quickly the majority of the failures could be captured. This cumulative failure curve should be compared to that shown in Figure 25 for the current failure capture rates. Figure 33 shows this comparison.

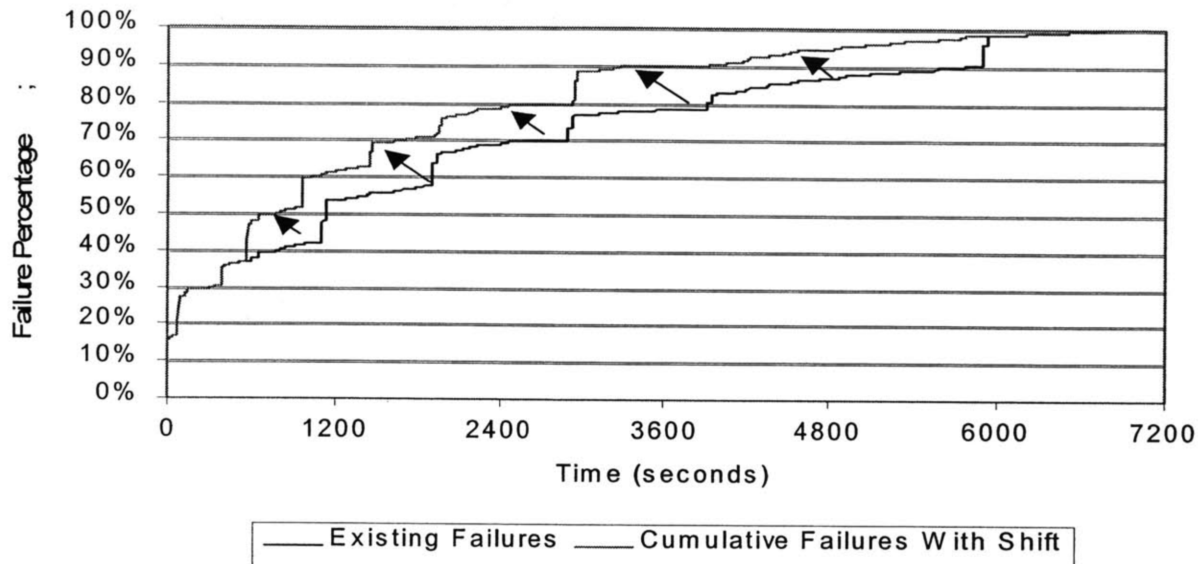

Figure 33 – Comparison Between Existing and Shifted Cumulative Curves

As Figure 33 shows, by shifting the failure spectrum to the left and inducing failures earlier in the Burn In test the expected cumulative failure curves shift correspondingly upwards and to the left. This expectation is crucial in shortening the test time while trying to capture the same or greater failure rates in house. Figure 33 reveals that fully 90% of the expected failures have been realized after 60 minutes in the shifted state whereas, in the unshifted state this milestone is not realized until after over 90 minutes.

It should be noted that only spikes I-V were shifted. As infant mortality appears to be the dominant affect early on, it was suspected that shifting would not necessarily have the corresponding effect that it would as the importance of infant mortality lessens later in the testing phase. In addition, since the shifting reduction factor is 50% there will be a decreasing influence on the overall time saved as the shifting gets earlier and earlier into the test. For instance, if a spike which occurs at t=3 minutes were shifted 50% to the right it would save 1.5 minutes. A spike at t=100 minutes saves 50 minutes. So there is are decreasing returns the earlier in the test the shifting is applied. For both of these factors only spikes occurring later in the test sequence (I-V) were shifted.

Chapter 7 Product Performance Analysis

The final and most critical aspect to redefining test philosophies was analyzing product performance at the customer site, also described as 'field performance'. Analyzing field performance data was thought to be the most accurate method of validating the current test system methodology, as prior to shipment every unit produced is subjected to the 4-phase test regime. Even with such aggressive test standards, units fail upon start-up and over the 24-month warranty period. Thus, in analyzing product returns, one might find a 'weak link' in the existing test system philosophy.

When a customer experiences problems with a product, the customer contacts technical field service. Technical service representatives are trained to minimize the impact to the customer by quickly providing the customer with a functional product (assuming a quick fix is not available). If the product is still under warranty, usually 24 months, technical field service sends a replacement unit to the customer and the customer returns the failed unit to the field service office. When the 'problem' unit is received at field service, the local experts attempt to repair the product but most often send the unit back to the factory for additional analysis. When the unit is received at the factory, it is subjected to a battery of tests to determine the root cause of failure. Formal records are maintained for all returned products. This data is seen as the most accurate data available and was used as the basis for the performance analysis. Obviously this data although thought to be the most comprehensive, lacks information regarding failures that the customer or field service technicians were able to repair and includes only units that were returned to the factory.

This chapter reviews field failure data and discusses innovative analysis techniques developed to analyze field data. Additionally, opportunities are identified to utilize marginal qualifier data gathered during the test sequence. Finally, recommendations are developed to more effectively utilize test performance data in conjunction with field failure data.

7.1. Initial Findings

When a failed unit is returned to the factory, a failure analysis/root cause analysis is performed. When the reason for failure is determined, a fault code is assigned to the unit. To maintain confidentiality, a letter has been assigned to each fault code as

opposed to a description of the failure. Figure 34 represents the root cause data collected at the internship site:

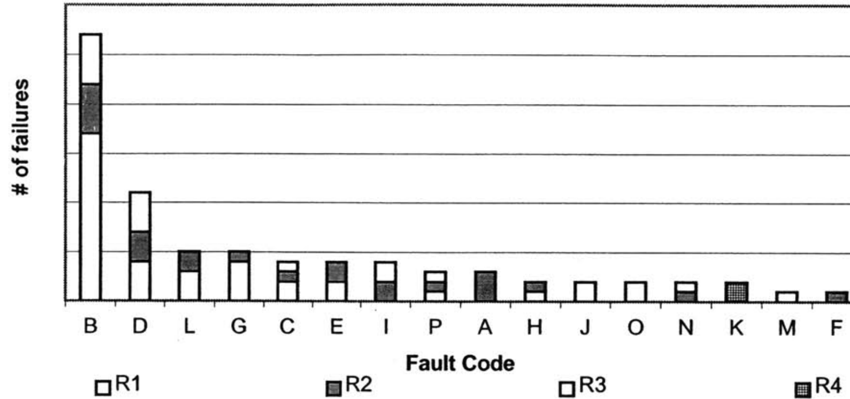

Figure 34 - Fault Code as a Function of Frame Size

This chart depicts failure code as a function of frame size. The numbers are nominal and not adjusted for volume. Adjusting for volume does not identify a trend between fault code and frame size. The data presented spans a 6-month period and 16 fault codes are represented in this chart. One issue became immediately clear after looking at the data – there wasn't enough data. Production commenced on this product line, approximately 9 months prior to this analysis. Given that the units are built to stock, it can take up to 2 months for a unit to reach a customer site. Another uncertain amount of time, at least 2 to 4 weeks, passes before the unit is installed. Consequently, it takes 3 to 4 months to receive data on product performance. An additional complication is the 24-month warranty period. Assuming failures occur throughout the warranty period, it takes approximately 27 months, after the date of manufacture, to obtain a complete picture of product and indirectly, test system performance. This data is acknowledged as incomplete and considered a lagging indicator. Any findings from this study will be subject to a sensitivity analysis to accommodate the above acknowledgment.

7.2. Performance Analysis

In an effort to further understand the relationship between field performance and test system performance, a query was written to extract test system performance on all returned units. Utilizing the product serial number as the search criteria, this was easily accomplished. The data revealed the following:

- 60% of returned units did not fail a production test, meaning they passed all production tests on sequence number 1 (a 100% cumulative first pass yield).

This number was unexpectedly high, as it was believed that units, which experienced problems during assembly, were more likely to fail in the field. Consequently, additional discussion and analysis was required to substantiate the finding. This 60% will herein be referred to as sample one. In the original analysis, units that received a marginal qualifier in System Test 1 or 2 were considered a passed unit. As mentioned in section 3.3.2, units receive a marginal test outcome in T1 or T2 when the unit is performing at the outer 10% of the specification limits. Given these units are regarded just as those that receive a pass on each test, it was important to investigate if a relationship existed between marginal units and returned units. Figure 35 was constructed to present the test results for units in sample 1; the data is nominal and not adjusted for volume

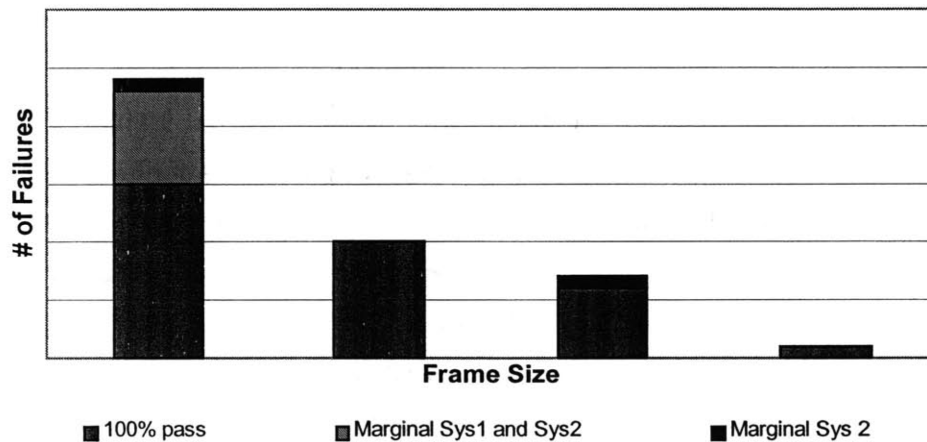

considerations.

Figure 35 - Test Results: Units that passed all quality tests

The figure above depicts three characteristics: units that received a pass on all tests, units that received a marginal pass on T1 and T2, and units that received a marginal pass on T2 only. No units were identified as having received a marginal pass on T1 only. As there were only a few units that received a marginal pass on T1 and/or T2, there does not seem to be a strong relationship between marginal units and field failure. At this time, marginal test performance was subsequently ruled out as a potential signal for field failure. However, as additional data is received, this relationship should continue to be monitored.

The next phase of the study was to examine the relationship between units in sample 1 and the assigned failure code. If a relationship were identified, it would become the basis for improvements and modifications to the test system. The intent was to identify a specific failure that the test system might be overlooking. The analysis is presented in Figure 36 below:

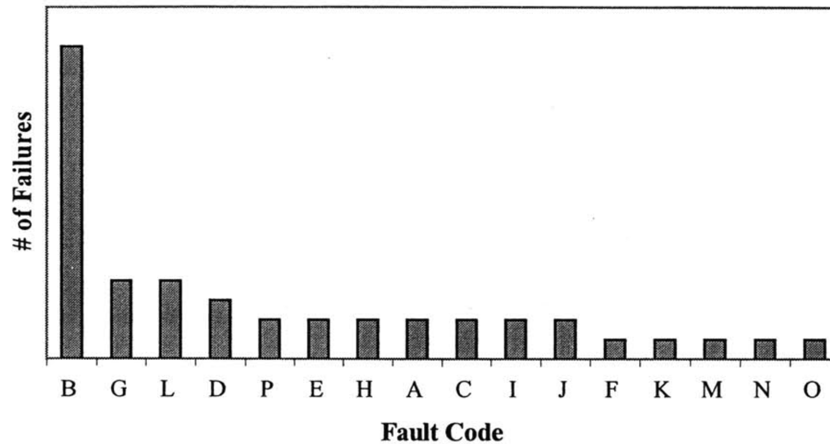

Figure 36 - Units with 100% Cumulative FPY

Each of the 16 failure categories is represented in the chart. There seems to be no relationship between units in sample 1 and the resultant failure mode.

Up to this point, no relationship had been established between units with 100% test performance (sample 1) and marginal units or fault codes. The next course of action was to analyze the remaining 40% of the sample, units that failed a production test, now termed sample 2. The first analysis was to identify which production test failed as a function of assigned fault code. For this see Figure 37 below:

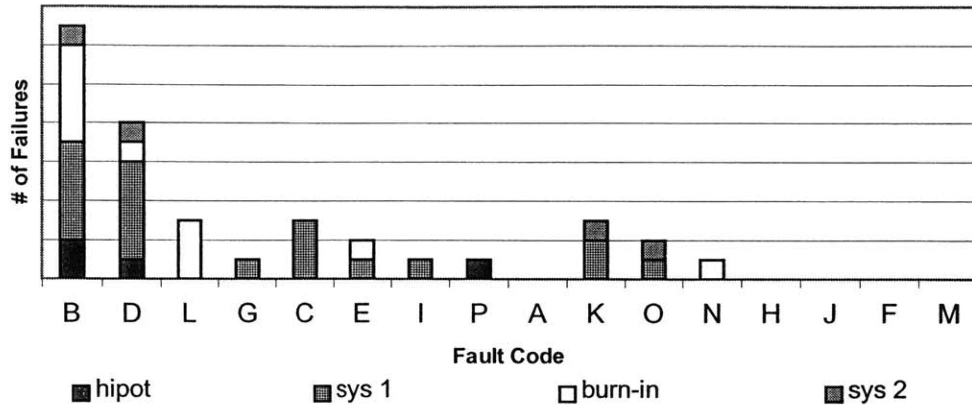

Figure 37 - Test Results: Units that Failed a Production Test

In this case, unlike the initial analysis, not all fault codes are represented in the data. No conclusive relationship was identified. Up to that point, the test performance of the units was examined. However, the database contained additional data, including the number of times a unit failed a specific test. In an attempt to ascertain a correlation between sample 2 and field performance, a subsequent analysis was conducted. It was believed that units that fail a test more once or fail more than one test, have a higher probability of failing in the field. If this were true, test performance could be used as a signal to identify potential field failures. The next step was to verify that the second sample contained a significant number of units exhibiting multiple failures. Figure 38 illustrates this data.

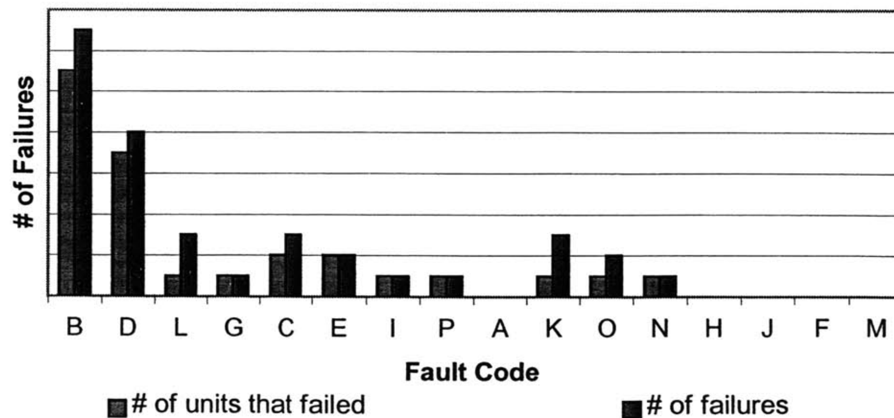

Figure 38 - Units that Failed Production Tests

Figure 38 depicts the following data: units returned from the field (sample 2) and number of test failures for those units. The instance where the number of failures is greater than the number of units indicates more than one test failure. For confidentiality purposes, the nominal failure data cannot be disclosed. It was determined that there were not a significant number of units that failed more than one test. It is acknowledged that this conclusion is not evident from the chart above, given that the nominal data is not disclosed.

An additional source of product performance data was discovered. The technical team performs a root cause analysis on all units that fail a production test. The data is collected by week for each test phase and is maintained in independent Excel spreadsheets. The analysis provides an additional level of interpretation on test system performance. This analysis distributes failures into additional categories such as: test system, component, assembly, and undefined. The data also compares product performance to the total number of units manufactured. In this case, field failure data would be examined by week of manufacture and compared to weekly test system performance. A query was created to provide week of manufacture data for all returned units. This data is shown in Figure 39 below:

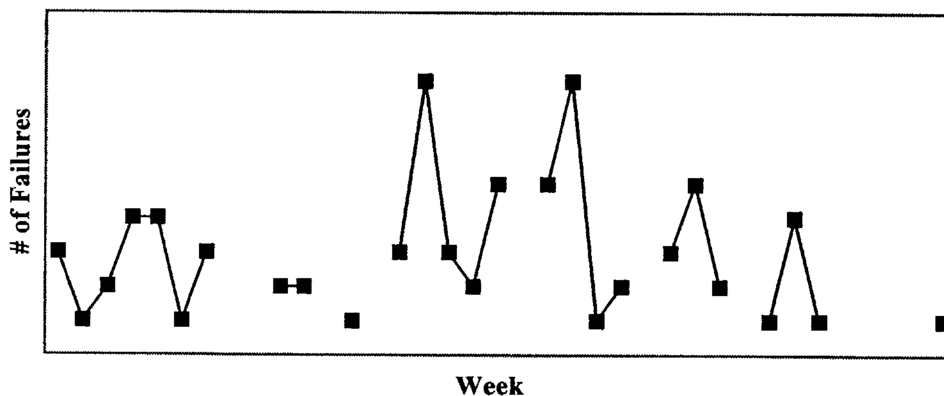

Figure 39 - Failures by Week of Manufacture

Figure 39 represents the number of field failures versus week of manufacture. Many weeks did not have any reported failure data and consequently the data is not continuous in the chart above. Most weeks show multiple field failures. It is important to recall that this data is incomplete until the 24-month warranty period has expired. Normalizing the data for volume yields the following data in Figure 40:

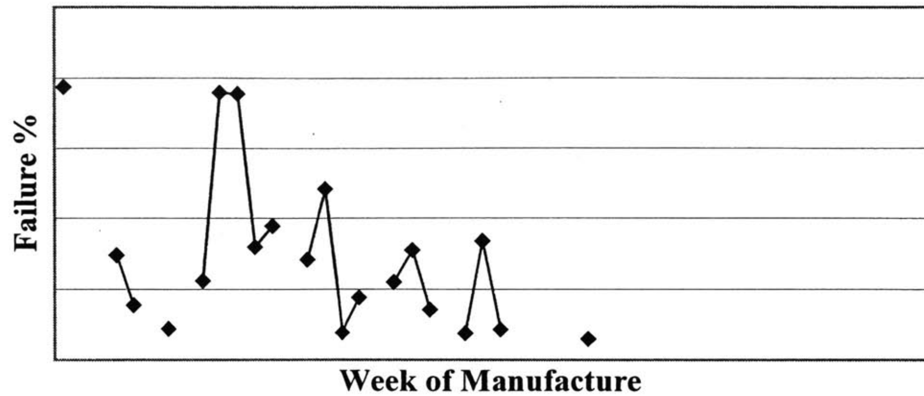

Figure 40 - Field Failure Percent by Week of Manufacture

The data collected by the technical team was compiled next. The team collected data by test phase. Figure 41 combines the data for all test phases:

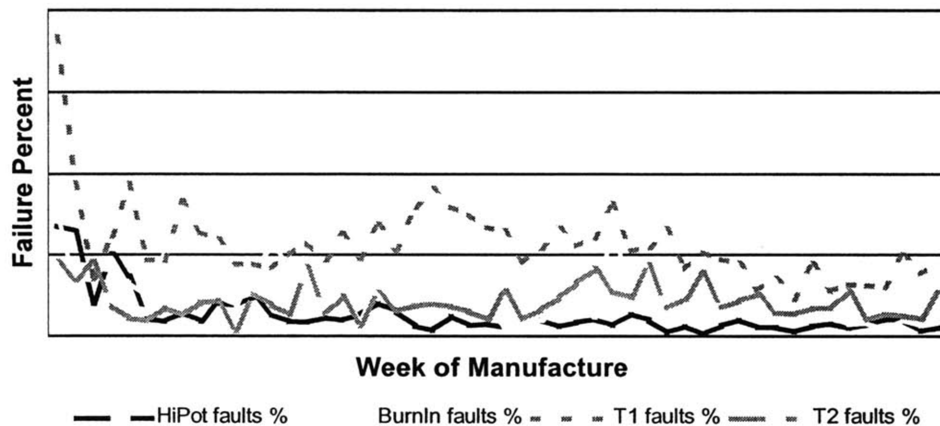

Figure 41 - Faults by Week of Manufacture

This chart allows us to analyze test performance as a function of manufacturing volume. Performance varies significantly each week, with T1 faults being the primary source of failures. Field failure data was then overlaid. This is shown in Figure 42 below. The purpose was to investigate if field returns varied with test performance. If failures varied with test performance, it might be possible to use weekly test performance data as an indicator for field failures.

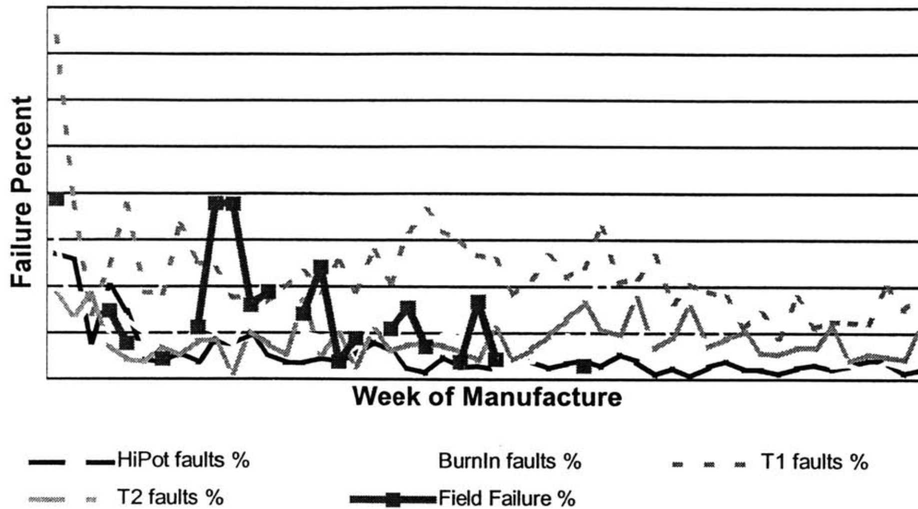

Figure 42 - Failure Percent by Week of Manufacture

As field failure percentages were significantly below test system percentages, the field failure data was adjusted by a nominal factor for ease of comparison. The relative magnitude of failure percentages is not vital in this case. Upon initial inspection, it looks as though field fàilures are trending similar to T2 performance, however, this assessment is based on limited field failure data and so an accurate assessment cannot be made. However, this technique exhibits the greatest potential in establishing a correlation between test system performance and field failures – assuming a relationship indeed exists. Unfortunately, it could take up to 15 months to gather enough field failure data to verify a relationship.

Several attempts have been made to correlate field failures with test system performance. The analyses have not been successful which can be attributed to various factors, the primary factor being field failure data. It has already been acknowledged that the data set is incomplete. The current data does not include failures repaired by the customer or field service representative. It is expected that failures will continue to occur over the warranty period, and finally, the data set utilized herein spans only 6 months – all of which happen to be during the production ramp-up phase. The recommendation at this time was to continue monitoring week of manufacture data for test system and field performance correlation.

7.3. *Additional Findings*

Upon closer inspection of the fault codes, it was apparent that there was not a predetermined standard for fault codes, codes were established as units were assigned a root cause. Accordingly, there was redundancy in failure codes. For example, there were three failure codes relating to a single component failure. The codes simply distinguish how or where the component failed. The most vital information was in the aggregation of the data, which identified the component failed on three occasions. Aggregation of the data reduces the failure codes from 16 to 5. A summary is shown in Table 2 below:

Table 2 - Aggregate Failure Types

Failure Type
No Failure Found
Customer Error
Software
Damaged or Defective Component
Manufacturing Error

Failures categorized as 'no failure found' and 'customer error' were known to be errors in which customers specified an incorrect unit for their application or the unit was incorrectly installed. If customers were not specifying the correct unit for their application, it could be a result of increasing complexity in the product, which required additional technical assistance. The sales department was to be contacted to raise awareness and improve customer education. In the case of improper installation, Research and Development was to be contacted so they might simplify or mistake-proof/fail-safe the installation procedure in existing and future product generations.

It was also recommended that Research and Development be contacted in the case of 'software' errors. The software errors discussed here were conflicts between the software structure and the operation of the unit. Such interdependencies should have been addressed during the design phase of the product, prior to releasing the product to manufacturing.

The two remaining categories: damaged or defective component and manufacturing errors are issues that can be addressed by manufacturing. Damaged or defective component failures were attributed to three things: supplier quality issues, damage at the customer site, and damage during the assembly process. Chapter 2 recommends several supplier programs that should alleviate supplier quality issues. Damage at the customer site is difficult to address. However, the recommendation to sales and marketing regarding improved customer education may help to prevent damage in the future. Finally, damage during the assembly process can be addressed by raising awareness on the manufacturing shop floor. Also within Chapter 2, are recommendations for storage and handling of the most delicate components, such as printed circuit boards.

Manufacturing errors comprised two failure modes: loose screws and soldering failures. Power tools were used to ensure screws were tightened to the proper torque, however, the tools either weren't operating properly or employees were neglecting to use them all together. As these failures were received, this information was relayed to shop floor supervision. Soldering failures on the other hand, as mentioned in chapter 3, were a complex issue. An automated soldering robot performs all soldering operations on the production line. Operators subsequently perform a visual inspection of the joints and are trained to hand solder if necessary. However, unsoldered pins continued to escape manufacturing and find their way into the customer's hands. A project team was assembled to address the complex soldering issues. Chapter 8 addresses these issues in detail.

7.4. Recommendations for Future Analysis

Although no specific concepts to re-define test methods were discovered in this analysis, there was much learning. Innovate methods to analyze data were developed and the process with which field failures are categorized was questioned. In summary, the field failure analysis highlighted the following issues:

- Cumulative first pass yield analysis showed the potential to re-define test system methods. As additional field failure data is received, production test performance should be analyzed to identify performance trends.
- The marginal test qualifier should be investigated as field failures are received. Due to time constraints the marginal analysis did not proceed to the next level but it is

perceived that marginal performance on explicit test specifications may be a signal for potential failure in the field.

- Compare week of manufacture data between test system performance (gathered by the technical team) and field failure data. The goal would be to identify a trend between test performance and field failures.
- Field failure data is incomplete due to the 2-3 month lag in information and the 24-month warranty period. Every effort should be made to obtain all repair data from customers and technical service representatives. Additional data would benefit the analysis process and would aid in identifying relationships between field failure data and the manufacturing process.
- Rationalize fault codes, in an effort to gain better understanding from field failure data and reduce redundancy.

Chapter 8: Solder Study

As discussed in Chapter 7, the field failure analysis indicated solder quality was a common problem in field failures. In response to the findings, a team was assembled to investigate the current soldering process and improve solder quality.

8.1 Solder Quality

The first task was to identify the magnitude of solder failures. This was accomplished utilizing three sources of data: field performance, weekly test system data and visual inspection. It was clear from this analysis that solder quality was a consistent failure mode throughout the production process and that determining the cause of solder failures would reduce test system failures, increase outgoing quality and reduce field failures.

8.1.1 Field Performance

As discussed in Chapter 7, when a unit is returned to the factory it is subjected to a root cause failure analysis. Of the units returned, 6.4% were attributed to unsoldered pins between the IGBT and the printed circuit board. Theoretically, the test system should identify these failures, as the unit should fail to perform if contact is not established between the IGBT and the printed circuit board. Utilizing the serial number from each unit, product performance in the test system was investigated. Each unit passed all System Tests upon initial presentation to the station; i.e. exhibited 100% cumulative first pass yield.

Further analysis of the failed units identified an additional source of information, clock time. Clock time identifies time in use prior to failure. In all cases, clock time on the unit was between 8 and 9 hours. Therefore, the IGBT was making contact with the board during initial operation and as identified above, throughout production testing. Since the pin is unsoldered the only means of contact between the two components is if the pin is leaning to one side (bent), making physical contact with the board. It was subsequently decided that after prolonged operation, either due to the dynamic thermal or vibrational environment, the pin migrates away from the board, causing the unit to fail. The test system can not be expected, nor is it designed to screen for such conditions.

8.1.2 Weekly Test System Data

The next source of data analyzed was the weekly data collected by the technical team. Each Monday the engineering staff generates a report from the test system to review the test failures from the prior week. Engineering specifically focuses on the comments entered by the operator at the repair station. From this report engineering creates a failure pareto for each test. Figure 43 below represents the weekly failures attributed to soldering for each test phase:

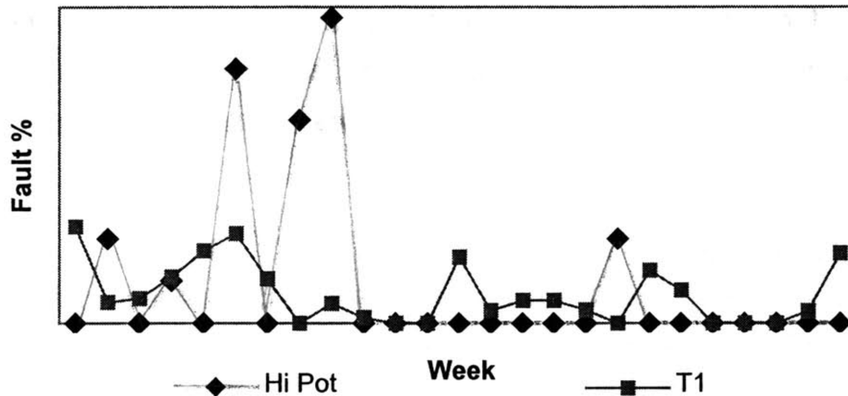

Figure 43 - Soldering Failures by Week

Soldering failures represent 5% and 7% of the failures in Hi Pot and System Test 1, respectively (12 week moving average). No soldering failures were identified in the Burn In or System Test 2 test phases.

8.1.3 Final Assembly

A final source of data was the solder quality check performed at the final assembly station. Operators are instructed to do the following:

1. Confirm all pins are soldered
2. Verify the solder material completely fills the hole and passes through to the other side of the board
3. Inspect the quality of the solder joint

A resource used by the operators during visual inspection is a photo of a printed circuit board. There are 4 photographs, one for each frame size, and each clearly identifies the solder joints to inspect. An example of this photo is shown in Figure 44:

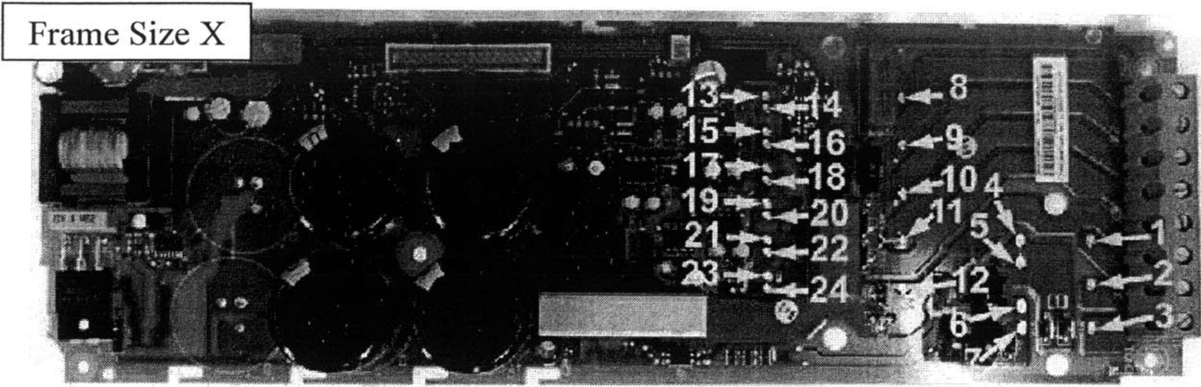

Figure 44 - Photo Identifying Solder Joints

Operators are also trained to hand solder, should they discover an inadequate joint or a pin that is not soldered. Each station is equipped with a soldering iron for these purposes.

If a board fails the visual inspection, the operators are instructed to record the frame size, pins or joints that failed, and the date and time on a data sheet. If soldering failures persist, the operator typically notifies the repairperson. Each week the data sheets are collected by the repairperson and recorded in an Excel spreadsheet. The spreadsheet calculates a weekly failure rate, based on the number of units that failed visual inspection divided by the total number of units produced. The data over an 8-week period is depicted in Figure 45 below.

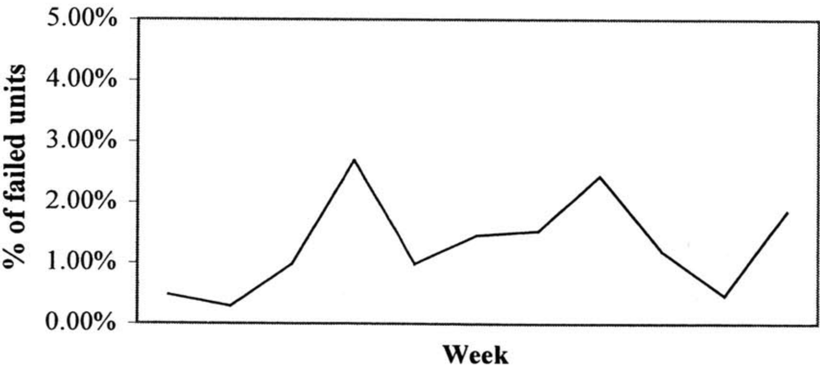

Figure 45 - Operator Inspection Data

The solder failure rate captured at the final assembly stage averages around 1% of units produced. In summary, solder quality is a consistent mode of failure throughout the production system and in field performance. The above data gathered from the final

assembly station, the on-line test system and field performance became the baseline for solder quality improvements.

8.2 The Soldering System

8.2.1 The Soldering Robot

A soldering robot performs the soldering operation. The robot has two soldering fixtures, also known as heads that support the soldering tip and solder wire. A second head is utilized to solder the smaller pins on the smallest frame size. This frame size is the only unit that requires the second solder head. A picture of the solder head fixture is shown in Figure 46 below:

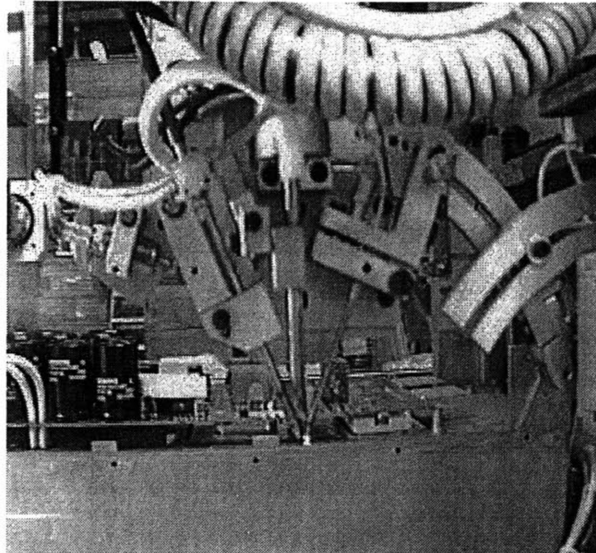

Figure 46 - Solder Fixture

A complex software program controls the robot. There are four distinct programs, one for each frame size. As a unit enters the robot, a signal containing frame size information is sent to the program. Each frame size program contains specific information regarding pin location, solder material requirements and cleaning frequency. Coordinates in the x (width), y (length) and z (height) directions are specified for each pin location on the printed circuit board (PCB). A time specification (in seconds) instructs the solder wire spindle to rotate, thus releasing a predetermined amount of solder wire to the solder tip. Coordinates in the x, y & z direction instruct the solder head to move to the cleaning station, where the tip is presented to a set of rollers which 'swipe' the tip to remove solder residue. The cleaning instruction is used only after multiple pins are

soldered. Additionally, the small device size program contains instructions on which head to use to solder specific pins. Head instructions such as 'park' head or 'get' head provide x, y and z coordinates for the head storage locations.

Upon entering the robot, the first instruction is alignment of the soldering head to the unit. This assures that the soldering head will accurately locate the pins to be soldered. The remaining instructions in the program vary by frame size and are a combination of the commands discussed above.

8.2.2 Soldering Tips

Two soldering tips, manufactured by Metcal are utilized in the soldering operation. The temperature of the soldering tips is maintained at 535°C. Metcal refused to disclose material composition of the tips, as it is patented and proprietary. However, SEM analysis[7] of the solder tip indicates the barrel material is composed of iron, chromium, and nickel; the shoulder material primarily of tin; and the tip itself a combination of tin and molybdenum, along with traces of iron and potentially silicon.

Metcal tips utilize Smart Heat® technology[8], which results in a high watt density heater that solders reliably and quickly and eliminates the need for calibration. Each tip is equipped with a self-regulating heater, which senses its own temperature and maintains idle temperature of ±2°F over the life of the cartridge. The soldering robot induces a high frequency current to the heater. The tip is comprised of two metals, one with high thermal and electrical conductivity and the other a magnetic material with relatively high resistance. As the frequency of the alternating current increases, the current flow becomes increasingly confined to the skin of the device, known as the skin effect relationship. This phenomenon serves the purpose of driving the current primarily through the high resistance magnetic layer, causing rapid heating. As the outer layer reaches a certain temperature, the Curie temperature, the layer loses its magnetic characteristics. The skin effect decreases, permitting migration of current into the highly conductive core. The effect of transitioning the material from a high resistance material to a low resistance material. Metcal controls the operating temperature of its tips by utilizing alloys with different Curie points, as the Curie temperature is set by the material structure and does not change over the life of the product. SEM analysis indicates that the Molybdenum is the secret to Metcal's success in maintaining temperature[9].

8.2.3 Solder Wire & Flux

A flux cored solder wire manufactured by Multicore is utilized in the soldering process. The material is a ternary alloy with a composition of 62% tin, 36% lead and 2% silver. The presence of silver is to aid in the creep strength of the solder joint and suppresses absorption of silver into the solder. The wire is a 5 core Crystal™ flux cored solder. The Crystal™ flux contains a balance of resins and activators formulated to leave clear residues and to resist fuming, with an approximate flux content of 2.7%. The wire has a melting temperature, solidus and liquidus, of 179°C. The minimum recommended soldering junction temperature is 239°C.

Two diameters of solder wire are utilized in the soldering system. The diameters are 0.7mm and 1.00mm. The smaller wire diameter, 0.7mm, is utilized for a subset of pins on the small device size only.

8.3 Soldering – The Existing System

Over the course of this study, two primary reasons for solder failures were identified. The first was board design vs. solder head/fixture geometry. The second was tip failure. The following sections discuss each failure mode in detail.

8.3.1 Board Design vs. Head Geometry

In interviewing the repairpersons it was discovered that component placement, pin location and head geometry are often in conflict. As can be seen in Figure 8.4, the solder head is a large fixture relative to component size and pin location. Often the solder head comes in close proximity to components during the soldering process. Upon closer inspection, it was noticed that skewed components could impede the motion of the solder head. In the case of a component that is not properly seated on the PCB, the solder head will make contact with the component, preventing the head from fully traversing in the z direction. Thus, the solder tip does not make contact with the pin and the joint is left unsoldered.

In the software program there is no process verification that indicates the solder head has achieved the specified xyz coordinates. Thus, the soldering robot proceeds with the soldering process, as though no failure has occurred. By incorporating this feedback loop in the software, one can add an alarm to notify the operator when the solder head has not reached the specified coordinates, preventing the unit from being released to the subsequent stage of production.

8.3.2 Tip Failure

As discussed in section 8.2.2, the Metcal tip has a self-regulating and calibrating temperature system. However, tips wear over time and in several cases, tips have caved-in upon themselves due to material failure. It is unlikely in these cases, that the self-regulating temperature system is any longer operative. In order to improve solder quality, it was thought that tips should be replaced prior to catastrophic material failure.

Thus, the operational maintenance schedule for the soldering robot was investigated. The findings were disturbing. Currently, the repairperson replaces a solder tip when:

a) Operators at the visual inspection stage complain about solder quality,
b) A tip has been in use for 48 hours, or
c) Maintenance is performed on the robot, typically a 2-day cycle.

This procedure yields a variety of outcomes. In the case of situation 'a', the tip is replaced after several failures have occurred. Under this situation, it is possible that many units have been exposed to this under-performing tip and there may be weak or missing solder joints present in units throughout the production system. This method also does not account for operator variability, as different operators will notify the repairperson at different failure levels. In situation 'b' and 'c', the tip can be changed prematurely. These situations do not account for the varying level of volumes the soldering robot might process in a 48 hour time period. At $100 per tip, changing a tip pre-maturely is a costly action and should be minimized.

In summary, there is no routine procedure for tip replacement, which could be a primary cause of solder failures. Developing a strict preventative maintenance schedule could increase solder quality and simultaneously reduce expenditures related to solder tips. However, there was no pre-determined time at which to change the tip. An effort was undertaken to define the 'right time' to change the tip. The preventative maintenance schedule for solder tips is discussed in section 8.4.

8.4 Solder Tip Preventative Maintenance

It was speculated that the appropriate time to change the solder tip was after a specified number of joints had been soldered. In an effort to establish this number, data was collected to determine the number of joints soldered per tip. As tips were replaced, the repairperson documented the date and time of replacement. This data would

establish the number of joints soldered per tip and aid in defining upper and lower limits of tip life. As the limits were identified, an IR camera would be utilized to model tip degradation across these boundaries. The resultant model would identify the number of joints at which the tip should be replaced. This specification would be incorporated into the on-line production system and/or the robot software to notify the repairperson when the solder tip has reached the solder joint boundaries.

8.4.1 Technique 1: Solder Tip Data Collection

As mentioned above, data was collected as tips were removed from the soldering robot. It was decided to collect data only on the large tip, as this tip was used across all frame sizes and accounted for more than 85% of all joints soldered.

Utilizing the time and date of replacement, it was possible to determine the number joints soldered over the life of the tip. A query was created to extract this data from the on-line production database. The data was collected from August 18th to December 1st. Over this period, tip life varied from 1500 to 12000 solder joints and ranged from 12 to 84 hours. The average number of joints soldered was 6449, with a standard deviation of 2588. This highly unstable performance further proves that the process to replace solder tips was inconsistent and extremely variable. The chart shown in Figure 47 was created from the data collected.

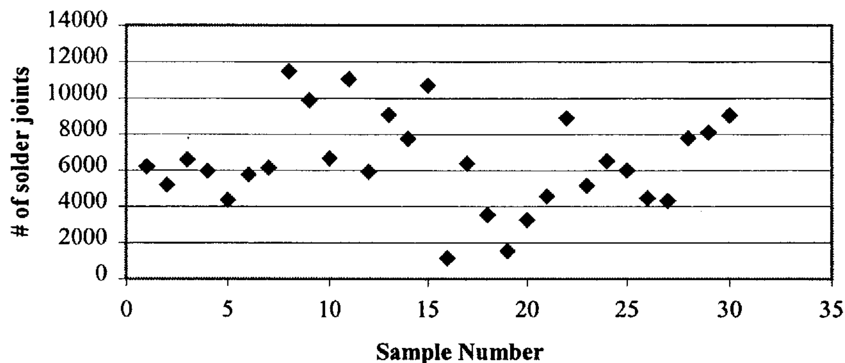

Figure 47 - Solder Tip Data

Early in the data collection process, the number of joints soldered varies around 6000 joints and is relatively consistent. This is attributed to increased repairperson awareness to the 48-hour replacement guideline. However, at sample 8, performance significantly shifts upward. It was at this time, that the 'summer intern' returned to Boston for mid-stream review. It is assumed that performance over this time period declined due to lack of attention on the system. Upon follow-up discussions with the repairperson's it

was also determined that some tips may not have been collected during this time, explaining the increase in the number of joints soldered. It is apparent at sample 16 that the 'intern' returned and so there is an increased awareness on tip changes – so much so, that tip's might have been prematurely changed during this time.

Utilizing this data is was impossible to establish tip life boundaries with any level of certainty. Thus, it was decided to utilize the IR camera to define average tip life, as it was a less subjective technique to monitor degradation in solder tip performance. Repairpersons continued to collect tip data and thus the chart above depicts data throughout the internship period.

8.4.2 Technique 2: Utilizing the IR Camera to define tip degradation

It was decided to model tip degradation over the 48-hour period, as opposed to utilizing the number of solder joints, as a basis for examination. Temperature data was gathered at 0, 24 and 48 hours of tip life. These increments were used as a basis to begin the study. Should evidence be found validating this technique, these time intervals would be reassessed. The intent was to identify thermal degradation due to material loss or failure. In the case where a tip outlasted the 48-hour period, subsequent data was gathered at 30-minute intervals.

Experiment Protocol:

Data was gathered utilizing an Inframetrics 600 IR camera. An example of camera output is shown in Figure 48. This image represents a new tip. As can be seen the output includes data regarding thermal variation, date and time of the image.

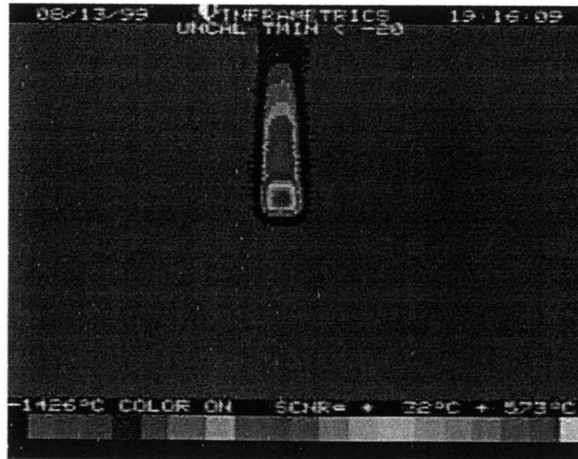

Figure 48 - IR Camera Output

The following conditions were utilized in the experiment:

- Temperature span: Set to extended range to capture center temperatures from 0 to +1000°C.
- Emissivity: Set to 0.83. A tip of known temperature was used to define the emissivity setting for the system. The tip was viewed utilizing the IR camera. When the temperature reading was equivalent to the known temperature of the tip, the emissivity was fixed.
- Tips were wiped clean, prior to collecting data in an attempt to minimize any impact solder wire residue might have on the above emissivity rating. It was acknowledged that this action could improve solder performance.
- The IR camera lens was held 12 inches from the solder tip for all measurements. The camera was then focused until the hot spot on the tip area became red and all white spots were eliminated. The temperature measurement was then considered to be accurate and could be compared across tips.
- In an effort to reduce measurement variability, the same two-person team took all measurements.
- Tips were replaced as they exhibited material wear or if solder joint performance was diminishing. The repairperson determined the appropriate time to replace the tip.

The data was collected over a 2-week period. Each day at approximately 14.30 the summer intern arrived at the soldering station to take the measurement. It is assumed that this interaction effected the outcome of the experimental results. Interestingly, most tips failed at or around the 48-hour time period during this 2 week span. The data collected is presented in Figure 49.

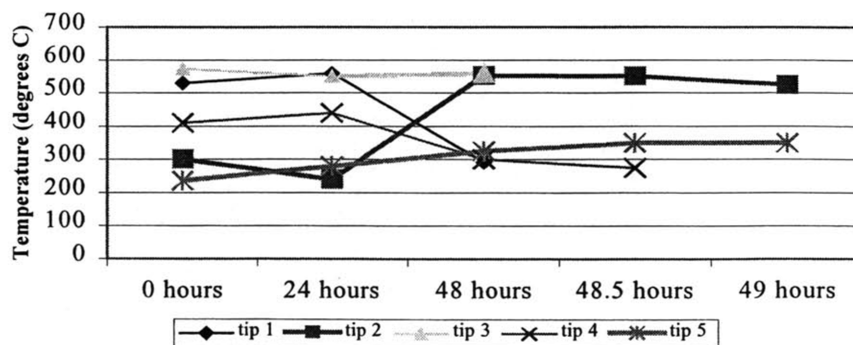

Figure 49 - Tip Temperature Over Time

Temperature data varied greatly across the samples. Tip temperature at 0 hours ranged from 250°C to 575°C. This variation was unusual and unexpected. Tips 1, 4, and 5 experienced an increase in temperature over the first 24 hours, while tips 2 and 3 experienced a slight degradation. Tips 1 and 4 continued to exhibit degradation over the second 24-hour period. This temperature was expected, as it was assumed that the tip would begin to lose temperature prior to failure. The remaining tips all experienced a temperature increase.

Finally for tips 2, 4 and 5, which surpassed the 48-hour time frame, additional measurements were taken at 30-minute intervals. Tip 5 exhibited a slight increase in temperature while tips 2 and 4 exhibited temperature degradation. In these three cases, as temperature measurements were taken at 30-minute intervals, it was expected that the tip would experience significant temperature loss prior to failure. In general, the temperature remained flat over this time period.

Tip temperature at failure varied from 275°C to 550°C. Again this outcome was unexpected. It was thought that the solder tips would fail at a temperature well below normal operating temperature. As some tips failed at 550°C, this was clearly not the case.

The data collected above produced unexpected results. Not only did all tips start at different temperatures and fail at different temperatures but they all exhibited varying temperature patterns over the 48-hour period. Clearly, this technique could not be used to model tip failure over the 48-hour time period.

8.4.3 Technique 3: Time constant evaluation

In a final attempt to determine the optimal time to replace the tip, re-generation of tip temperature was investigated. It is known that the tip experiences temperature degradation due to conduction during the soldering process. It was expected that over the life of the tip, the time required for the tip to regenerate to the soldering temperature increases. As this time increases, the tip may be unable to adequately recover temperature prior to soldering subsequent joints. If this were the case, joint quality would decrease or joints would be cold soldered.

Throughout the internship process, tips removed from the soldering station were collected and stored. Thus a representative sample of soldering tips existed, which could be utilized in the regeneration study. Each of these tips was assumed to exhibit an insufficient regeneration time, based on the number of joints soldered. The tips utilized ranged in tip life from 1,500 to 11,500 joints. Each tip was mounted in the solder fixture, the power to the unit was turned on and the IR camera was utilized to capture a thermal video of the regeneration time[10]. Upon a frame by frame analysis of the video, two measurements were taken: the time at which the tip reached 535ºC and the time at which the tip achieved a constant temperature, thus capturing the thermal signature of each tip. Each of these measurements is subject to variation, as the camera and solder unit had to be simultaneously powered to achieve an accurate time measurement. Measurements should be accurate to ±1 second. Figure 50 and Figure 51 represent these measurements, respectively:

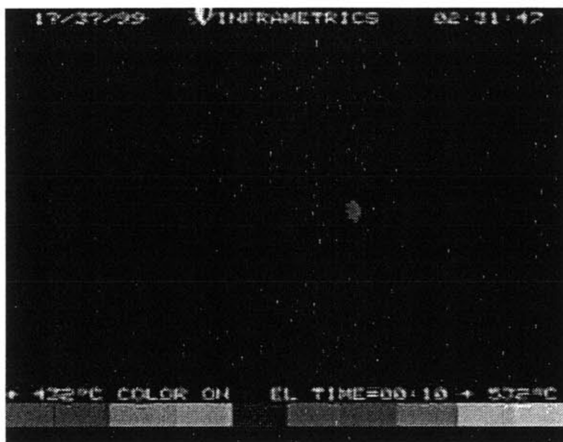

Figure 50 - Initial Sign of Heat

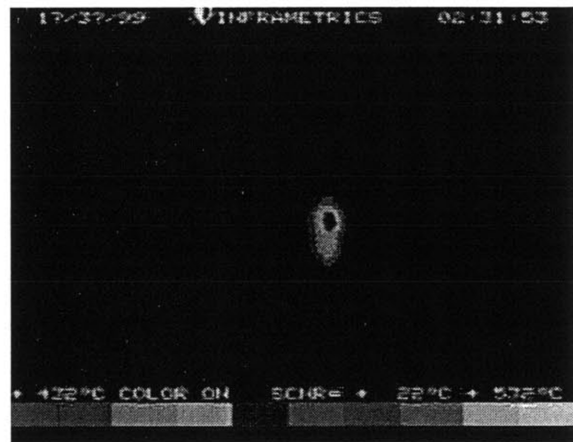

Figure 51 - Thermal Signature of Tip at Steady State

Figure 50 represents the initial sign of heat. Figure 51 represents the thermal signature of the tip at steady state. It should be noted that at steady state the tip does not exhibit the temperatures at the higher end of the temperature scale. This can be seen by examining the center of the image, which is the hottest area of the tip.

Figure 52 represents the times at which each tip reached 535ºC. The average time measured was 12.8 seconds.

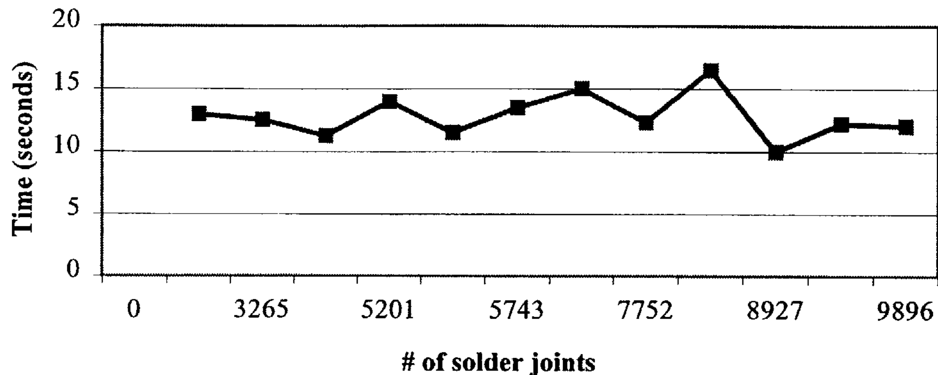

Figure 52 - Initial Sign of Heat

The results were unexpected. It was assumed that the time to reach 535°C would be more consistent across the tips, as each tip exhibited failure characteristics. Given the measurement variation discussed above, some of this 'noise' can be factored out. However, the varying degree of tip failure could play a significant factor. It is presumed that as a representative sample of the current tip maintenance schedule, a portion of these tips were removed prematurely while others were removed after exhibiting significant failure characteristics. At this time it is impossible to distinguish accurately between these tips, as this information was not collected.

Figure 53 depicts the time each tip achieved a stable temperature:

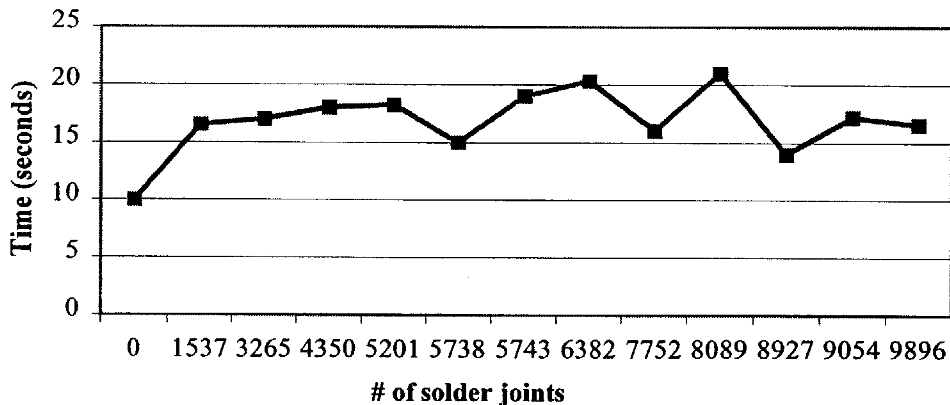

Figure 53 - Stable Temperature

These results were consistent with expectations, since it was assumed that as the number of joints soldered increased the time required to regenerate would also increase. However, not all of the data points fit the line above. It was determined that more data was necessary to better fit the curve and to rule out suspect data points.

8.4.4 Future Data Collection

It was recommended that the engineering team maintain a technical focus on the solder robot, to improve solder quality. The regeneration analysis seemed to be the most promising technique to define the solder tip replacement schedule. However, more data is needed to model the regeneration profile. This technique should be used as a basis to establish a strict preventative maintenance schedule for the repairperson. Additionally, the technical team should verify the maximum regeneration time that continues to yield acceptable solder joints. It is suspected that this time may vary by frame size. The product variation with the largest number of joints should be used to define this requirement, as the tip is exposed to the greatest number of regenerations during this cycle. Any preventative maintenance schedule that is based on hours of operation or joints soldered should be incorporated into the on-line production system or the software program, for the most accurate implementation. In this way, the procedure is less subjective and is simplified for the repairperson.

8.5 Summary

Failures due to solder quality were common throughout the production system. Two common failure modes for soldering were identified: board design vs. solder head geometry and tip failure due to an erratic replacement schedule. The primary issue to be dealt with at this point is implementing a maintenance policy that the repairpersons will follow on a consistent basis.

Board design vs. head geometry failures can be addressed in the following ways:

- Modify the software program to alarm if the solder head does not reach the specified xyz coordinates.
- Work with research & development on future PCB designs to account for solder head geometry.
- Work with the equipment provider to reduce the physical size of the solder head fixture.

Several attempts to correlate tip failure to product and process characteristics were unsuccessful. Although regeneration time is believed to be the most promising characteristic in predicting tip failure, it is believed that material failure is the ultimate driver of tip performance. The original SEM analysis[9] included a cross-sectional analysis of the soldering tip. However, the attempt to cross-section the sample caused the tip to separate from the solder barrel. Consequently it is believed that over time the solder tip

and solder barrel bond weakens. This was evident during the cross-sectioning process. Separation of the tip/heating element from the barrel could account for many of the tip failures, as separation of the materials would prevent the heating element from receiving the current necessary to generate heat. A time in use measure, accounting for the number of joints soldered, and the number of cool down and heat-up cycles should be a suitable predictor of this mode of failure.

Additionally, during the cross-sectioning process it was discovered that the section of the tip that protrudes into the barrel has a heating element wound around it. Visual inspection of the heating element detected non-uniform winding of the heating element. It is believed that this variability in winding is common across solder tips and non-uniform winding could potentially generate hot spots during operation of the tip. Generation of hot spots could explain why experimental results using the IR camera were variable among tips and ultimately inconclusive.

An established preventative maintenance procedure can prevent such failures from impacting solder joint quality. At a minimum, the 48-hour replacement procedure should be adhered to. An alarm could be written into the robot software to inform the repairperson when the time specification has been reached. This would enable a consistent preventative maintenance schedule and could help to establish whether 48 hours is the appropriate time frame.

Additional issues that can be addressed to optimize the soldering process, improve solder quality and reduce cost are listed below:

- The visual inspection process does not adequately screen solder failures. However, the data collected on the log sheets should be used to support test system data. Operator log sheets should be modified to include more detailed information, such as the serial number of the failed unit. In this way, the log can be used to supplement the weekly analysis process of the technical team.
- Reduce tip dwell time, the time the tip is hot but not in use, to reduce variability across tips and extend tip life. As tip utilization is based on product mix, minimizing dwell time could eliminate another variable effecting tip life. This is especially critical at breaks, lunches and over weekends when the tip is warm and not in use for extended periods of time. This can be accomplished by incorporating a 'park' head command in the software program, ultimately eliminating operator intervention. A minimum time, 5 to 10 minutes, of idleness

should be used as a basis for the command. This should increase tip life by as much as 25%.

- Clean solder tips so they can be reused. Tips with less than 6500 solder joints that do not exhibit material wear may be candidates for re-use. All the tips collected during the internship could be candidates for an analysis to verify that a cleaning schedule can extend tip life.
- Leverage the software capability to implement additional automation. Three ideas suggested above were: addition of an alarm to indicate if the solder head does not travel to the expected coordinates, an alarm for solder joint count, and a 'park' head command to reduce dwell time.
- Wherever possible, utilize the on-line production system or software capabilities to automate and reduce operator intervention. As product mix continually varies, the on-line system is the most accurate method of tracking the solder robot performance.

Chapter 9 Results and Conclusions

The results of this study have been broken down into two categories based upon the study's primary objectives. They were originally defined in Chapter 1 but are recounted below for reference purposes as follows:

> **Priority 1**: Benchmark industry best practices amongst competition and equivalent companies to determine optimal quality control procedures with minimal cost impact
>
> **Priority 2**: Determine optimal test methods to achieve desired quality goals and decrease overall testing times by 50%

For the first priority a benchmarking study was conducted, the results from which are shown in detail in Chapter 2. The main learnings from this study are two-fold and are summarized below. Additionally, a production test system for The Company was analyzed in accord with the second priority. The results and conclusions from this analysis are listed in a later section.

9.1 Benchmarking Study Conclusions and Summary

The learnings from this benchmarking study are best aggregated into two ideas: HAST testing during the product design phase and better overall supplier interaction to support the goals of the entire enterprise. In the first case, three of the four companies investigated were actively using the HAST testing philosophy, or a derivative thereof, of incorporating an iterative phase of destructive testing to explore their product's predominant failure mechanisms. After each iteration, the mechanism is investigated to understand the underlying failure components and, employing a quality feedback loop system, the product and/or system design is reevaluated to correct the fundamental problems or issues. Such a system, if employed correctly appears to provide a more robust design and relegates the traditional burn in testing to a cursory run in test greatly reducing manufacturing cycle time and subsequently, product cost.

In the second case, several of the companies investigated were conducting extensive and bold supplier relations. To the extent that these companies were comfortable invading their supplier's manufacturing lines and even their product design groups, they were quite successful in building mutually serving relationships. These

relationships served both sides by improving incoming component quality to The Company's manufacturing lines and ensuring future business for the suppliers who showed a willingness to grow and adapt to process improvements. Additionally, by evaluating the capabilities of their supplier's processes the companies were able to successfully manage future product and process changes. Specifically, they were able to reduce incoming product variability reducing the need for more extensive in-line testing. Effectively, such testing was successfully removed from The Company's manufacturing and pushed upstream, where appropriate, to its suppliers.

By employing both more extensive design phase testing, such as HAST, and intensive supplier relations it is expected that The Company may be able to improve overall product quality and reduce existing in-line product testing. The Company is currently conducting rather lengthy product testing, specifically burn in type testing on the high volume products to verify component quality is satisfactory and ensure outgoing quality is high. HAST testing, being specifically designed to reduce failure mechanisms central to early life failure, has the capability to specifically reduce burn in type testing.

Additionally, intensive management of supplier relations throughout the product life cycle also has the possibility of reducing product testing. Such relations need to be mutually beneficial to both parties ensuring long-term relationship success and the gains that such a long-term relationship can garner. Currently, a large part of both functional and burn in testing appears to be utilized to verify supplier quality and design compliance. This system appears to draw extensive company and division assets away from the main goal of burn in testing, which is to precipitate and capture early life failures. By introducing certain measures of both HAST testing and more extensive supplier testing and relations The Company may be able, over the long term, to reduce the assets currently allocated to these endeavors.

Finally, Table 3 below is a compilation of the benchmarking study across some important parameters. As is evidenced by the results from this table there are many testing areas in which the industry has embarked in new directions, which may prove fruitful to The Company in the future. The two areas, which are highlighted above are, central to this theme.

PROGRAMS	COMPANY				
	A	B	C	D	Company
Suppliers					
Joint partnership to achieve process control	✔	✔	✔	✔	
PCB's					
In-house manufacturing		✔	✔	✔	
Contracts include quality specs	✔	✔	✔	✔	
Environmental Testing		✔	✔		
Joint partnership to achieve process control	✔	✔	✔	✔	
Product Design Phase					
100% Test	✔	✔	✔	✔	✔
Highly Accelerated Life Testing (HALT)		✔	✔		
Product Testing					
Hipot Test	✔		✔	✔	✔
Burn-In					✔
Highly Accelerated Stress Screening (HASS)		✔		✔	
Run-In Testing		✔	✔		
100% TEST	✔		✔		✔
Sampling		✔	✔	✔	

Table 3 - Benchmarking Results

9.2 Burn In Testing Reduction Results and Conclusions

At the time of project inception the overall burn in time the devices were subjected to was 120 minutes. The overall testing time including all four test phases was about 135 minutes. The stated goal of the project was to reduce the overall test time by 50% as stated in the second project priority statement. As Burn In testing accounted for a large majority of the overall testing time it had the greatest possibilities for overall test time reduction to achieve this goal. Therefore, the ambitious goal for Burn In was to achieve a testing time of 60 minutes.

After careful analysis, a significant failure mechanism associated with the 'Power On' command was discovered. In these cases there was a very significant spike event when the device was commanded 'ON' after a period of downtime. These periods were there to simulate a suspected real world operating condition and would allow device cooling to occur. For every such period followed by a power 'ON' command, there was subsequently a very large spike in the failure rates.

It was hypothesized that if these events could be moved closer to the beginning of the test it might be possible to reduce the Burn In test by 50%. While interesting, a concern regarding internal heating had to be addressed. In The Company's opinion, internal device temperature was a significant failure inducing mechanism; hence the use of the sauna. After reviewing internal temperature data it was determined that the longest thermal constant for any investigated component was approximately 50 minutes. Interestingly, this component was a critical power control board which had very high failure rates both internally and in the field. Such failure rates were due to its delicate nature. Based upon this internal temperature profile the goal of a 60 minute Burn In test appeared to be, at least, feasible.

To reduce the Burn In test time a new test sequence was proposed and implemented. The goal of the new test sequence was to force all of the power on events to well within the first 60 minutes while maintaining the existing ratio between On/Off events during the speed cycle. In addition, the new sequence would have an overall length of 90 minutes to correspond with the length of the Burn In test. Here it is important to note that the Burn In test length was shortened to 90 minutes from the initial 120 minutes at project initiation due to outside analysis. It was important to maintain the existing test time length to ensure that failures which were being caught in the last 30 minutes of testing might also be caught by the new sequence. The old (repeated here from Figure 4 for ease of reference) and new burn in cycles are shown in Figure 54 and Figure 55 respectively. Note, however, the time scale difference between the two (120 minutes vs. 90 minutes).

Figure 54 - Original 2 Hour Burn In Cycle

Figure 55 - New 1.5 Hour Burn In Test Cycle

As noted in Chapter 6 the cumulative failure curves should be pushed up and to the left requisite with the goal of faster failure inception rates. It was speculated in Chapter 6 that fully 90% of the existing failures attributed to the existing testing sequence could be realized within the first 60 minutes of testing as opposed to 90 minutes for the previous sequence. Additionally, if only Component and Assembly failures are included the expected captured failures at 60 minutes jumps to 95%. This is a very good measure of the existing system as these failures represent true device failures unencumbered by the overlying complexity of including all the Test System failures. Failures of these types are not good measures of the product robustness and actually cloud the data collection. As was seen in Chapter 6 such failures represent up to 77% of all failures recorded.

Using the new test sequence data was collected over several months and the following failure curve was generated from the new data. In Figure 56 the new test sequence has been overlaid on the failure rates for reference.

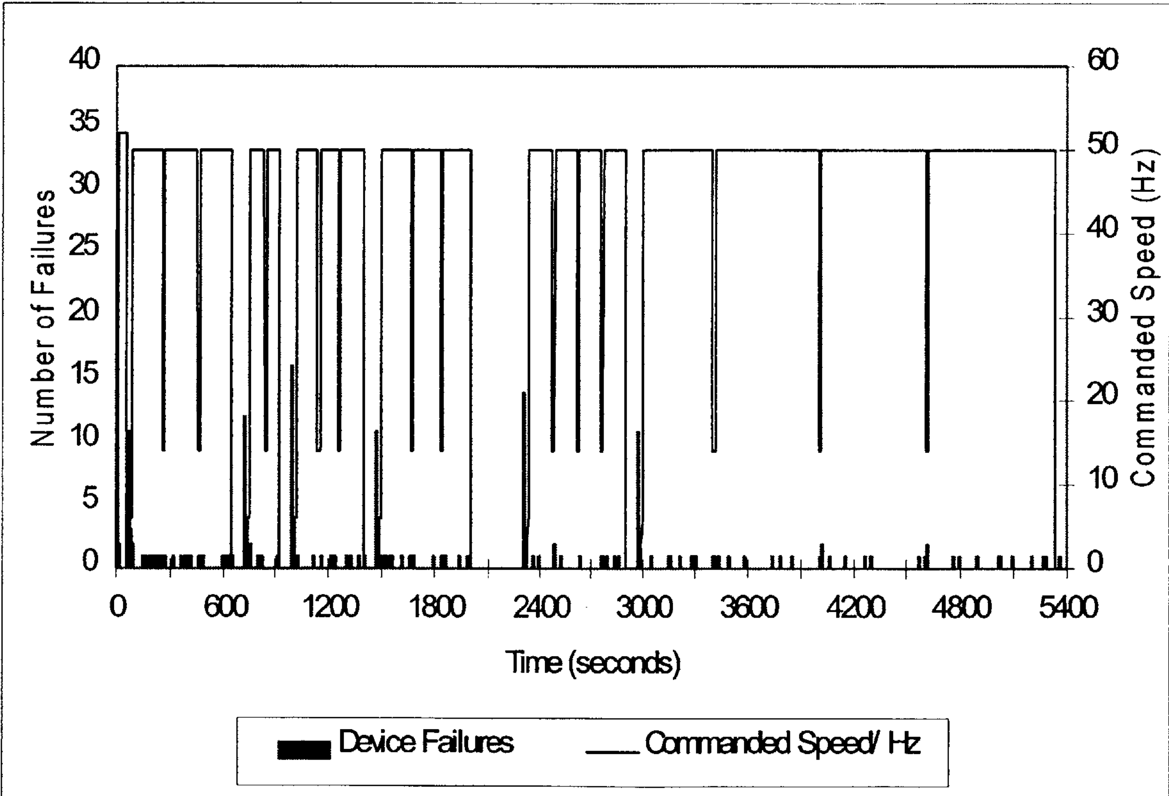

Figure 56 - Failure Rate for New Burn In Sequence

As can be seen in Figure 56 the spikes associated with the Power On condition followed the new sequence perfectly, strongly validating the hypothesized failure mechanism. In addition, the cumulative failure curves are shown in Figure 57

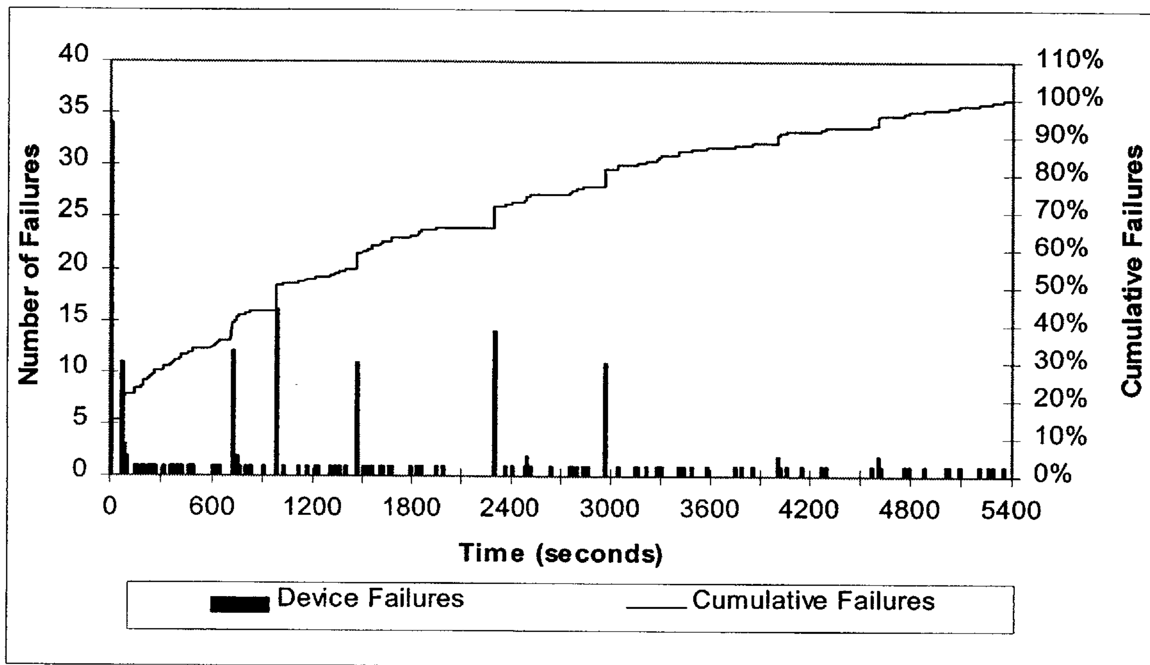

Figure 57 - New Burn In Test Cumulative Failures

At 60 minutes the failure captured from all failure types is 87%. This value is very near the failure capture rate of 89% that was expected from shifting the original data the requisite amount to match the new sequence. Such close correlation between the actual and expected values provides a high level of confidence in the relevance of the new test sequence and the data it generates.

Of course, the remaining 13% of the failures that occur after 60 minutes and prior to the end of the test at 90 minutes would not be captured and could represent escapements. This is the cost of reducing the production test sequence. However, it should be noted a similar cost was acceptable when the original test sequence was reduced from 120 minutes to 90 minutes during the middle of the investigation.

Additionally, all of the failures, which are represented here, would not necessarily show up as field failures. Such failures may never precipitate. Lastly, when 77% of the failures in aggregate are the result of the test system itself or not attributable to any known cause the remaining failures are small in number.

It is surmised from this analysis and the investigation of the test results that The Company may be able to reduce the overall Burn In test cycle to 60 minutes from its existing 90 minutes using the new test cycle proposed herein. This is only acceptable, however, if the company is willing to accept similar possible failure escapements into the field as was found acceptable in going from 120 minutes down to 90 minutes.

As field data becomes available during the next several years, The Company should evaluate it to determine if significantly more failures exist after the test change than before. As field data takes at least one year to obtain, this analysis is not feasible as part of this study.

9.3 Concluding Remarks

The 30 minute reduction of the Burn In test suggested by this study when combined with the time savings found in the two System Tests allows for an overall reduction from the 90 Burn In test of 33 minutes. The overall testing time would be reduced from 135 minutes down to approximately 70 minutes achieving a 48% test time reduction in accord with the stated project goal of 50%.

While this was an ambitious goal it does not come without costs as detailed in the preceding section. The benefits, however, could be far reaching for The Company. The shorter test time of course allows for greater throughput, as the test time is a significant percentage of cycle time. With the flexibility such a production system allows, The Company may be able to work more to a make to order system instead of a make to stock system as they currently do. This move could allow for drastic cuts in inventories increasing the savings. In addition, The Company may be able to cut back on its need for night shift work possibly creating better employee relations and the subsequent goodwill it engenders. Such analysis and implementations, however, must be the focus of future studies.

Bibliography

Eagar, T. W. 3.371 Welding and Joining Processes, Teaching Note, 1998.

Eppinger, S. D. and Ulrich, K. T. *Product Design and Development.* New York: McGraw-Hill, Inc., 1995.

Fine, C. H. *Clock Speed.* Reading: Perseus, 1998.

Harry, M. J. *Six Sigma: The Breakthrough Management Strategy Revolutionizing the World's Top Corporations.* New York: Currency, 2000.

Hopp, W. J. and Spearman, M. L. *Factory Physics: Foundations of Manufacturing Management.* Boston: Irwin, 1996.

Manko, H. *Solders and Soldering,* (2nd Edition). Reinhold: Van Nostrand, 1977.

Nahmias, S. *Production and Operations Analysis* (3rd Edition). Boston: Irwin, 1997.

Oster, S. M. *Modern Competitive Analysis* (2nd Edition). New York: Oxford University Press, 1994.

Shiba, S., Graham, A., and Walden, D. *A New American TQM: Four Practical Revolutions in Management.* Portland: Productivity Press, 1990.

Vining, G. G. *Statistical Methods for Engineers.* Boston: Duxbury Press, 1998.

End Notes

[1] Automation Research Corporation, 1997
[2] Grant and Leavenworth, "Statistical Quality Control", McGraw-Hill, New York, 1996.
[3] "Bed of Nails" refers to the test boards utilized. Such boards utilize many narrow probes which are positioned down onto specific circuit board points to verify conductivity within the circuitry.
[4] For confidentiality purposes wherever possible specific data relating to volumes or yield levels have been removed from charts.
[5] For confidentiality purposes letters represent phase descriptions.
[6] Unless otherwise noted the data used throughout this analysis was extracted from 1/1/99 – 7/19/99 production data.
[7] Mr. Don Galler of MIT's Department of Materials Science & Engineering performed SEM analysis, under the direction of Professor Tom W. Eagar.
[8] Metcal Product Information, www.metcal.com
[9] "M. F. Simington Soldering Ironing Work," MIT DMSE Don Galler, January 14, 2000.
[10] Note: this is regeneration of the solder tip from a cold start