Aligning Netbook Microprocessor Reliability to Market Demands

By

Anuj Bhardwaj

B.S. 2002 Mechanical Engineering
The Pennsylvania State University

Submitted to the MIT Sloan School of Management and the Department of Mechanical Engineering in Partial Fulfillment of the Requirements for the Degrees of

Master of Business Administration
AND
Master of Science in Mechanical Engineering

In conjunction with the Leaders for Global Operations Program at the

Massachusetts Institute of Technology

June 2010

Signature of Author ______________________
May 7, 2010
Mechanical Engineering, MIT Sloan School of Management

Certified by ______________________
Duane Boning, Thesis Supervisor
Professor and Associate Head of Electrical Engineering and Computer Science

Certified by ______________________
Roy Welsch, Thesis Supervisor
Professor of Statistics and Management Science, MIT Sloan School of Management

Certified by ______________________
David E. Hardt, Thesis Reader
Department of Mechanical Engineering

Accepted by ______________________
David E. Hardt, Professor and Chairman, Committee on Graduate Students
Department of Mechanical Engineering

Accepted by ______________________
Debbie Berechman, Executive Director of MBA Program
MIT Sloan School of Management

This page has intentionally been left blank.

Aligning Netbook Microprocessor Reliability to Market Demands

By

Anuj M. Bhardwaj

Submitted to the MIT Sloan School of Management and Department of Mechanical Engineering on May 7, 2010 in partial fulfillment of the Requirements for the Degrees of Master of Business Administration and Master of Science in Mechanical Engineering

Abstract

In 2008 Intel released a low-power and low-cost microprocessor that opened up a new market for smaller mobile computers, commonly known as netbooks. During the subsequent two years, netbooks have grown to be a substantial component of the mobile computer industry, capturing a major share of mobile computer sales in 2009. As Intel considers its growth opportunities, nothing attracts more attention than the processor for netbook computers. However, low-cost processors are a new undertaking for Intel, in contrast to an historically strong presence in the mainstream and high-end personal computer and server industries. In order to compete in the low-cost sector, the design of the microprocessor cannot include any unneeded capability. This is certainly true for aspects based heavily on customer usage characteristics, such as design for reliability.

The package-level CPU design includes solder joints which mount the CPU to the computer motherboard. Most of these solder joints serve the dual function of an electronic bridge between the silicon microprocessor and the computer as well as providing structural stability for the connection between the two. In recent years, Intel has added additional solder joints for certain products that are solely meant to provide additional stability. This is required to meet stringent reliability standards based upon estimated customer use patterns. In the case of the netbook processor, the usage patterns were originally assumed to be the same as those used for notebooks, influencing package-level design and requiring additional solder joints. However, after being in the consumer market for only a short period, evidence points towards different usage patterns for netbooks, which in turn affect the reliability targets and overall package-level design.

Through recent studies and analysis, it has become evident that consumer usage of netbooks is indeed different from notebooks and that Intel's existing set of assumptions may need to be revised. Specifically, the device lifetime, percentage of active time, power cycle frequency, and application use vary significantly between notebook and netbook devices. By leveraging Intel's extensive experience in relating usage patterns to reliability targets, it is possible to estimate the effect that such differences may have and conclude that additional solder joints are not required. These analyses estimate that by using more realistic usage assumptions, it is possible to avoid an additional $24 million in production costs over a two-year period for the upcoming microprocessor design to be released in 2011.

Thesis Supervisor:	Duane Boning
Title:	Professor of Electrical Engineering and Computer Science
Thesis Supervisor:	Roy Welsch
Title:	Professor of Statistics and Management Science

This page has intentionally been left blank.

Table of Contents

List of Figures

List of Tables

Acknowledgements

This project would not have been possible without the assistance of many people. First, I would like to thank the Leaders for Global Operations program at the Massachusetts Institute of Technology for giving me the opportunity to experience and learn a tremendous amount about business and operations. My advisors, Duane Boning and Roy Welsch, have been especially helpful to me during my internship. My classmates in this program were a major asset during this project and they continue to guide me as I attempt new challenges.

I would also like to thank Intel for sponsoring me as an LGO intern. The Mobile Platforms Group, Assembly Technology Development, and the Quality and Reliability organizations were exceptionally supportive and provided me with the tools needed to complete my analysis. I am especially grateful to my supervisor, Jason Ku, who gave me guidance and direction throughout my internship. Cindy Ng, Shubhada Sahasrabudhe, and Andreh Janian were particularly helpful and I certainly could not have attempted this project without their experience and insight.

My family has always been there for me and I have been able to join and complete this program with their unyielding support. I am especially grateful to my wife Avni who has been my advisor and friend throughout.

1 Introduction

Intel Corporation is the world's largest manufacturer of microprocessors for personal computers. The company has embarked on numerous new growth opportunities in recent years, which include opening up new markets for embedding microprocessors into commercial and industrial equipment as well as increasing the computing capability of small and handheld computing products. This study will focus on the design for reliability required for the recent smaller personal computers that have been made possible by a microprocessor product innovation by Intel.

1.1 Intel and the Netbook

Intel designs its microprocessor products to meet the requirements of both end-users and original equipment manufacturers. In order to ensure proper functionality and longevity, all microprocessors must meet a certain level of reliability which considers how end-users will employ their devices. Much of this analysis is based upon historical data gathered over a substantial period. There is significant data that can be applied to understand customer use for Intel's core products: desktop, notebook, and server computers, all of which have been in the market for at least 15 years. This has allowed Intel to make important decisions to optimize the design of its microprocessors for adequate reliability.

In early 2008 Intel released a low-cost and low-power processor which spawned a new set of devices known as netbook computers. Similar to notebook computers, netbooks are ultra-portable, characterized by a small form factor which is measured by a display with a diagonal of 8 to 10 inches. With the low-power processor, netbooks can be used for simple functions which allow them to boot up quickly but limit their capability for heavy computing applications. This puts them in a separate category of devices with functionality above a PDA (personal digital assistant) or smartphone but below a full-size notebook or desktop computer.

With the launch of this new processor, netbook sales have grown dramatically, and are currently estimated to represent 20% of the mobile PC market's annual sales.* This has placed enormous cost pressure upon the mobile industry. The price point for the Intel netbook processor is a fraction of the price of microprocessors used for notebooks. Intel must now put additional effort on cost reduction of manufacturing this product. For this reason, netbooks must achieve significantly higher volumes in order to produce revenues comparable to the notebook business.

1.2 Assumptions about Netbook Usage

Reliability for microprocessors is heavily dependent upon the assumptions made about customers' usage of their computer devices. For example, if it is assumed that customers will keep their device for 10 years as opposed to 5 years, an additional amount of reliability must be provided to ensure adequate performance throughout the remaining 5 years. This one factor alone shows that customer use estimates are extremely important for the purpose of determining reliability design.

As previously stated, the markets for netbooks and notebooks differ distinctly from each other. Since the market for netbooks was initially unknown, the use condition assumptions for netbooks were conjectured to be the same as those used for notebooks. Shortly after the release of netbooks, research showed that people were using the devices in a manner significantly different from the supposition. Specifically, the design for netbook reliability appears to be higher than what might be acceptable. This is not surprising since notebooks are used consistently for corporate purposes and for heavier computing applications, where netbooks are

* Mobile PC market includes notebook and netbook products.

not. For Intel, this means that unnecessary money is being spent to design and build reliability for their netbook products.

1.3 Cost of Reliability

There are multiple aspects of the microprocessor design, including the silicon die, that are affected by reliability specifications. This study focuses on the assembly-level design, specifically looking at the backside of the package, which is the area mounted to a computer's motherboard. The actual connection between the microprocessor and the motherboard is made by an array of solder balls (or joints) which adhere the two units together. Throughout the lifetime of the device, the solder joints can crack due to thermal expansion and the performance of the microprocessor may be compromised if signals are transmitted through a broken solder joint. Solder joints which transmit electricity or information are known as critical to function (CTF). Thermal-mechanical reliability is defined by the amount of thermal expansion or contraction that can be withstood by the solder joints before the CTF joints are compromised. This effect is different from the reliability required for the silicon die and its transistors. A separate set of criteria and design is in place to deal with the reliability of the die, however, this study will focus on the thermal-mechanical reliability of the solder joints.

One method of avoiding the failure of CTF solder joints is to add additional joints which are non-critical to function (NCTF) onto the backside of the CPU package. These joints can be strategically placed in order to ensure that they fail before a CTF joint, allowing the microprocessor to perform adequately. On the other hand, adding NCTF solder joints may increase the necessary area required on the package backside, thus increasing the overall size of the package.

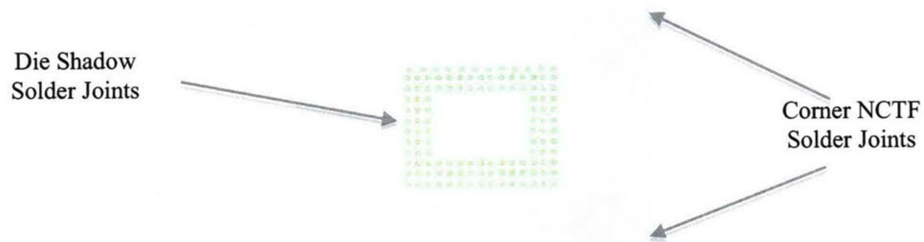

Figure 1: Drawing of Microprocessor Package Backside.

As can be seen in Figure 1 above, the backside of the microprocessor package consists of solder joints in two key areas. The center region is known as the die shadow, because it resides directly underneath the silicon die. The corner regions may contain solder joints if additional stability is required. In this case, NCTF joints may be added, possibly increasing the package size. An increased package size has a direct impact upon the cost per part since it requires additional substrate. As we will see, this is a significant cost which can be avoided using adequate assumptions regarding reliability.

1.4 Advantages of Superior Reliability

We have reviewed the direct cost of adding increased reliability which comes in the form of package* cost. It is also necessary to discuss the strategic decision of increasing or decreasing reliability targets. It has been Intel's historical policy to provide a product with reliability superior to that of its competitors'. Over time, this has led Intel's customers to expect certain reliability standards for its notebook products which have not changed for many years. Some

* The microprocessor package is the printed circuit board that provides the circuitry to transfer processor information from the silicon die to the motherboard.

believe that this has created a situation where changing those specifications can alert an OEM customer to think that a major design change has taken place which reduces reliability of the product. For the purposes of this study, we will look at the design-based assumptions and use data to justify any recommendations. As we will see, none of the recommended changes should come as a surprise to OEMs since some of them have already adopted different assumptions for netbook and notebook computers.

It is also worth noting that few (if any) of Intel microprocessors actually fail during use. It is far more common for other components of a netbook or notebook system--such as the display or keyboard--to fail well before the microprocessor. This allows Intel to maintain its reputation of high quality and reliability. The intangible value of that reputation and brand is difficult to quantify, but we can assume that it is far greater than the savings associated with changing reliability targets. However, it is the hypothesis of this study that with respect to netbooks, Intel is building reliability well beyond that which is required to protect its brand. It is also important to better understand the market for netbooks. For these reasons, we will look deeper into the manner in which netbooks are used and try to determine an optimal set of reliability targets that balance high performance with reasonable expectations.

1.5 Thesis Overview

During this study, we will discuss the background of Intel and its new microprocessor product for the mobile industry, going into further detail behind the origins of the product and its design for reliability. We will also review in detail the conditions that affect reliability in order to understand the factors that influence reliability targets and subsequently the design. Using the tools that Intel has developed over a significant period of time, we can then estimate the reliability targets required to ensure that the microprocessor design is adequate for meeting the given set of use conditions.

Finally, we will estimate the design changes that would reflect a more customer-aligned set of use conditions. Calculating the financial costs or benefits of making such a change to use condition assumptions will help to understand the strategic impact of making such a decision. We will also discuss the challenges of making such changes as well as other factors that impact Intel and its OEM microprocessor customers.

2 Background

Intel has one of the strongest brands in the world and is well known for its microprocessors which are used in most personal computers. To attain its current level of success, Intel has had to optimize its product design and set up an organizational structure that stays close to its customers, while managing cost and providing adequate reliability. Because the microprocessor functions as the brain of any computer, its failure is unacceptable. For that reason Intel takes a conservative approach by building in enough reliability to ensure that failure is either non-existent or very rare.

2.1 Intel Corporation

With annual revenue of $35 billion,[1] Intel is the most successful semiconductor manufacturing company in the world, their core products being the central processing units for computing devices. Although their overall market share dropped slightly towards the end of 2009, they currently have 71.1% of the desktop market, 87.3% of the mobile market, and 89.8% of the server market.[2] Their dominance in the manufacture of microprocessors is largely created by the ability to innovate manufacturing methods that produce chips with smaller transistors, in turn enhancing processing power without increasing the size of the unit. They are able to stay ahead of their competition by consistently producing new manufacturing technology on a two-year time frame by their coined tick-tock method, where every other year generates either a new platform design or utilizes a smaller transistor. This has created a microprocessor market where many of Intel's competitors offer lower performance products with older transistor technology at a discounted price. For this reason, Intel has been able to charge a significant premium for their microprocessors and to extract a large portion from the computer product value chain.

2.2 Organizational Roles with Respect to Netbook microprocessors

Intel has successfully handled the responsibilities associated with leading the microprocessor industry by carefully balancing the duties of different internal organizations. The Mobile Platforms Group (MPG) is comprised of engineering, sales, and marketing groups. The thermal group in the Client Platform Engineering department of MPG is responsible for understanding some end-user operating conditions for netbook products which may impact design. For example, determining the application run by an end-user is the job of the thermal group while the Quality and Reliability Group (Q&R) determines other parameters such as the daily time of use for that application. MPG oversees the business of netbook products, including marketing, business strategy, and the development of customer relationships. Although in many ways MPG owns the microprocessor products for netbook devices, it must adhere to requirements set by Q&R, while relying on the Assembly Technology Development group (ATD) to help develop its package-level product design.

As mentioned above, the Q&R organization has a corporate responsibility to ensure that all products meet necessary requirements. There is a significant amount of effort needed to adequately complete the task of determining these requirements. The business decision to provide a certain level of reliability for a product is based upon finding a balance between reducing costs and minimizing the number of failures. Intel has a reputation for providing high-reliability products which requires allowing very few failures in the end-user market. However, the estimation of how many failures may occur based upon certain reliability targets depends upon the accuracy of the assumed use conditions. As mentioned previously, for netbook processors the use conditions have been assumed to be the same as for notebook products. We believe that this assumption is not valid and that Intel should adjust use condition assumptions to be adequate but not inappropriately costly.

The overall package-level design is created by ATD and constructed to meet the reliability targets set by Q&R within certain design specifications determined by MPG. ATD serves Intel as the group that optimizes the assembly and testing processes required for high-volume manufacturing. Through their main development facility at the Chandler, AZ, site, they determine manufacturing processes that are transferred to high-volume centers in Malaysia and Costa Rica. ATD is a valuable organization with the broad responsibility of both determining design parameters and creating manufacturing processes.

2.3 Notebook Products and Reliability Design

Notebooks computers provided a tremendous amount of growth for Intel in the late 1990s and early 2000s, mostly due to their superior performance which pushed the limits of computing power. Intel's notebook microprocessors have had historically high prices and they have been very profitable at the level of margin normally attained for Intel's performance premium. Intel has been producing notebook microprocessors for over 15 years, so there is a wealth of data that has been collected to understand customer usage for design and marketing considerations. This information has also been used to set targets for reliability. As mentioned above, Intel has chosen to take the standpoint that it is best to stay on the high end of reliability in order to build and protect its strong reputation. For that reason, notebook microprocessors have had a strong history of low failure rates, especially in the thermal-mechanical mode.

3 Netbook Microprocessor

The main assumption that drives the reliability design for microprocessors is the set of conditions assumed for consumer use. The reliability of a microprocessor could vary significantly based on how an end-user utilizes a computer. In order to provide high reliability for all end-users, Intel must make conservative assumptions regarding usage and study closely all of the different use conditions that affect reliability. We will look further into the current set of use condition assumptions and consider how more realistic assumptions may be different.

3.1 Origins of Product

Historically prices for notebook products have been relatively high, allowing for a significant margin to be extracted for the microprocessor ingredient product. However, in recent years, the growth for notebooks has been stagnant, which some attribute to what appears to be a nearly saturated market. Recognizing this, Intel made a strategic decision to produce processors targeted towards the growing market of smaller mobile devices, such as PDAs, smartphones, and netbooks. The microprocessor created for this market is distinctly different from that used for the mainstream mobile notebook market. It consumes less power and has limited computing capability but takes advantage of Intel's newest wafer fabrication technology. The cost is below mainstream mobile devices and thus it is offered at a significantly lower price.

This product was officially released in early 2008 and quickly witnessed very fast growth. Many originally anticipated that the largest portion of sales for the new microprocessor would be in PDAs and smartphones. Netbooks were a new product, targeted at a market of consumers who either preferred a smaller form factor or were only willing to pay a lower price for a computing device. It was also expected that the initial major market for netbooks would be in emerging economies, where the willingness to pay would be lower than the prices offered for notebook computers. It was projected that families who could not afford laptops at higher prices would purchase netbooks as their primary computing device.

3.2 Present Market

As we now know, netbook sales growth has been highest in countries with established economies where consumers are using them as secondary devices.[3] From Figure 2, we can see that the regions of North America and Western Europe together currently comprise the majority of netbook sales, while Asia-Pacific has the highest total sales of any individual region. Based on this data alone, it is apparent that established economies represent the majority of sales.

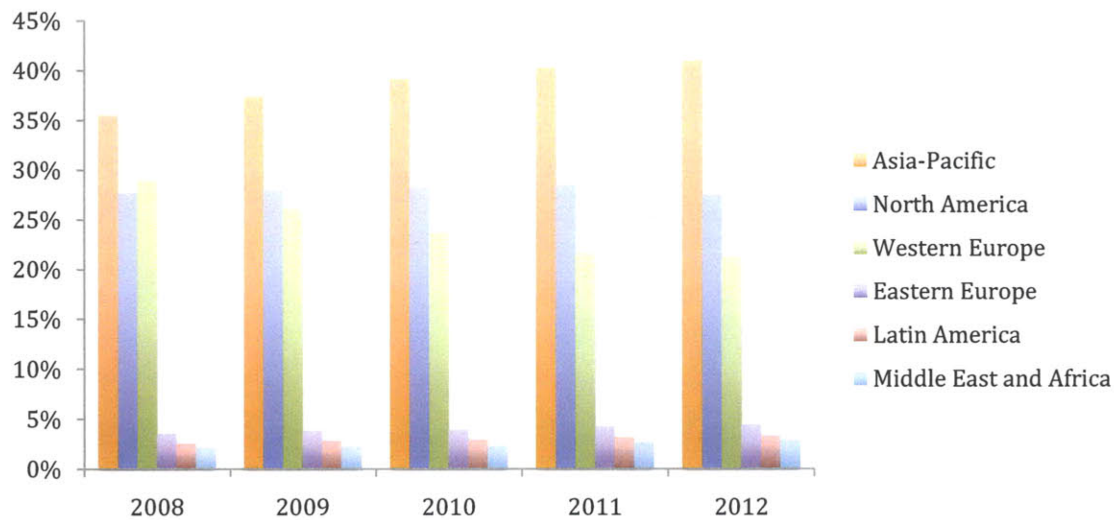

Figure 2: Worldwide Netbook Unit Sales by Region.[3]

As mentioned above, netbook sales are strongest as secondary computing devices. They are almost solely used for personal purposes, which is significantly different from the notebook market wherein a sizeable portion of devices is for corporate use. Current data show that only 3% of netbook sales are utilized for corporate usage. As we will see, these characteristics have a strong bearing upon the use conditions affecting reliability.

3.2.1 Use Conditions Affecting Reliability

Table 1 lists the use conditions that are employed to determine thermal-mechanical reliability and states the currently assumed values for netbooks.

Use Condition	Currently Assumed Value	Percentile User*
Lifetime	5 years	90^{th}
Ratio of Active Time to Idle Time	23 %	90^{th}
Power Cycle Frequency	6 cycles/day	90^{th}
Application Used	High Intensity†	50^{th}

Table 1: Current Customer Use Condition Assumptions‡

Looking further, we will examine each of these use conditions to determine their appropriate values for the netbook market.

3.2.1.1 Lifetime

While netbook products have only been on the market for approximately one year, it is possible to make some conclusions regarding the average total lifetime for each device. We will attempt to estimate the lifetime for netbook devices based upon a comparison with notebook computers while analyzing the differences between the consumer use patterns for the two devices.

* Represents the percentile of users that assumption is meant to capture. For example, the lifetime assumption is that 90% of netbooks will be used for 5 years or less.

†Assumes that a high intensity application is running when the device is active. This is one of three components used to calculate the temperature of the solder joint, an input in the reliability model.

‡ Data is based upon assumptions made by Intel regarding customer usage patterns.

Intel's 90th percentile lifetime estimate for netbook microprocessors is currently 5 years. This allows for numerous considerations, and assumes that 90% of devices will not be in use after 5 years. The actual average lifetime for computers varies from year to year based upon several factors, such as the state of the economy. It was predicted that the replacement time of notebook computers will begin decreasing after 2009, due to forecasts that the global economy would slowly strengthen following the recession that began in 2008.[4] The economic downturn caused many users to delay upgrading their devices in order to save money, but while markets across the globe begin to strengthen, the lifetime of computers should decrease again as individuals accumulate greater buying power. There are also studies that show significant savings to companies that replace their corporate notebooks earlier and thus avoid costly repair and servicing.[5] This is a trend that more firms may be adopting. Altogether, the average replacement time is not expected to fluctuate by any significant value with respect to the use condition assumptions utilized in this study.

Figure 3 below is a distribution of the percentage of devices in use for a given period of time.[6] The data shows that over 90% of devices will be replaced in 5 years or less.

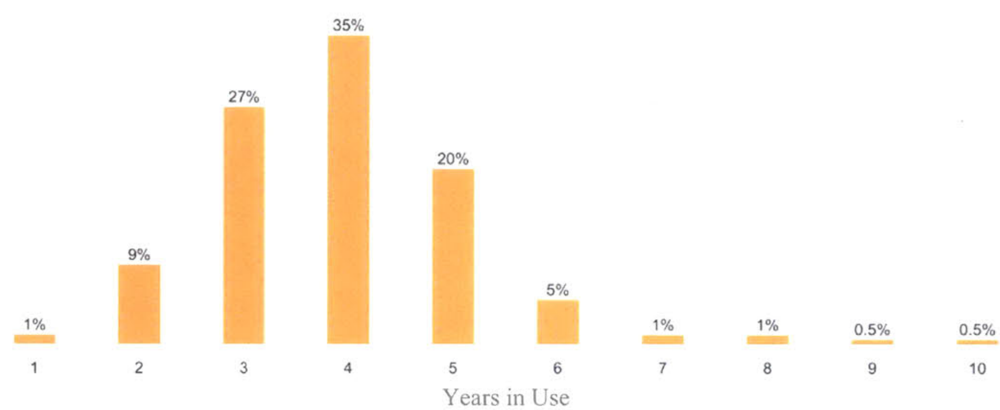

Figure 3: Percentage of Notebooks Over Years from Initial Purchase.*

* Data in figure is based upon Intel's estimates for average lifetime of notebook computers.

It is worth noting that a secondary user may operate a small proportion of a computer's lifetime after the primary user has given up ownership. For example, if a device is donated after the primary lifetime, it may be used for a secondary period. The current average secondary lifetime is estimated to be 30% of the primary life,* extending the average lifetime for these devices. This data is captured in the above distribution. Although this extended lifetime has been noted, it will not be considered in the analysis for netbook computers since most users do not expect longevity during the second life of a device and very few netbooks are expected to be in use beyond their primary life.

One of the major driving factors behind notebook replacement time is the initial price of the device. Mainstream notebook computers generally cost between $800 and $1700 for normal consumer models. Due to this relatively high price point, consumers are less likely to purchase a new device until the hardware technology has become obsolete, the device will no longer support software upgrades, or components begin to fail. On occasion, consumers will choose to repair or replace individual components if the cost is lower than the purchase of a new device. This prolongs the length of time for keeping a notebook. In many cases, the cost of purchasing new software is also much less than the cost of a new device, in which case consumers will upgrade to new software instead of buying a new device with pre-installed software.

The initial price of netbooks is much lower, most retailing from between $200 to $500 dollars, a fraction of the cost of notebook computers. Consumers will likely choose to buy a new device instead of repairing the much cheaper netbook in the event that a component fails. It may also make sense to purchase a new device with preinstalled software, especially an operating system. Perhaps one of the major reasons that a netbook computer is expected to have a shorter lifetime is suggested by the inherent manner in which the device is being used. In this study we will see that netbooks are used in a manner similar to smart phones or PDAs. As a result, they

* 30% value based upon forecasts for 2009 usage.

will increasingly be stored in small bags or purses and incur more physical damage than their notebook counterparts. This may cause the more fragile components, such as the LCD display, to fail earlier. At an average replacement price of over $100, consumers are less likely to replace a display, marking the end of a netbook's life. This is commonly seen in smart phones and PDAs, both of which have an average lifetime of less than two years.

Another major factor in the lifespan of a device is the demographic profile of its market. Netbook sales are currently stronger in established economies where the devices are used to accompany another notebook or desktop computer. Increased usage in emerging economies may produce different effects. One argument suggests that if a device is used as a primary computer for an individual or a household, the computer will be kept for a longer period in an economy where people have lower disposable income. This tends to be the case for notebooks since the initial price is much higher and it may be more economical to replace failed components rather than the entire device. Figure 4 depicts the relative predicted average lifetime for notebook computers in several emerging and developed economies. The chart also shows a breakdown depicting the difference between enterprise (corporate) and personal use.

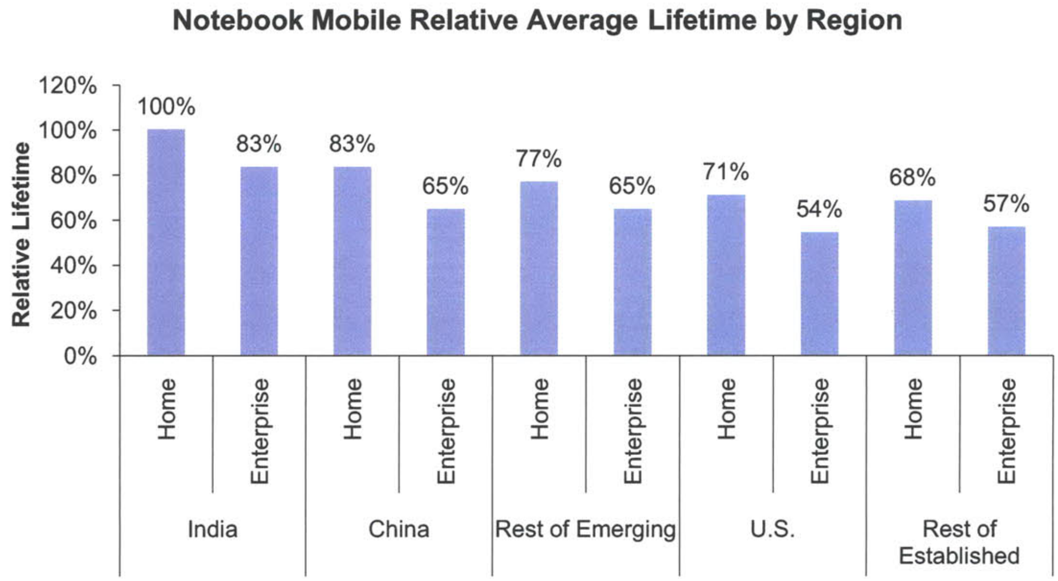

Figure 4: Comparison of Lifetime by Region.*

As can be seen from this data, emerging economies--especially India and China--have longer replacement periods than established economies. For reasons mentioned above, this is logical since disposable income is lower in these countries and people try to avoid spending the large upfront cost to replace a computer. If the market increases in emerging economies for netbooks as primary devices, then this same trend should be expected with netbooks. However, use of netbooks as primary devices in emerging economies has not occurred as originally anticipated.

It was initially predicted that the fastest growing markets for netbook computers would be in emerging economies where many households--previously unable to afford a more expensive computer--would be able to purchase a netbook as their primary device. So far established markets have been responsible for the majority of sales, and there are two potential reasons. First, netbooks proved to be attractive as secondary devices in established markets where

* Comparison estimates from Intel Market Sizing and Forecasting department. Data is normalized against India home market, which is equal to 100%.

disposable income is higher. Secondly, many households in emerging markets have opted to spend slightly more for larger inexpensive notebooks, often powered by low-end microprocessors such as the Intel® Celeron™ processor. This second point is intriguing as it indicates that many users in emerging markets will ultimately choose not to purchase netbook computers as primary devices, as had been previously thought. This has a major impact on assumptions such as the average device lifetime. It could subsequently be argued that even emerging markets will not use a netbook for the same period that they would a primary computer such as a large notebook.

Another driving factor behind the average lifetime for netbook computers is that telecom companies--such as Sprint, Verizon, and AT&T--are offering netbooks at no up-front cost with two-year subscriptions to wireless data service packages. For consumers who purchase and maintain these subscriptions, almost all will limit the primary lifetime of their notebooks to two years since they would be offered a new one with continuation of service. It is estimated that mobile operator-issued netbooks will reach sales of over 60 million units by 2014. The same forecast estimates that 60 million unit sales in 2014 will be approximately 40% of the netbook market. This is a substantial volume that could play a major role in setting the standard for how long a netbook computer is kept and how often the technology is refreshed since new models are released to meet continuing service provider offers.

These factors provide a compelling case for netbook devices to have significantly shorter lifetimes than their notebook counterparts. In order to predict a reasonable average lifetime for netbook computers, it is necessary to take into consideration the factors mentioned above: low price point reducing the number of repaired devices, low adoption as primary device in established as well as emerging markets, and two-year offers with telecom service contracts. It can be estimated that the average netbook lifetime may be greater than two years, but less than three. For the purposes of this study, the average life span will be estimated at 2.5 years, with a distribution function equivalent to that used for notebook devices. Intel Quality and Reliability uses a cumulative normal distribution, requiring an average value and standard deviation, to plot the lifetime for notebooks. In the case of notebooks, the average lifetime value found by the Q&R group at Intel is approximately 3.75 years. Since our estimate for average lifetime is 33% lower for netbooks than for notebooks, an equivalent difference will be chosen for the 99.9th percentile value for notebooks, currently estimated to be 7 years. This makes the 99.9th percentile value for netbooks approximately equal to 5 years.

Assuming that the 99.9th percentile user keeps a notebook for 5 years and the mean user for 2.5 years, we can distribute the lifetime of netbooks as shown below in Figure 5.

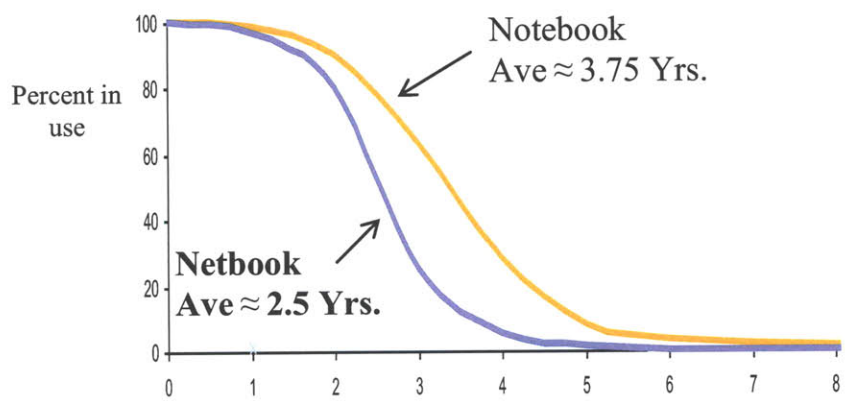

Figure 5: Cumulative Normal Distribution for Netbook and Notebook Lifetime.*

Observing the same distribution for notebook products we can see that approximately 8% of devices are used beyond 5 years. This is near the target of 10%, which Intel chooses to incorporate into the reliability targets. For netbook computers, the 10% distribution mark is near 3.5 years. It is this value of 3.5 years that we recommend for trying to estimate the 90th percentile value for netbook lifetime.

* The trend of this figure is based upon what is used by Intel's Quality and Reliability department, but the data shown here is for conceptual purposes. The netbook lifetime trend is purely speculative.

3.2.1.2 Warranty Information and the Implications for Device Lifetime

Another way to understand the expected lifetime for netbooks is to gain information from OEM manufacturers who design, build, and sell them. Unfortunately, these companies do not publish their expectations. However, they do provide warranty information which we can compare directly to warranties they offer for notebook computers. The fact that we understand well the lifetime statistics and expectations for notebooks allows us to compare the warranty information and infer OEM expectations for netbooks.

OEM	**Standard**	**Extended (total life)**	**Extended Costs**	**% Original Cost**
Dell	1 year	3 years	$179	~20%
Lenovo	1 year	3 years	$269 or $179	~25%
HP	1 year	4 years	$399.99 or $319.99	~40%
Acer	1 year	3 years	$200.00	~25%
Average	**1 year**	**3.25 years**	**$240**	**~30%**

Table 2: Notebook OEM Warranty Information.

OEM	**Standard**	**Extended (total life)**	**Extended Costs**	**% Original Cost**
Dell	1 year	3 years	$99	~35%
Lenovo	1 year	3 years	$200	~60%
HP	1 year	3 years	$179.99	~55%
Asus	1 year	3 years	$145	~50%
Acer	1 year	2 or 3 years	$129 (2 years)	~40%
Average	**1 year**	**3 years**	**$160**	**~55%**

Table 3: Netbook OEM Warranty Information.

Tables 2 and 3 break down the warranty periods and costs for both netbook and notebook products.* While the cost to extend the warranty for netbooks beyond the first year are comparable to notebooks, the proportion of that expense relative to the original device cost is 55%, almost twice as high as for netbooks. From this information, it can be inferred that the OEMs either expect a smaller proportion of people to extend the warranties for their netbooks or that the OEMs prefer to extract a larger sum of money for the netbook extended warranty since they expect these devices to fail earlier.

In either of the cases mentioned above, the warranty information suggests that OEMs expect netbook devices to have a significantly lower lifetime on average than notebooks. This supports the findings from the customer use studies that we reviewed earlier, and may play a role in the dynamics affecting the netbook device life span.

3.2.1.3 Device Active Time

The number of hours that a device is in use during a day traditionally incorporates active time as well as idle time. In the past, notebook devices targeted for a 7-year life were assumed to be in use (active or idle) for approximately 17 hours per day. Recently, the 90^{th} percentile target for notebook computer lifetime was reduced to 5 years, but the amount of time that the system is in use is now assumed to be 24 hours per day, leaving the total hours in use to be the same value as previously used for the 7 year lifetime design or 43,800 hours.

While it can be argued that computers are turned off for at least portion of their lifetime, the difference between the idle state and the off state is negligible for the purpose of determining necessary solder-joint reliability. This is because the temperature of the microprocessor during the idle or hibernating state is very low. For this reason, it will be acceptable to adopt a simplified

---

* Both Tables 2 and 3 are based upon information gathered from public sources, including the websites of the OEM companies that create netbook and notebook computers.

version by considering the device as in use for 24 hours per day, while assuming that non-active time is equivalent to being in the standby mode.

As we did previously, it is necessary to consider the case of notebook computers for comparison when attempting to determine the amount of time that a netbook device is active. The current assumption is that a notebook is active 21% of the time that it is in use, the same value assigned to netbooks. This is equivalent to approximately 5 hours per day of active time and is meant to represent a 90th percentile user.

For notebooks, corporate use of a large proportion of devices is a major factor in driving up the percentage of active time. Most netbook use will be personal or at home, not in a corporate environment. This is supported by a recent study which recorded that 63% of its participants reduced the use of their personal desktop computers when they began using a netbook device.[7] The same study also collected data showing that on average netbook devices were in use for a period of slightly less than 2 hours per day. Approximately 50% of the 2 hour per day average was recorded to be active while the other 50% was in a ready-to-use condition and not in stand-by. For estimating reliability, we will look only at the portion of time when the device is actively being used. One reason for a large proportion of the time in an on, but non-active, state is that the users of netbooks in this study were often found to be preoccupied simultaneously with other tasks such as watching television.

While this study recorded an average active time of approximately 1 hour per day, it is necessary to consider a majority of the user population. For this purpose, it will be necessary to count the equivalent percentage of users included in the analysis performed for notebook computers. As mentioned before, the assumption is 6 hours per day of active usage for notebook devices. This estimate attempts to represent 90% of the population of notebook users. If we consider the error associated with performing only a single study with a limited sample size, we should state that actual results may vary from these findings. Therefore, with our currently limited data, we will take a conservative standpoint and estimate that the average netbook usage could be assumed to be around 2 hours per day, with the 90th percentile value being 4 hours per day, translating to 17% active.

We believe that 17% active time is a conservative estimate of the 90th percentile user, however, with a limited amount of information we do not want to recommend a lower value until further data is gathered. Active time is something that can currently be accurately measured, unlike lifetime. Intel has conducted only a single study to understand a variety of use conditions and it is highly recommended that further effort be made to understand how long netbooks are being used actively.

3.2.2 Power Cycle Frequency

A power cycle can occur each time a device is cycled from the on state to an off, hibernating, or a standby state. However, in order for it to have an effect upon solder-joint reliability, the microprocessor must have enough time to cool or heat up to a near-steady state temperature, which requires approximately 20 minutes. For this reason, a change in state must occur in the manner mentioned above and the new state must be held for at least 20 minutes. For instance, restarting a computer or shutting down and immediately rebooting do not constitute power cycles. Power cycle frequency is a significant contribution to solder-joint reliability since it causes thermal cyclic loading on the microprocessor. This creates stress on the solder joints between the package and the motherboard due to thermal expansion of the microprocessor package.

Market research has indicated that netbook devices are cycled more often between power states than notebook devices. Some of these studies have shown that people are more likely to turn their device off even if they think there is a possibility it might be used again shortly. This behavior is similar to that used with mobile phones. In contrast, users of notebook computers will often decide to leave the device on if they anticipate that it may be used again before long.[8] This is intuitive when one considers that netbooks are a medium between PDAs and notebooks and that PDAs are turned on and off at a much higher frequency than most computer devices.

The current supposition is that the power cycle frequency for netbooks is 6 times per day. Like the other use condition assumptions, this estimate was adopted from the notebook value. Further analyzing actual power cycle frequency, evidence suggested that the 6 cycle per day assumption is likely conservative when predicting the 90th percentile notebook user and perhaps more appropriate for the 90th percentile netbook user.

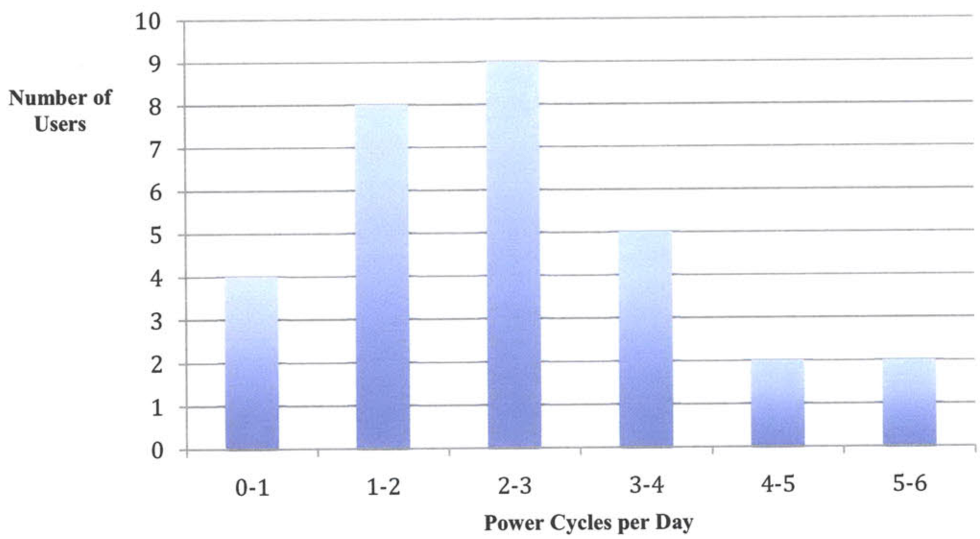

Figure 6: Power Cycle Frequency Histogram for Notebook Usage.

One recent study measured that the 90th percentile notebook users cycled their device 2.65 times per day.[9] The data represented from Figure 6 depicts a histogram from this study. It is evident that the majority of users cycle between 2 and 3 times per day, with only a small group residing above 4 cycles per day. The 2.65 cycle per day average for notebooks actually measures the number of times that a device is changed from on to an off, hibernate, or standby state without evaluating the amount of time that the device is in the changed state. Therefore, the 2.65 cycle per day measurement is higher than the actual power cycle frequency, as previously defined.

The netbook study was used to measure the amount of active time by netbook users and found that the same devices were cycled between the on, standby, hibernate, or off state on average 2 times per day.[10] Again, this does not consider the number of times that a change in state occurred with at least 20 minutes for the temperature of the microprocessor to settle. However, the study did contain biases, as the number of users of the devices was limited to 60 and their handling of the devices may have been influenced by the fact that they knew they would not be keeping them long-term. Nonetheless, the 90th percentile power cycling value for netbooks has not been measured to be as high as the currently assumed 6 cycles per day.

Based upon this information, we recommend that the current power cycle frequency assumption of 6 cycles per day not be increased. At this point in time, 6 cycles per day is an adequate (possibly conservative) estimate even for netbooks. As we will see later, power cycle frequency has a major impact upon solder-joint reliability and for that reason we highly recommend that further data be gathered to better understand this use condition parameter.

3.2.3 Solder Ball Temperature and Application Use Assumption

In order to understand how application use affects solder-joint reliability, it is necessary to discuss the parameters that affect the solder ball temperature. The device solder ball temperature is estimated by the Client Platform Engineering organization within Intel's Mobile Platforms Group. The method of determining the solder ball temperature is dependent upon the following factors:[11*]

* CFD analysis is performed based upon this information as the main input. The solder ball temperature calculation is the final output from this analysis.

1. Estimation of the thermal solution used by an OEM. The thermal solution evaluates the kind of cooling systems which may be in place within the netbook. A 3D computer aided design (CAD) model is used for simulation, but is simplified in order to allow for efficient analysis.
2. Boundary conditions assumptions which are input for a computational fluid dynamics (CFD) analysis using the 3D netbook CAD model.
3. Microprocessor load (measured in watts) which is based upon an active steady-state use condition assumption, a function of what kind of application a user is running on the device.
4. Power usage of each other component in the netbook CAD model design. These assumptions are based upon estimates of power dissipation from numerous components within the netbook device.

The thermal solution of the netbook device is a major input that depends upon what the OEM customer chooses to incorporate into a design. There are a number of decisions that can affect the product's reliability. Intel provides guidance to the OEM customers, but it is ultimately the decision of the OEM whether or not a design solution will be based upon Intel's suggestions.

The simulated netbook design is based upon a model Intel has created for the purpose of thermal analysis, which is founded upon an actual system. However, it also incorporates Intel's judgment of what upcoming products may look like. Figure 7 below is a depiction of this model.*

* Model is based upon Intel's simplified estimate of a netbook computer assembly.

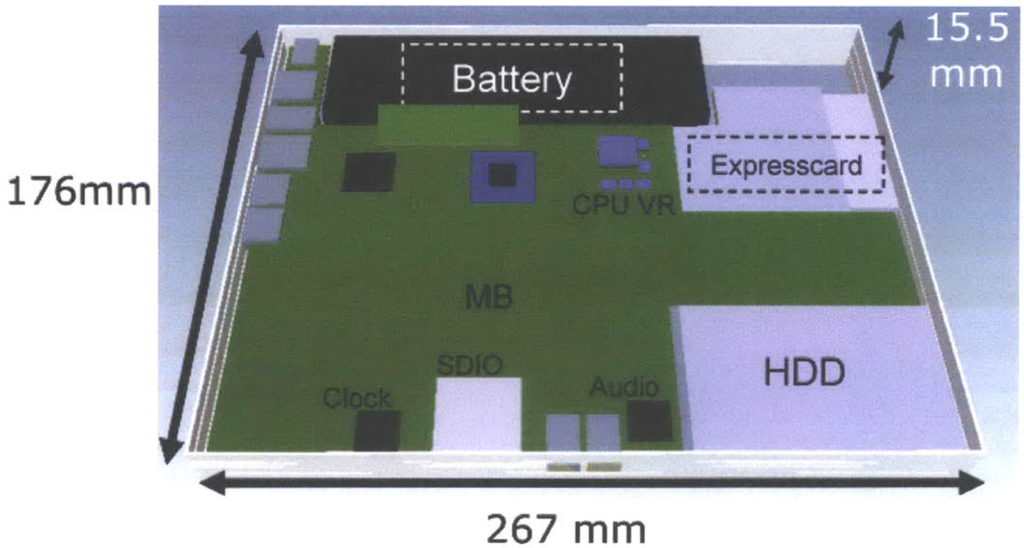

Figure 7: Layout of Netbook Design Used for Thermal Analysis.

For the purpose of simulating actual use conditions, the model incorporates a fanless design. While many netbook products currently have systems with cooling fans, it is thought that most future models will not include this feature. This will create thermal effects that have an adverse impact upon solder-joint reliability.

In the current design, the microprocessor load is based upon the use condition assumption that a high intensity workload application represents the 50th percentile user. For netbooks, this is an intense assumption that causes the microprocessor load to be unusually high during the simulation.

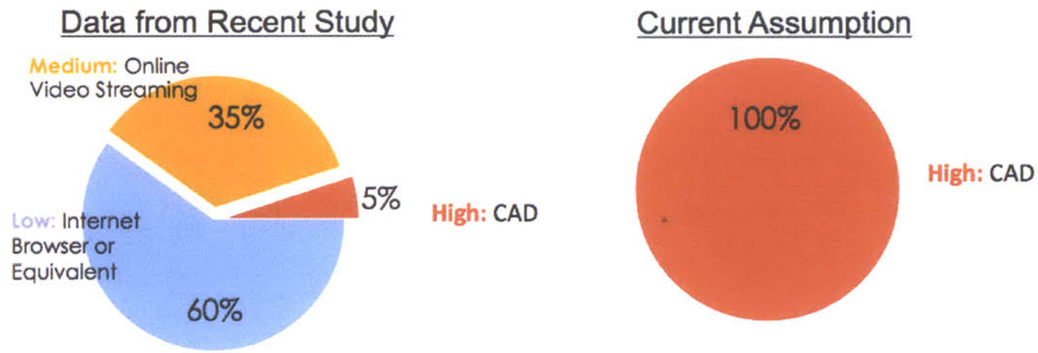

Figure 8: Netbook Application Use Comparison of Recent Findings to Current assumption.[7]

From Figure 8, we see that a recent study shows netbook computers are largely utilized for internet browsing and that high microprocessor usage is rare.[*] In fact, microprocessor usage was found to be less than 30% of total consumption for over 85% of the time that the device was in operation during the study. In the same study it was found that while the device was on and active, the mean microprocessor load was approximately 16%; most of the activity was attributed to web browsing. Running the simulated high intensity application requires a microprocessor load of above 50%, and microprocessor loads at that level lasted for one second or less for more than half of the occurrences. Figure 9 below is a chart depicting the power usage in watts for the microprocessor platform when running a number of different applications.

[*] The CPU load measurements from the study were categorized into three groups which we have equated to low, medium, and high intensity applications.

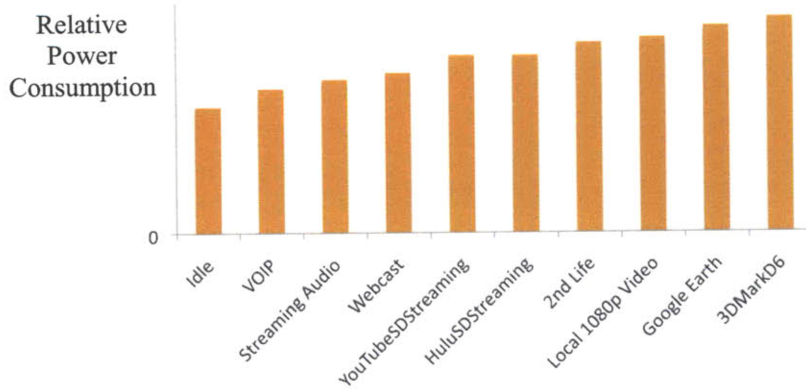

Figure 9: Estimated Relative Power Consumption for Various Applications.[12]

Figure 9 is a relative comparison of power consumption for different applications when run on a netbook computer. All netbook products to date have been designed under the assumption that the median device is being used to run the high intensity applications, which are similar to those towards the right-hand side of the figure. As can be seen in Figure 10 below, recent studies have demonstrated that the majority of netbook use tends to be for the purpose of internet browsing and email, with a portion attributed to watching streaming video online. The belief that a netbook is being used to run streaming online video all the time would likely capture the most strenuous of conditions for most netbook users and would allow for a slightly lower power usage assumption for this netbook processor design. This has a significant effect upon the thermal-mechanical reliability since the microprocessor load requires additional power and in turn dissipates greater energy. Therefore, the assumed application plays a significant role in affecting the reliability targets and it may be appropriate to adopt more accurate deductions for this parameter.

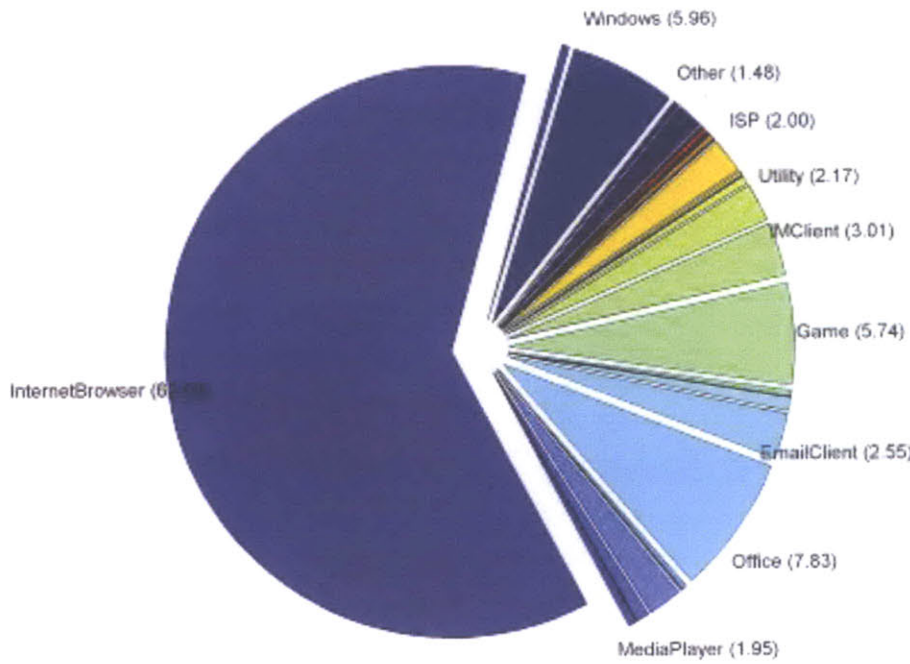

Figure 10: Percentage Breakdown of Netbook Usage by Application.[13]

As mentioned previously, the power usage of each component in the model is also an input for the CFD analysis performed. This is based on a 60-second moving average of each component's projected power. The output from the design and assumptions mentioned above has yielded a predicted solder ball temperature of 77° C. If we theorize that the 50th percentile user is running streaming online video, this value would be reduced to 73° C. In our analysis we will see the effect upon solder-joint reliability of the solder ball temperature and the assumptions used to determine the solder ball temperature.

4 Reliability for Central Processing Units and Intel Design

The Quality and Reliability organization at Intel has created sophisticated programming tools to convert use condition suppositions into reliability targets that can be used in the package-level design. Their methods will be employed in this study and we will use the programming tools to understand the effect of different use condition assumptions and their impact upon the overall target.

4.1 Conception of Reliability Targets

While it seems logical to alter use conditions based upon the best market research, there are internal pressures as well as external costs to be considered. It is important to understand the impact of making changes. For example, if aligning the application use parameters has only a marginal change upon the reliability target, the overall return for making such a change could be negative if the internal cost is high.

An expense associated with changing use conditions is created by altering agreements previously made with OEM customers. Specifically, there is a risk associated with changing parameters that may affect the customer's perceived quality of the product. With respect to the agreed-upon use condition parameters, OEMs may place more emphasis on some items and less on others. Table 4 is an example of usage assumptions that Intel shares with OEMs.

Category	Parameter	Value
Durations	Operating Time	5 yrs (43800 hrs)
Active Operation	Time (%)	23
Inactive Operation	Time (%)	77
Power Cycling	Cycles/Day	6

Table 4: Reliability Information Conveyed to OEM Customers.*

The actual reliability target used to determine if a design will meet the necessary conditions is based upon a test performed by Q&R and ATD. This test places a number of sample parts into a chamber that thermally cycles the parts between two extreme temperatures. The test is specifically designed to simulate the total aggregate effect of all use conditions and OEM design parameters that influence solder-joint reliability. We will use a tool created by the Q&R department in order to analyze the effect of changing use conditions upon the reliability target. This process has been historically used to determine reliability targets based upon a given set of use conditions. In order to estimate the reliability target, we must convert a set of use conditions into the temperature cycle qualification test with the correct parameters. We can ultimately determine the number of cycles to failure which the package must withstand in order to qualify. A certain number of microprocessor samples are chosen for the test and the design is deemed adequate or not depending upon the number of samples that pass.

4.2 Transposing Use Conditions into a Reliability Target

In order to analyze the effect of changing use conditions upon the reliability target (measured in temperature cycles to failure), we will use a program created by the Quality and Reliability department, known as the Speculative High Density Interconnect Package Simulator (SHPS). This tool has been historically used to determine reliability targets based upon a given

* Information is based upon Intel's use condition assumptions.

set of use conditions and design parameters. One of the main outputs is the number of cycles to failure which the package must withstand. As previously mentioned, it is for this target that the solder joint pattern on the backside of the package must be designed in order to meet qualification of assumed use conditions.

The SHPS program was designed based upon extensive research and analysis by Intel's Q&R department. Through careful understanding of the notebook market, the microprocessor design, and the temperature cycle test, experts at Intel were able to put together a simulator, which enhances and speeds up the design process for the microprocessor package. By being able to estimate the conditions that the package must withstand, designers have a far greater chance of providing an optimal solution before qualification tests begin. This saves significant time and money for Intel.

The method used to determine customer use condition sensitivity requires running a series of trials through SHPS and analyzing the effect upon the provided temperature cycle requirement. In order to successfully achieve an adequate analysis, a space-filling uniform design experiment was created that would predict the correlation between each use condition and the reliability target. It is necessary to take a closer look at SHPS in order to understand how it can be used before we discuss the impact upon the use conditions.

4.2.1 SHPS Interface

The SHPS program was created to run simulations based upon a variety of design and use conditions. However, for the purpose of this study, we are concerned with only a limited number of those use conditions. Thus it was more efficient to create an interface which could convert the use conditions from values obtained through market studies into data that is easily used within the SHPS system. To better understand the SHPS interface and the conversion spreadsheets, please refer to Appendix A which contains examples of those spreadsheets.

4.2.2 Statistical Analysis of Targets Prediction

We must complete a statistical analysis within the SHPS program to determine the effect that each of four variables has upon overall reliability. While the cost of running an analytical trial through the software is relatively low (requiring only additional runs through the program), it is still helpful to minimize the number of runs since SHPS is only capable of six runs at a single time. In order to do this, we used a design of experiment which varied the following use conditions: lifetime, active time, power cycle frequency, and microprocessor load.

We need to consider the actual failure mechanism that it affects for the purpose of measuring the influence of microprocessor load. Because CPU load increases the temperature of the microprocessor package, it has a direct effect upon the solder ball temperature. The SHPS program considers this and other inputs which affect the thermal design of the computer. Because the analysis that we were performing was able to also measure the impacts of the thermal design, the MPG thermal group requested that two other variables be included in the study. Neither of these variables are functions of use conditions; they are based upon netbook design parameters and their estimates are subject to change, depending upon the assumptions made regarding OEM designs. For this reason, it is practical for us to include them in our statistical analysis. The first variable, known as T_{rise} is based upon the heat generation from all components within the computer other than the microprocessor. It acts as a bias to the solder ball temperature and has a major impact as we will see. The other variable, denoted by θ, has units of $\frac{Temperature}{Power}$ and is a term that converts CPU Load (measured in watts) into a temperature. Through these variables, the solder ball temperature is calculated as follows:

$$T_{sb} = T_{rise} + (\theta * CPULoad)$$

Within the DOE, we vary all three variables, T_{rise}, θ, and CPU Load in order to see the effect of each. However, with respect to understanding customer use conditions, the CPU Load variable is the only one that concerns us. As mentioned previously, studying the effect of the other two variables has other uses, such as helping the MPG group understand the reliability impact caused by those design parameters. Ultimately, within SHPS, the solder ball temperature is the actual input for the design. However, for further analyses, the solder ball temperature will be represented by the three components that comprise it.

The statistical analysis software known as JMP was used to create a 60-run spce filling uniform design of experiment (DOE). The purpose of the DOE is to include enough runs to adequately estimate the use condition effects within the ranges chosen for each. Every run was put through SHPS to receive a corresponding output. The output from SHPS was then entered back into the JMP, so that the software could perform a regression analysis and create a prediction formula based upon this output. The prediction formula is a second-order polynomial response surface model. Following its formulation it was tested against a new random set of 30 runs in order to determine its accuracy in predicting the output of SHPS.* The JMP prediction and the actual SHPS values were identical up to an average error of approximately 1.20%. This suggests that the prediction formula is a very accurate representation of SHPS and can be used to determine the reliability target for any set of use conditions within the ranges utilized to create the DOE.

4.3 Study of Use Condition Relation to Reliability Targets

A graphic representation is a helpful way to intuitively understand the prediction formula. Within JMP this was made possible with the profiler tool. Figure 11 is a single set of use conditions in the profiler. The profiler picture below depicts a single set of use conditions and the slope of each curve shows the impact that each condition has if all others remain static at their

* See Appendix B for 30 Sample Prediction Formula Test Runs.

listed values. For example, for the set of values shown in Figure 11, power cycle frequency has the largest influence on the number of cycles to failure since its slope is the highest. However, the impact of power cycle frequency may also change if any of the other values are changed.

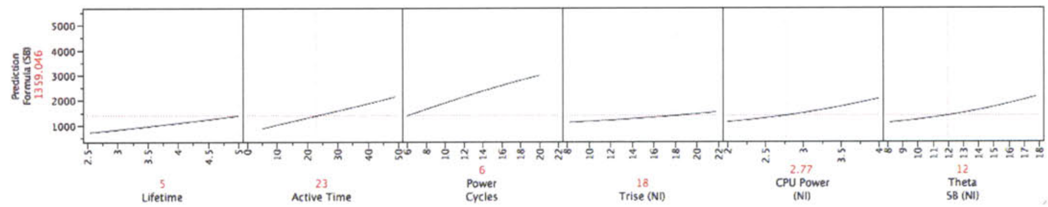

Figure 11: Profiler View of Current Use Conditions.

It can be seen that the current set of conditions produce a reliability target that requires the package to withstand 1359 temperature cycles in the accelerated lifetime test. In other words, this is the number of accelerated temperature cycles that the design would need to meet to pass qualification. It is worth noting that in order to make continued use of the JMP profiler, it is necessary to use the tool itself. The figure only shows us the values, sensitivity, and linearity of each variable based upon a single set of conditions, but the real use of the tool is its ability to show changes as the values of one or more variables is altered. There is no way to graphically represent the sensitivity of all variables at the same time because this is a multivariate problem with six dimensions. However, we can change values in order to understand the impact of such changes.

Another main point of interest is the set of conditions that we have recommended previously. Table 5 is a summary of those use condition values.

Use Condition	Value
Lifetime	3.75 Years
Active Time	17%
Power Cycle Frequency	6 cycles per day
Application Use (microprocessor Load)	2.77 Watts*

Table 5: Recommended Possible Use Conditions.

When we enter these data into our JMP profiler model, we get the result shown in Figure 12 below.

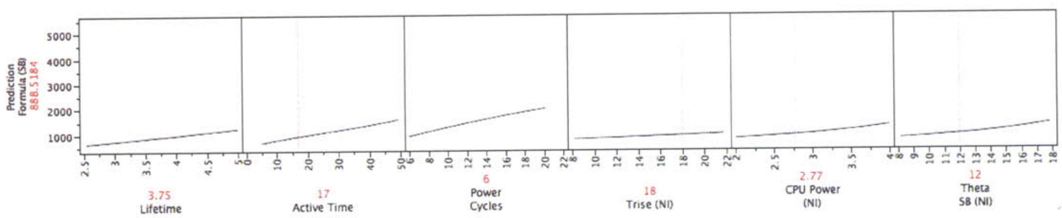

Figure 12: Profiler Prediction with Customer-aligned Use Conditions.

From this figure, we can see that the reliability target predicted by SHPS reduces significantly from the current prediction. With the set of conditions from Table 5, the number of temperature cycles that the package would need to withstand is 888 in order to meet qualification. This represents a 42% reduction in required solder-joint reliability, which means that the package-level design would be required to withstand over 450 less thermal cycles during the qualification tests. If fewer cycles are required for qualification, then less NCTF solder joints may be needed, allowing the package size to be decreased. Later we will see the effects of this kind of reduction.

* We have chosen not to vary the CPU load, as the change is relatively minimal with currently optimized processor conditions. However, it is recommended that the future standard be changed from high intensity to medium intensity, which has been determined to be a more appropriate assumption for 50th percentile netbook use.

It is worth noting that a change to any specific use condition not only has an effect upon the reliability target, but it also changes the sensitivity of other use conditions. For instance, if the power cycle frequency is reduced, then the impact that lifetime has upon overall reliability also decreases. The profiler tool helps us understand this relationship by reducing the slope of the lifetime curve when the power cycle frequency value decreases. The reason for this relationship is that the algorithm that runs the SHPS program is based upon the assumption that power cycle frequency and magnitude are the major levers of solder-joint reliability. The magnitude of a power cycle can be thought of as the temperature range that the solder joint must undergo when the state of the microprocessor changes. If the temperature range between an on and off state is high, then the magnitude of that cycle is also high, and the effect upon the solder-joint is more pronounced. This value is captured in the solder ball temperature calculation that considers the load on the microprocessor as well as the design parameters of the netbook that have an effect upon thermal performance. As we have discussed earlier, power cycle frequency is a function of the netbook user. If the frequency increases, then an increased lifetime will have a greater overall impact upon the required reliability since the microprocessor will see a larger number of total cycles during the additional years of use.

Since we now have some understanding of the effect of use conditions upon the overall reliability target, we can focus our attention on the microprocessor package design to see where there are opportunities to optimize the design.

5 Microprocessor Package Design Impact

The solder-joint reliability target has a direct impact upon the package-level design. In order to understand these effects, we consulted with the Assembly Technology Development group at Intel and leveraged their experience with design for reliability in microprocessor products. The findings show that varying use condition assumptions to create more realistic targets has an effect upon the package size. We will look further into the causes and effects of such reliability target variations.

5.1 Overall Package Design

We have spoken about the need for non-critical to function (NCTF) solder joints, but it is worth taking a closer look at the layout of the microprocessor package. The front-side of the microprocessor consists of a silicon die, which is the heart of the microprocessor, containing billions of transistors in an area as small as a few square millimeters. The die is mounted to a printed circuit board (PCB) package that is designed to extend the network of the silicon transistors to the motherboard of the computer.

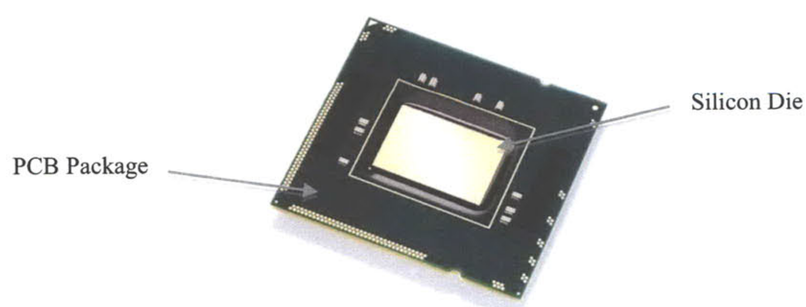

Figure 13: Image of Microprocessor Package Front-side.

From Figure 13, we see the silicon die sitting atop the PCB substrate. The printed circuit board consists of multiple layers the quantity which varies depending upon the complexity of the design. A significant amount of effort goes into the design consideration for the package substrate since it is an expensive and sizeable component of the microprocessor. Intel designs its own packages but sources them from outside vendors. One impact on substrate size is the amount of circuitry required to transfer information from the die to the motherboard; another is the array of solder joints that reside on the backside of the package.

From Figure 14 we see a side view of the microprocessor and the solder joint pattern that adheres the package to the computer motherboard.[14] This pattern consists of a number of CTF and NCTF solder joints between the die and package.

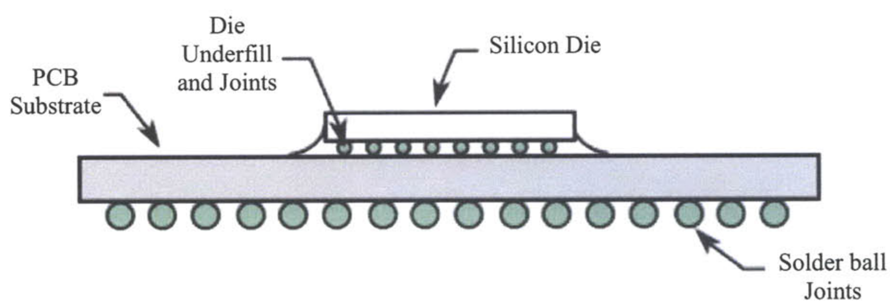

Figure 14: Side View Depiction of CPU Package.

The overall thickness of the package is determined by the thickness of the die, the bond area between the die and the substrate, the thickness of the substrate, and the solder joints between the package and the motherboard. The die is attached to the package using solder joints and strengthened with an epoxy underfill.

5.2 Critical to Function and Non-Critical to Function Solder Joints

Critical to function (CTF) solder joints, like their title suggests, are necessary in order for the microprocessor to function correctly and drive the computer. They have the dual purpose of transferring information and providing stability to the microprocessor-motherboard bond. In order to meet increasing demand for reliability, Intel has included solder joints solely for the purpose of providing structural stability, which are non-critical to function. These solder joints have an impact upon certain aspects of the design, including the size of required substrate. Specifically, adding NCTF solder joints can require other features to be moved away from the center of the package or can create signal routing difficulty, in both cases causing the package size to increase in order to accommodate the additional joints.

Recalling Figure 1, which is a depiction of the solder joint array on the backside of the microprocessor package, there are solder joints underneath the die as well as under the package corners. Most of the joints underneath the die are CTF while the corner joints are NCTF. The NCTF joints on the package corners are mainly for the purpose of structural stability to counter the effects of mechanical shock due to catastrophic failure. These are modes separate from thermal reliability and include situations where a netbook device is exposed to physical impact, such as being dropped. For protection of CTF joints against thermal-mechanical effects, NCTF joints are needed around the border of CTF joints in the die shadow region.[15] The stress from thermal expansion--which is caused by the expansion coefficient mismatch between the silicon die, the PCB package, and the system motherboard--is strongest in the die shadow region. It is this area where NCTF joints can be avoided if the CTF solder joints provide enough stability to withstand the thermal reliability target.

5.3 Analysis of Reliability Effect on Solder Joint Pattern

We have mentioned that customer use conditions have an effect upon the overall microprocessor package design and it is now necessary to look further into the actual impact upon the design. The Assembly Technology Development and the Quality and Reliability organizations at Intel have done extensive work to understand the relationship between reliability and package-level design. Furthermore, recent tests with an upcoming netbook microprocessor platform have produced results that help us understand this relationship.

5.4 Experiments with New Netbook Processor Design

A series of tests was carried out under differing conditions to determine the level of reliability the current design is capable of meeting. These tests were performed in a manner similar to that used for qualification of new products. The samples chosen for these tests included ones with different substrate thickness as well as those under varying compression loads. The reason for placing the microprocessor package under compression during this test is for the purpose of simulating the stress that it may undergo when assembled into the netbook computer. Under certain conditions, usually when a cooling fan is used, the microprocessor package can be sandwiched between other assembly components. This adds to the stress caused during thermal cyclic loading since the solder joints are forced further into compression during a period when thermal expansion is occurring.

The qualification test included different levels at 890, 1250, 1500, 1850, 2150, 2450, 2750, and 3000 cycles, which means that the samples were analyzed at each of these check points. The conditions included 8 lbs of compressive load and the samples were mounted to a PCB which was 40 thousandths of an inch thick. This represents conditions similar to actual use, but it is important to note that OEM designs will mostly be fanless in the near future. For fanless designs, it is possible that the compressive load on the microprocessor may be lower than 8 lbs, which means that these test conditions may be slightly more harsh than necessary. This provides a high level of confidence that the design will meet the necessary requirements in actual use.

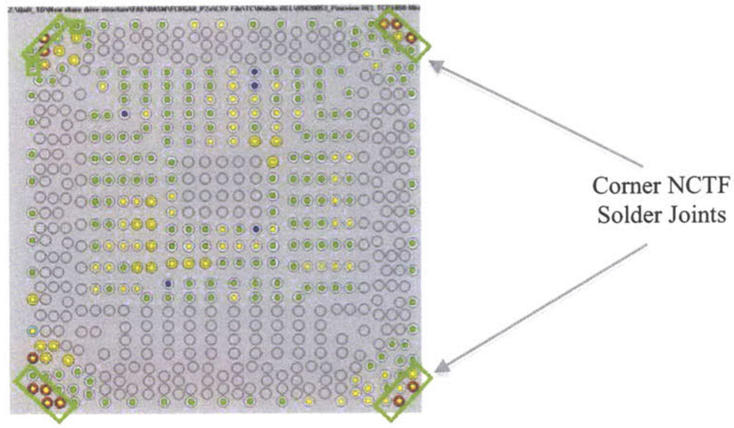

Figure 15: Depiction of Failure Recorded during Tests.

Following the test an extensive analysis of the results following the qualification trial was completed. While the performance of the solder-joint adhesion was relatively strong, their findings indicate that critical failures were noticed in specific regions between two checkpoints during the test. Altogether, their analysis concludes that a first-level failure occurs at 1250 cycles and a second-level at 1850 cycles. This means that one row of NCTF solder joints is required in order to surpass 1250 cycles and a second to meet an 1850 cycle target.

5.5 Design Impact

Under specific circumstances, it is possible to add rows of NCTF joints without increasing the size of the substrate and microprocessor package. This is done by rearranging some of the CTF solder joints through creative methods employed by designers familiar with the microprocessor package layout. In general, each additional row of NCTF solder joints will increase the overall package size by approximately 1 mm x 1 mm.[16]

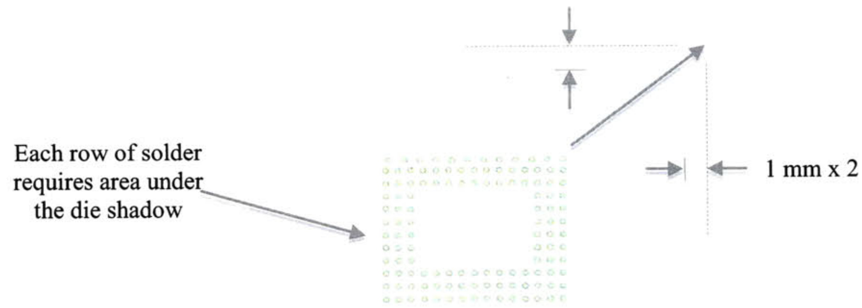

Figure 16: Impact of Adding Rows of NCTF.

Figure 16, re-representing Figure 1, allows us to see that adding rows of NCTF increases the amount of area occupied underneath the die shadow. While increasing the amount of solder joints on the package, there is less area available for other features and on average the package size will increase by 1 mm x 1 mm for each additional row of NCTF joints under the die shadow. In some cases, this can increase the package size by the same amount of area required for the additional solder. Ultimately, the additional PCB for the package is a significant cost to Intel.

6 Financial Analysis

The design impact we have discussed has direct cost implications for microprocessors. The actual cost impact for a specific design change can only be determined on an individual basis, which is very difficult to predict ahead of time. For this reason, the cost increase per part can be estimated but must be considered on an average size increase basis for prediction purposes. In this chapter we will review the direct cost implications to determine what kinds of benefits may exist for re-aligning customer use conditions.

6.1 Substrate Cost Impact

By combining the analysis from the test results, the tools created to predict reliability targets, and our use condition studies, we can estimate what costs will be incurred to maintain the current set of use conditions. As mentioned above, designs using this package technology will each require an additional NCTF row at 1250 and 1850 temperature cycles. Each additional row above the baseline substrate size will on average increase the substrate by 1 mm in both length and width dimensions. The cost of adding PCB package area to the microprocessor package is variable and is based upon the manufacturing costs incurred by Intel's substrate suppliers. The PCBs are manufactured in fixed panel sizes, which are then separated into smaller pieces, based upon the specified design. For instance, a PCB from one supplier might be converted into 100 packages when the substrate size is equal to 30 mm x 40 mm. This may happen with some remaining amount of leftover PCB whose dimensions are incapable of producing any more 30 mm x 40 mm packages. For this reason, a slight increase to the 30 mm x 40 mm substrate size may yield 101 packages, but with less leftover PCB. In such a circumstance, the slight substrate size increase may affect cost significantly depending upon whether or not the size increases to a point where the yield drops. For instance, if the substrate size increase now creates a situation where only 95 packages can be created from the parent PCB, then the cost per package will increase. Furthermore, each supplier of PCB has different manufacturing conditions which will yield cost differences for varying sizes and shapes. Intel sources from multiple vendors in order to reduce cost as much as possible. One can imagine that the costs incurred by a substrate size increase can vary significantly on a case-by-case basis. However, for the purpose of project costing and NPV analysis, the ATD finance team is able to estimate the average cost increase for packages larger than the size which is currently in use.

Substrate Size (mm)	Per Unit Cost Increase Above Baseline Substrate Size[17]
Baseline	–
+ 1 mm x 1 mm	$0.12
+ 2 mm x 2 mm	$0.19
+ 3 mm x 3 mm	$0.31
+ 4 mm x 4 mm	$0.38
+ 5 mm x 5 mm	$0.44

Table 6: Per Unit Cost for Pertinent Substrate Size Increases.*

Table 6 consolidates the estimates for the package sizes that concern us. It is evident that the cost increase per unit is not directly proportional to the increase in area for each substrate. As mentioned previously, this is largely due to the fact that major substrate production limits will be reached only for certain size increases. However, the table does average cost increases over multiple vendors while smoothing out the effect caused by major step increases in production costs.

* Data for table is based upon an Intel's averaged estimate of incremental cost increases that could be incurred for larger substrate sizes.

6.2 Prediction of Overall Cost Impact

With the previous test results, we can determine when we will require new rows of NCTF joints, and subsequently when we will need to add substrate to the package. If we consider the market forecasts for upcoming unit sales, it is possible to estimate the total cost impact of using current use conditions. We can then consider what benefits might arise by using more appropriate use parameters, such as those discussed earlier.

Any design changes to Intel's netbook microprocessor will be made on the platform released in 2011 since the next design model has already been released to market.* Forecasts for this model have been obtained and will be used to estimate the total cost benefit for only the two-year period in which this model will be produced. From the previous analysis, we determined that the current set of use conditions requires a reliability target of 1359 temperature cycles. This falls above the failure limit of 1250 cycles, but below the second failure limit of 1850 cycles. However, because the value is greater than 1250, it would require an additional row of NCTF joints and will increase the substrate size by 1 mm x 1 mm.

* The earlier tests analyzed in section 5.3.1 were performed on the model which has already been introduced to market. However, the financial analysis is based upon the impact to the following model.

Reliability (Cycles to Failure)	Substrate Size (mm)	Per Unit Cost Increase Above Baseline Substrate Size	Predicted Increase Over 2-year Period*
0 – 1250	Baseline	–	–
1250 – 1850	+ 1 mm x 1 mm	$0.12	$24 Million
1850 -	+ 2 mm x 2 mm	$0.19	$38 Million
	+ 3 mm x 3 mm	$0.31	$62 Million
	+ 4 mm x 4 mm	$0.38	$76 Million
	+ 5 mm x 5 mm	$0.44	$88 Million

Table 7: Summary of Cost Impact Based upon Reliability Target Increase.

Table 7 is a summary of the total cost impact over a two-year period based upon an increase in reliability targets for Intel's netbook microprocessor. The highlighted line is the estimate for the cost which will be incurred if the current set of use conditions is maintained. As shown in the table, the current conditions will cause a $24 million increase in total substrate costs if not changed. This is obviously a significant cost for Intel and can be avoided by adopting appropriate netbook usage assumptions.

* Financial benefit based upon total volume estimated during the timeframe following product release.

7 Recommendations and Further Steps

In the future, Intel can avoid additional cost by better aligning its netbook use condition assumptions to those actually observed in the market. As we have previously shown, there is strong reason to believe that the current set of assumptions is unrealistic. We have also seen that there are strategic and financial benefits to making realistic changes. There are numerous ways to implement such changes and we will discuss some of the favorable actions Intel can take in the near and long term.

7.1 Initial Levers

It was observed that the current use conditions place the reliability target estimate at 1388 temperature cycles while the estimated threshold requiring increased package size occurs at 1250 cycles. The analysis performed on the SHPS program allows us to see the effect of certain assumptions upon overall reliability. While the recommended use conditions discussed earlier provide an estimated target of 888 temperature cycles, there are other sets of conditions that predict a value lower than the 1250 cycle threshold. The percent of active time fits well in these categories when considering which use conditions can be easily both challenged and measured. Data has already been gathered to strongly support the fact that the amount of time a netbook is actively used is far lower than the current assumption. Most studies of netbooks have also shown that the usage patterns of netbooks are further away from notebooks and towards the direction of smartphones and PDAs. All of this data directly measures or suggests that active time usage for netbooks is lower than the 23% assumption. Earlier, we suggested using a lower but still conservative estimate of 17%. Although evidence points to an even lower 90^{th} percentile value, 17% could be adopted as the new standard for netbook active usage time. Let us take a further look at what this does to our reliability target estimate.

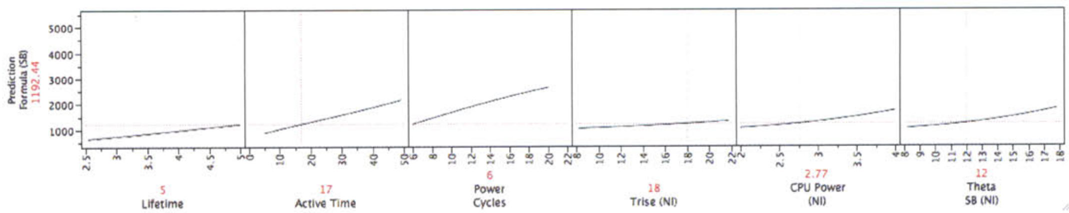

Figure 17: Profiler Reliability Target Prediction for 17% Active Time Usage.*

From the JMP profiler, the predicted reliability target is slightly less than 1200, at 1192 temperature cycles, if all use conditions are held to their existing values but active time is reduced from 23% to 17% (see Figure 17 above). By changing the active time usage assumption, we are able to reduce the target below the first threshold where a row of NCTF joints would be required. As detailed in our financial analysis, this amounts to a cost reduction of $24 million. While it is strongly recommended that further data be gathered in upcoming months, this is a reasonable change which should be made before qualification of the next processor model.

7.2 Data Gathering and Next Steps

Since netbooks are a very recent product in the consumer market, it is understandable that Intel and other OEMs do not yet fully understand usage patterns. Currently at Intel, total device lifetime is often assumed to be the main factor in designing for reliability. As device lifetime requires significant time in market for making decisive arguments as to the correct value, it is not possible for netbook lifetime to be accurately estimated at present since the product has been on the consumer market for slightly over two years. Although we can strongly conclude that the lifetime is shorter than the currently assumed value of 5 years, it will be difficult to convince major stakeholders at Intel that it is worth changing this value while conclusive evidence is not yet available.

* All other use conditions are held to their currently assumed values.

However, as we have seen, there are other major factors that significantly affect solder-joint reliability, which can be measured at this time. Among these are power cycle frequency and active time usage. Conclusions from this study strongly recommend that the MPG marketing and Q&R groups at Intel work together to gather data which can mutually benefit both groups. MPG marketing is mainly concerned with learning more about Intel's end-users for marketing purposes and active time usage and power cycle frequency are patterns that directly fit within the metrics of their studies. Q&R has also conducted numerous studies to gather such information, but their resources in terms of labor and finances are often limited. If the two organizations join together to implement studies that capture both groups' needs, they can obtain data for active time and power cycling relatively quickly, before the next processor model is qualified. The Insights and Market Research Group and the People and Practices Research groups at Intel have already done an excellent job of gathering pertinent information and their understanding of reliability concerns is also strong. MPG Marketing and Q&R can work with these organizations to conduct studies that will help to accurately determine active time usage and power cycle frequency for netbooks. In the current resource-constrained environment, it may be difficult for Q&R to justify further research. For this reason, Q&R may benefit by using MPG and IMR analyses in order to save money and avoid replication of data gathering.

7.3 Perception of OEM Customers

Since Intel's direct customers are actually the original equipment manufacturers (OEMs), it is worth noting their possible perception of changing reliability targets. Generally, Intel's OEMs have not been overly concerned with some details of reliability specifications. We make this argument since recent alterations were made to reflect a change in specific details regarding Intel's predicted reliability performance for mobile products. It appeared that OEM customers did not have any concerns since it did not affect their core assumptions. With respect to netbooks, we have already seen that OEMs do not expect the reliability of the entire product to meet that of notebooks. This evidence was apparent in the warranty information that was reviewed earlier. It is difficult to think that they would have strong concerns if targets such as active time usage and power cycle frequency were changed to reflect realistic usage.

Considering the fact that smaller form factors continue to be a major design parameter, especially for the case of smaller netbooks, OEMs are largely concerned with having flexibility to create new designs. A smaller microprocessor package is a key parameter for making that possible.[18] For this reason, one of Intel's major concerns is being able to provide a smaller package which is more attractive to OEMs. While it is difficult to quantify the financial benefit, smaller sizes are preferred. This makes the design parameter a point of competitive advantage for the microprocessor manufacturer that can make a product with the desired form factor. It is important for Intel to understand the tradeoff between higher reliability and smaller size. OEMs tend to be discreet and do not directly offer their opinions on this matter. As we have shown, however, Intel is already designing well above the necessary level of reliability required, and should not need any additional increase in package substrate size. This makes the decision relatively straightforward, yet a situation may arise where there is a direct tradeoff between size and reliability.

7.4 Recency of Netbook Market and Future Considerations

Although we have underscored the fact that the netbook market is young, making lifetime specifications difficult to confirm, there are other factors regarding this new product that should be considered. Originally, the major breakout market for netbooks was thought to be emerging economies where price points would be low enough that consumers unwilling to pay the higher price for notebook computers would purchase a netbook instead. These were projected to be customers who do not currently own a computer but would use a netbook as their primary device. Thus far the main consumers of netbooks have been users who already own a primary device.

However, this trend could change since the product is fairly new. Whether or not that has an effect on reliability remains to be seen, but this pattern should be considered and further tests to determine consumer usage patterns should try to capture emerging markets and potential primary device users.

7.5 Implementing Change at Intel

Intel has been producing microprocessors for almost 40 years.[19] It has a proud history of providing products with a high level of reliability. The Quality and Reliability organization has the core responsibility of ensuring that all of Intel's products meet the highest standards possible in order to protect the brand and reputation of the company. Intel's brand remains one of its strongest assets. Other business units within Intel, such as the Mobile Platforms Group, manage their own profit and loss statements and are concerned with maximizing their financial returns. This creates a situation where organizations are incentivized in different ways while trying to maintain the common goal of maximizing value for Intel's shareholders.

It appears on the surface that each division involved with the netbook microprocessor package has distinctly different goals. In reality, there is overlap in certain areas. There is no group closer to the customers and the market for netbook devices than MPG; they are in constant contact with end-users and OEM customers of their products. As a result, MPG has access to specific and unique information which may be relevant to the work that Q&R performs when determining optimal reliability targets for netbook processors. For example, MPG conducts market research to determine the average amount of time that a netbook may be used. This information is important for calculating economic value of the product. It is also relevant to the work that Q&R does to determine lifetime reliability needs. Currently, there is no system in place for ensuring that such information is transferred from one group to another. Q&R in the past has used different sources and conducted their own initiatives to determine information similar to what MPG marketing produced. A critical piece missing from this strategic structure is a method for direct communication between MPG Marketing and Q&R with respect to determination of customer use conditions.

Another area of overlap is Intel's interaction with its OEM customers. Members of MPG work on product strategy and interact closely with the OEMs on a variety of market-related subjects. These discussions can include technical issues relating to product specifications. In parallel, there are groups within ATD that have close interaction regarding technical specifications for Intel's products. Although different groups have different goals, it may help to have direct points of contact between Intel and its customers in order to streamline the process of gathering information for optimizing the netbook processor.

The Q&R group is sometimes called the policeman of the company. This is due to its role as an organization that ensures adequate quality for each product even when business units are satisfied with less. This function is especially important for a company involved in the fast-moving semiconductor industry where the instinct can be to shift quickly at any cost. Other corporations have lost their competitive advantage due to lower quality products while Intel has maintained its status as a company with high standards. The Intel brand, consistently rated as one of the strongest in the world, relies heavily on this reputation. This makes the quality organization very important and gives it a significant amount of power and influence. It has also created an environment where some members of Q&R are protective of certain practices. This can contrast with the business-directed mindset of MPG, which focuses on efficient time-to-market and low-cost products. The discussion of altering use condition assumptions has in the past become a battle between one organization which feels it is protecting the reputation of Intel by maintaining high-quality products and another group which is heavily measured by its financial performance.

With respect to managing the manufacturing units at Intel, there are different viewpoints and incentives for ATD and MPG. MPG is close to Intel's customers and is concerned with providing different stock keeping units (SKUs) for some products which appear relatively similar. ATD is in charge of assembly-level manufacturing and works to ensure that low cost and efficient processes are used wherever possible. For this reason they prefer fewer SKUs, in order to reduce manufacturing costs, sometimes at the expense of increased costs for the OEM customer. This creates a political struggle in which two organizations prefer different models in order to maximize value. Each group has strong reasons for maintaining its core set of values because each is important to Intel as a whole.

7.6 Looking Forward

We have already discussed how data gathering should be the immediate next step. However, it may also be necessary to consider whether or not Intel is well-positioned to enter the low-cost microprocessor industry where it will compete against companies that obtain much lower margin for their products. Intel has enjoyed substantial profits through its core businesses over several decades and has held onto an enormous share of the mobile, desktop, and server markets. With its newest processor Intel is moving into markets where margins are much lower and manufacturing costs are a larger factor.

The current organizational structure is set up to create an optimal balance in power between MPG and Q&R to allow for adequate reliability, low cost, and quick time to market. However, the current operation of this system may not be efficient enough to allow quick changes in reaction to volatile and rapid movements in the specially low-cost and fast-moving netbook market. It may be necessary to look into modifications of the system that will permit more direct collaboration between the groups that share responsibility for understanding end-user behavior. Although the notebook, desktop, and server industries have each grown steadily and remained substantial for long periods of time, new markets such as netbooks may arise and evolve much more quickly. In order to be agile and remain ahead of the competition, Intel's stakeholders of package-level design require more direct collaboration. It is hoped that this study has identified areas for direct communication and awareness for methods by which reliability can be determined and quantified.

7.7 Key Takeaways

Intel's use condition assumptions affecting thermal-mechanical reliability for netbook microprocessor packages are relatively conservative and represent a smaller minority of users than what they are attempting to target. While all four discussed use conditions should be analyzed more carefully, it is worth investing time and effort in the near term to understand active time usage and power cycle frequency. Several studies with appropriate sample sizes will observe actual usage behavior that can directly measure both of these parameters. Based upon recent tests, there is reason to believe that the next processor model may not qualify under the current set of assumptions. More conclusive data for at least one or two use conditions will be very beneficial in order to be prepared for such a situation. If it can be determined with more certainty that the reliability target may be reduced based upon the adjustment of some use condition assumptions, then less NCTF solder joints may be necessary, allowing the package size to remain as is. This effect should allow Intel to avoid an additional cost of $24 million over a two-year period.

Appendix A: SHPS Conversion Spreadsheet

User input conversion for varying lifetime, percent active time, power cycle frequency, and active CPU Load.

		Market Segment					
Use Conditions		**1**	**2**	**3**	**4**	**5**	**6**
Total Lifetime		5	4.5	4	3.5	3	2.5
Percent of Time Active		23	20	17	14	11	8
Power Cycle Frequency (cycles/day)		10	9	8	7	6	5
CPU Load (Active)		3	2.7	2.43	2.187	1.9683	1.77147
CPU Load (Idle)		1.5	1.5	1.5	1.5	1.5	1.5
Hours per Day in Use		24	24	24	24	24	24
End-User Input		**1**	**2**	**3**	**4**	**5**	**6**
Operating hours (Run State)	Hrs	**43800.00**	**39420.00**	**35040.00**	**30660.00**	**26280.00**	**21900.00**
Non-Operating hours (Off State)	Hrs	**17520.00**	**21900.00**	**26280.00**	**30660.00**	**35040.00**	**39420.00**
Tair Climate Mean (Off State)	°C	**24.00**	**24.00**	**24.00**	**24.00**	**24.00**	**24.00**
Run-Off power cycle mean	cycles/hr	**0.42**	**0.38**	**0.33**	**0.29**	**0.25**	**0.21**

Time in state input for same conditions as above. The State 1 column represents solder ball temperature while active and State 2 represents the idle condition.

Market Segment 1	**State1**	**State2**
Time in State [Fraction]	0.2300	0.7700
Temperature Mean °C	**80.03**	61.0000
Market Segment 2		
Time in State [Fraction]	0.2000	0.8000
Temperature Mean °C	**76.427**	61.0000
Market Segment 3		
Time in State [Fraction]	0.1700	0.8300
Temperature Mean °C	**73.1843**	61.0000
Market Segment 4		
Time in State [Fraction]	0.1400	0.8600
Temperature Mean °C	**70.2659**	61.0000
Market Segment 5		
Time in State [Fraction]	0.1100	0.8900
Temperature Mean °C	**67.6393**	61.0000
Market Segment 6		
Time in State [Fraction]	0.0800	0.9200
Temperature Mean °C	**65.2754**	61.0000

Appendix B: Results from Test of Prediction Formula

SHPS Output	Prediction	Difference	% Diff
1298	1280.396	18	1.37%
1042	1031.739	10	0.98%
2923	2955.756	-32	1.10%
4170	4174.077	-4	0.11%
1991	2091.82	-101	5.09%
1324	1295.944	28	2.10%
1853	1832.988	20	1.07%
1099	1091.948	7	0.62%
2164	2150.153	14	0.65%
580	617.9618	-38	6.54%
2885	2849.804	35	1.21%
3614	3602.106	12	0.33%
2645	2644.506	0	0.00%
3559	3620.202	-61	1.72%
2161	2160.129	0	0.02%
1846	1851.082	-5	0.25%
4605	4460.784	145	3.14%
2836	2828.608	7	0.25%
1632	1625.545	6	0.38%
976	982.9856	-7	0.71%
2774	2815.073	-41	1.48%
4402	4435.792	-33	0.76%
1504	1500.694	3	0.19%
1263	1257.597	5	0.42%
2700	2702.301	-2	0.07%
2044	2006.949	37	1.79%
1586	1583.581	3	0.18%
1001	995.021	6	0.65%
2517	2590.116	-73	2.89%
916	916.9526	-1	0.08%
	Average		**1.20%**
	Standard Deviation		**1.50%**

Additional Readings

DisplaySearch. "Mini-Note PC (Netbook) Shipments Grow at Twice the Rate of Notebook PCs in Q2'09." www.displaysearch.com/cps/rde/xchg/displaysearch/hs.xsl

"CES 2009: The price of falling prices, and a path to Intel's potential demise - 1/8/2009 - EDN." http://www.edn.com/article/ca6627620.html?text=CES+2009.

Copeland, Michael V., Partanen, Anu. "Intel's Secret Plan." Fortune Magazine; Volume 159 Issue 11, p52-57, May 25, 2009.

Intel Solution Services. "Proactive PC Replacement Offers Insurance Company $4.6 Million in Potential Savings." Intel Corporation. www.intel.com/internetservices/intelsolutionsservices/

Dumenco, Simon. "How the $0 netbook might just help save the media industry." Advertising Age; volume 80 issue 21, p13-13, June 8, 2009.

Beavers, Brad. "The Story behind the Intel Atom Processor Success." IEEE Design & Test of Computers. Issued Date: March 2009, p8-13.

Gladstone, Darren. "Next-Generation Netbooks: The New Ultra-portables." PCWorld magazine. www.pcworld.com. April 2009, p78-84.

References Cited

[1] "INTC: Key Statistics for Intel Corporation - Yahoo! Finance." http://finance.yahoo.com/q/ks?s=INTC.

[2] "IDC: AMD takes processor market share from Intel - TechSpot News." http://www.techspot.com/news/37715-idc-amd-takes-processor-market-share-from-intel.html.

[3] "Netbooks, MIDs, and Mobile and CE Device Data." ABI Research, Oyster Bay, NY. 2009.

[4] David Daniel Pierce, Market Sizing and Forecasting, Intel Corporation.

[5] "Proactive PC Replacement Offers Insurance Company $4.6 Million in Potential Savings." Case Study, PC Best Practices, Intel Corporation. 2003.

[6] "7 to 5 year lifetime WP – Mobile." Shubhada, Sahasrabudhe, Assembly Technology Development, Quality and Reliability, Intel Corporation.

[7] "Netbook Usage Study." Ken Anderson, Gianluca Iannaccone, and Tye Rattenbury, People and Practices Research, Intel Corporation.

[8] Tye Rattenbury, People and Practices Research, Intel Corporation.

[9] End-user Power Cycle Frequency Analysis. Tye Rattenbury and Eric Baugh, People and Practices Research, Intel Corporation.

[10] Anderson, Iannaccone, Rattenbury. Op. cit.

[11] Carin Ruiz, Client Platform Engineering, Mobile Platforms Group, Intel Corporation.

[12] Adjoavi Worjoloh. Platform Power-Performance Forecasting, Intel Corporation.

[13] Anderson, Iannaccone, Rattenbury. Op. cit.

[14] Intel Corporation. Ball Grid Array (BGA) Packaging. Chapter 14, 2000 Packaging Databook. www.intel.com/design/packtech/ch_14.pdf

[15] Kayleen Helms and Andreh Janian, Assembly Technology Development, Intel Corporation.

[16] Andreh Janian, Assembly Technology Development, Intel Corporation.

[17] Mike Masci, Assembly Technology Development, Finance, Intel Corporation.

[18] Ladd Campbell, Patty Brusso, Xiao Ma, Assembly Technology Development, Intel Corporation

[19] "Intel's First Microprocessor—the Intel 4004." http://www.intel.com/museum/archives/4004.htm.