Enabling Technology for Mobile Computing

by

Vivek Phanse

B.S. Chemical Engineering
Yale University 1996

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

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

in Conjunction with the Leaders for Manufacturing Program

at the
Massachusetts Institute of Technology
June 2000

Signature of Author:______________________________________
MIT Sloan School of Management
Department of Materials Science and Engineering
May 5, 2000

Certified by:__________________________
Lionel C. Kimerling
Thomas Lord Professor of Materials Science and Engineering
Thesis Supervisor

Certified by:__________________________
Roy Welsch
Professor of Statistics and Management Science
Thesis Supervisor

Accepted by:______________________
Carl V. Thompson
Stavros Salapatas Professor of Materials Science and Engineering
Chair, Department Committee on Graduate Students

Accepted by:______________________
Margaret Andrews,
Executive Director of Masters Program
Sloan School of Management

Enabling Technology for Mobile Computing

by

Vivek Phanse

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

ABSTRACT

A strategic, technical, and organizational examination of the customer enabling effort at Intel Corporation was undertaken. The context of this study was the launch of the first boxed mobile processor in the company's history. The technical component of the work consisted of the design and development of a thermal solution to complement the boxed processor. The organizational component involved an examination of the way the team was structured to execute the project. The strategic component looked at the effect of the enabling effort on the mobile computing industry at large.

Thermal solutions that met the technical specifications of the boxed mobile processor were successfully designed and prototyped. The implication of a specific thermal solution design on the level of standardization of the notebook computer platform is discussed. The conclusion is that the level of standardization in the industry is already high and is unlikely to be influenced by the boxed mobile processor. Tactics for organizational structure design are suggested for future project staffing.

Thesis Supervisor: Lionel C. Kimerling
Title: Thomas Lord Professor of Materials Science and Engineering
Thesis Supervisor: Roy Welsch
Title: Professor of Statistics and Management Science

Acknowledgements

I would like to acknowledge Jason Ku, my supervisor at Intel, for his support through my internship and beyond.

I would like to acknowledge Professors Kimerling and Welsch for their constructive ideas and criticism.

I would like to acknowledge the Leaders for Manufacturing Program for its generous support of my studies and the research conducted in this thesis.

Finally, I would like to acknowledge my family, without whom none of my success would have been possible.

TABLE OF CONTENTS

List of Figures

List of Tables

1 Chapter 1: Overview

1.1 Background

Intel Corporation is the leading supplier of microprocessors to the computing industry. Intel is also a leading supplier of microprocessors within the mobile computing market segment. Recently, the company decided to launch a boxed mobile processor for the first time. Intel typically delivers processors in bulk to its high volume Original Equipment Manufacturer (OEM) customers. However, some customers, such as smaller resellers and retail outlets, receive individual processors "in a box". To launch this boxed product, Intel must “enable” the notebook computer value chain so that reseller customers can mate processors with bare-bones notebooks and assemble them into a working computer system.

In mobile computing devices, technical issues prevent the simple "plug and play" installation of processors; hence, special customer enabling is required. The most important technical challenge is the power that is dissipated from the processor in the form of heat. For a microprocessor to function correctly, a device known generically as a Thermal Solution must be attached to the processor to dissipate this thermal energy. The thermal solution is constrained by the form factor, namely the size and shape, of the notebook computer. The confinement imposed by the typical notebook computer creates serious challenges for thermal solution designers. In this thesis, the design and development of a thermal solution for the boxed mobile processor are discussed.

The decision on how to enable the boxed processor thermal solution is a complex one, involving both business and technical issues. The choices include whether to sell the thermal solution separately or with the processor and whether to design a specific thermal solution or to create a reference design. Technical limitations and goals favor certain designs while customer and vendor interests favor other features. This thesis discusses Intel’s decision regarding how to enable the thermal solution for the boxed mobile processor. Further, enabling activities are viewed as a cost of doing business; yet, no way to value these activities has been articulated. This thesis examines the general concept of enabling and how one might value such activities.

Finally, the enabling effort for the boxed mobile processor has implications about standardization in the notebook computer industry. By developing a specific design for a thermal solution, Intel is specifying a significant area around the microprocessor. Given that a single thermal solution will fit into several different computer systems, a de facto standard is created. Unlike in the desktop market, where Intel developed motherboards and standards, the company has not participated in the design or manufacture of notebook computer motherboards. This thesis investigates Intel's strategy and the implications of the boxed mobile processor on industry standardization.

1.2 Approach and Scope

The boxed mobile processor project involves two phases: (1) An early launch of the product for use in existing system designs (2) Development of a Design Guide to influence system designs for the future. This thesis examines the technical and business issues associated with Phase 2: the Design Guide.

From a technical perspective, the focus of the Design Guide was the development of a thermal solution designed for the boxed mobile processor. Using technical expertise present in the thermal-mechanical enabling group and the Platform Architecture groups at Intel, designs were conceived, prototyped, and tested.

From a business perspective, the project involved determining the enabling model to follow for the boxed mobile processor. Intel has certain enabling models in place for its desktop products and others in place for mobile products. The boxed mobile processor incorporates elements from both enabling models and as well as features not typically found in either. The larger question of valuing enabling is also explored.

1.3 Project Goals

The initial project goals were as follows:

- Launch of the boxed mobile processor
- Decision on which processor/package to use in early launch
- Technology to enable customers for current and future thermal/mechanical solutions associated with a boxed processor
- Published Design Guide for future systems supporting the boxed mobile processor

Because of the evolving nature of the technology and the time constraints of the internship, the final project goals were identified as:

- Develop technology to enable customers for current and future thermal/mechanical solutions associated with a boxed processor
- Contribute to publication of Design Guide for boxed mobile processor
- Ensure appropriate organizational support is in place for Design Guide activities
- Analyze the impact of the boxed mobile processor on standardization of the mobile computing industry

2 Chapter 2: Project Setting and Background

2.1 Company Background, Position, and Outlook

Intel Corporation is the leading producer of microprocessors in the world. Fittingly, the company's mission is to be "the preeminent building block supplier to the Internet industry". With 1998 revenues of $26.3 billion and profits of over $6 billion, Intel ranks among the largest and most successful semiconductor companies in the world. To promote further growth, Intel continues to expand its business into new products and markets.

The Mobile Computing Group (MCG) within Intel has a mission analogous to the corporate mission. MCG exists to "grow the mobile computing industry and be the pre-eminent IA mobile building block supplier to the Internet economy" (IA - Intel Architecture)[2]. With the proliferation of cellular phones, personal digital assistants, and related devices, the opportunities for growth in the mobile arena are impressive. However, because the average microprocessor price has been declining, MCG has sought volume growth to boost revenues.

2.2 Intel Divisions and Groups and the Boxed Mobile Processor Opportunity

One opportunity for MCG to enter a new market comes via the Channel. At Intel, the Channel refers to a portion of the personal computer value chain that is pictured in Figure 2.1

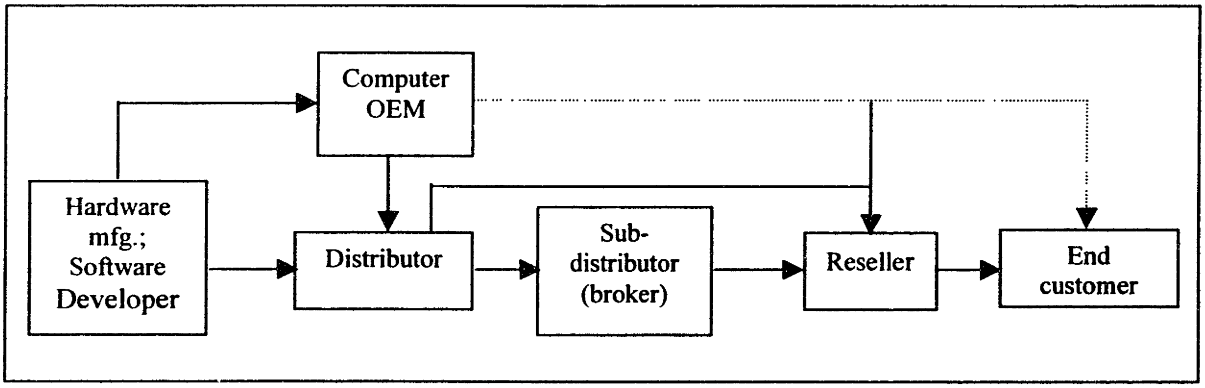

Figure 2.1: Value chain of the personal computer industry. The Channel refers to the distributor-sub distributor-reseller linkage. The dashed link reflects the emergence of the "Direct" model made popular by Dell. Intel is one of the component manufacturers at the head of the value chain.

Intel sells the majority of its processors to computer OEM customers. However, a fraction of its processor sales are made to the channel: distributors, dealers, and integrators known inside Intel as Intel Processor Dealers (IPDs), Intel Processor Integrators (IPIs), or Genuine Intel Dealers

(GIDs). Channel sales are made via individual processors "in a box". In contrast, Intel sells larger bulk volumes of processors to computer Original Equipment Manufacturers (OEMs). Boxed processors are packaged one per box while bulk orders are a minimum of 50 processors. In either case, the computer OEM or the dealer takes the processor and the other components of the computer and assembles them into a working system, branding that system with the name of the final assembler. In the mobile processor market, all of Intel's processor sales are made to OEM customers who manufacture notebooks. Though resellers may carry notebook computers, these systems have been OEM branded. Thus, there is a potential opportunity to sell mobile processors in a box to the channel, giving the dealers the opportunity to assemble and brand their own notebook computer systems.

The idea for this boxed mobile processor grew out of customer visits conducted by Intel's Reseller Products Division (RPD). RPD is the Intel business group that is chartered with growing "Intel Architecture Businesses by enabling the dealer channel to successfully compete in the Internet economy"[3]. RPD is essentially a market development organization that is responsible for Intel's sales made through the dealer channel. Marketing studies have shown that the reseller is the primary sales outlet for notebook computers targeted at small to medium businesses, especially in Asia. However, notebook computers make up only a small fraction of the business that resellers currently do. The graphs in Figure 2.2 below present some of the evidence of the market opportunity for the boxed mobile processor.

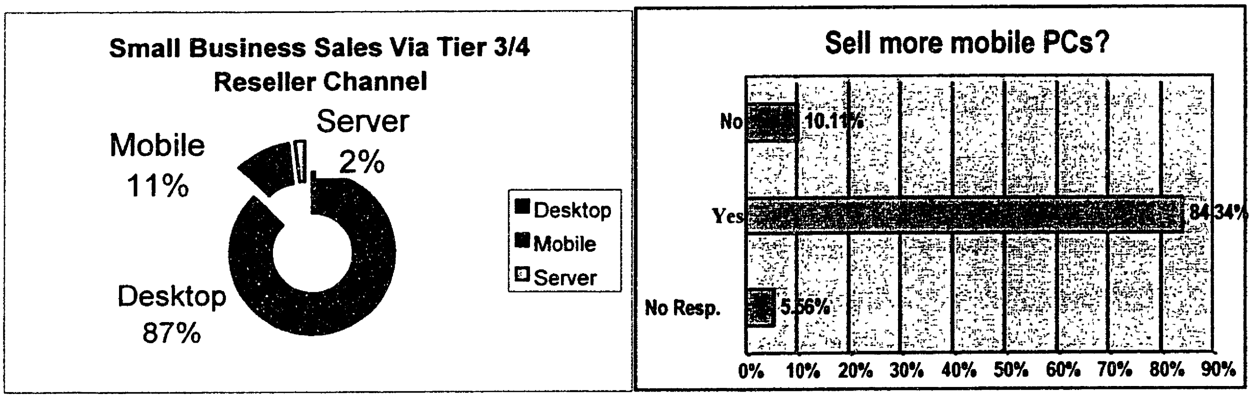

(a) (b)

Figure 2.2[4]: (a) Computer sales via the reseller channel to small businesses worldwide. (b) Survey of 1000 Genuine Intel Dealers from the Asia-Pacific region on their willingness to sell more mobile computers if support existed.

The boxed mobile processor is especially targeted at the Configure to Order/Build to Order market. The physical concept is illustrated in Figure 2.3.

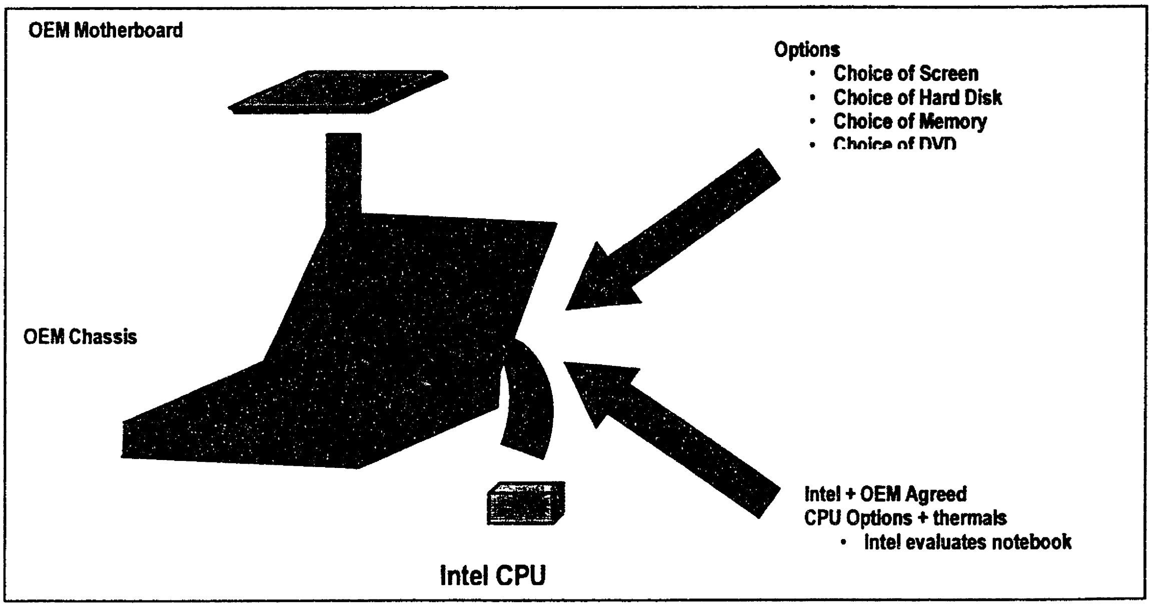

Figure 2.3: Illustration of the boxed mobile processor program. The OEM chassis + motherboard is termed a "Whitebox" and the open feature set of the computer serves the Configure to Order/Build to Order (CTO/BTO) market.

The notebook computer that is mated with the boxed mobile processor is called a Whitebox. A Whitebox is defined as a computer built by a manufacturer with everything except the hard drive, processor and memory. These options can be configured by a dealer/integrator as desired by the end user, hence serving the Configure to Order/Build to Order (CTO/BTO) market. One could also imagine that the screen could be a CTO option for notebook computers.

The Mobile Whitebox project capitalizes on the business opportunity for the dealer channel as follows[5]. The dealers are small stores, typically a five to fifty-person operation, that provide a high level of service to local customers. Small to medium businesses, which are not large enough for a major OEM to support, go to these dealers because they are local and able to provide immediate support if needed. The dealers, by having a Whitebox system, are able to expand their product line by offering more notebook computers. These dealers can also use Whitebox systems to develop their own brand name. Further, because the dealers are able to configure the Whitebox as needed, they can offer a wider variety of products at more flexible prices, creating the potential for higher margins. The OEM computer makers who participate as suppliers of whitebox notebook computers are able to minimize their inventory risk by reducing the number of fully configured systems in their warehouses. Intel, by enabling the Whitebox industry, will take certain steps to assure the quality of Whitebox notebook systems in the channel, adding to the incentive. The benefit to Intel is an expansion of the total available market for mobile processors.

2.3 Challenges of the Boxed Mobile Processor

There are several reasons why there is currently no boxed processor or whitebox system available to the channel. First, the lack of standardization in the mobile computer industry has historically been a barrier. Very few dealers have the technical expertise needed to assemble a notebook system, and with each OEM building a somewhat different computer, the problem is amplified. Having system-specific hard drives, CD-ROMs, and memory further complicates the situation. Second, many notebook computer systems use microprocessors that are surface mounted to the motherboard. This means that the processors are permanently soldered to the board, effectively eliminating the need or use of a boxed mobile processor. Another problem facing a boxed processor has to do with dissipating the heat generated by the processor. Heat dissipation requires the application of an item called a Thermal Solution. The thermal solution is usually a complex and fragile piece of equipment that requires special care to install and is unique to the system design. Again, the dealer channel simply does not have the resources or expertise to deal with this kind of complexity.

Fortunately, the industry is evolving, and this has increased the feasibility of a boxed mobile processor. First, as the average end-user might notice, most every notebook computer has certain common features: the keyboard is nearly standard, the displays are similar, most hard and floppy drives are now removable, and the PCMCIA card has become standard. The size of notebooks, while shrinking, has evolved into a few noticeable categories: 1) the full-feature desktop replacement, which has all of the features of a desktop 2) the mini-note, which is very small and mobile 3) the "thin and light", which balances between the features of the desktop replacement and the very small system. These trends have led CTO component suppliers (hard drives, memory, and floppy drives) to develop products that work in multiple brands of notebook computer.

Intel also recently introduced a mobile processor package (package refers to the material surrounding the silicon die and the scheme for connecting to the motherboard) that encourages the CTO market, the Micro-Pin Grid Array (Micro-PGA) package. The key characteristic about this processor package is that it is socketable (see figure 2.4 below). This means that that the processor can be inserted and removed from a socket that is mounted to the main board; it is held in the socket by a mechanical force instead of being permanently soldered to the board. The ease of installation and replacement of a socketed processor opens the door for a boxed product.

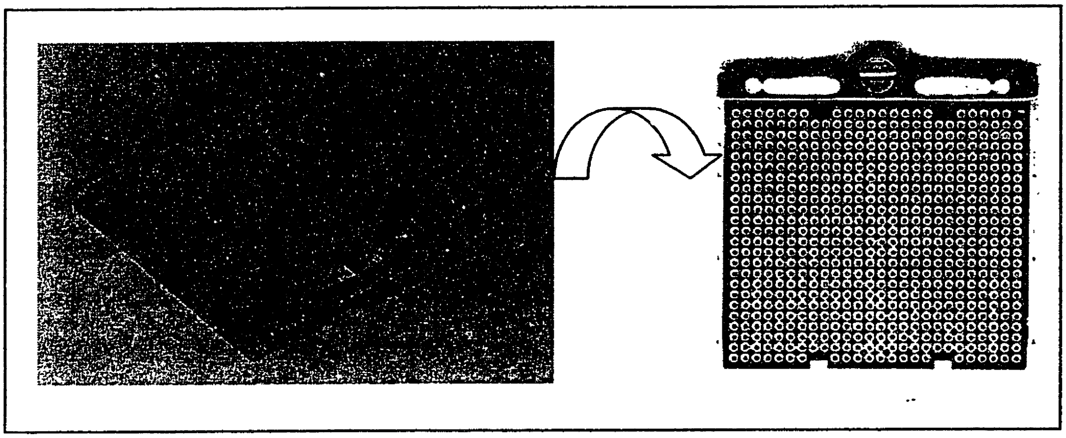

Figure 2.4: Micro-PGA processor and socket. The processor is easily inserted and removed from the socket, enabling the CTO/BTO dealer channel.

Unfortunately, the most serious technical challenge facing the boxed processor is thermal power management. The difficulty is twofold: continuously increasing processor performance has tended to increase the processor power; this is complicated by the fact that each notebook has its own way of cooling the processor, a unique Thermal Solution. The problem is illustrated well by an example.

Consider what a typical dealer would have to do to sell his own Whitebox notebook computer. It would be simple to obtain each of the parts: processors, whitebox notebooks, and CTO/BTO components. Putting them together is the challenge. In each whitebox notebook, installing the hard drive, the floppy or CD-ROM, and the memory is relatively straightforward. On the other hand, each notebook model has a unique thermal solution that might be attached to the processor in a slightly different way. This is complicated by the fact that the processor might be in different places in different models. Add in complexities like thermal interface materials and the fragility of thermal solutions, and the problem is magnified. The way that Intel went about addressing each of the challenges of the boxed mobile processor is addressed in the following chapters.

Chapters 3 and 4 will discuss how the thermal power issue is being resolved to enable the market for the boxed mobile processor. This includes enabling activities, design, and development of a thermal solution. Chapter 5 will discuss larger industry-wide implications of enabling this thermal solution design and Chapter 6 will reflect on organizational issues associated with the Whitebox project.

3 Chapter 3: Design of the Mobile Whitebox Thermal Solution

3.1 General Problem

One of the major challenges in mobile computing today is dissipating the heat emitted from the various components that make up a portable system. It is important to keep the electronic components relatively cool to maintain their reliability. For example, Intel processors have a temperature specification over which the reliability of the component is compromised. Yet, the largest generator of heat in a notebook system is the central processing unit, and as more and more transistors are added to the die (the core silicon of the processor), power levels continue to rise. Figure 3.1 illustrates the trends in power dissipation for mobile computing.

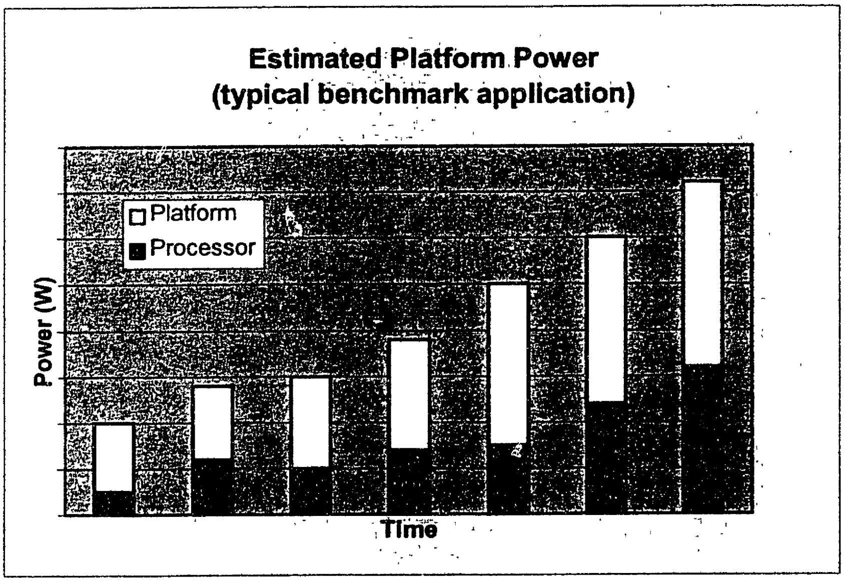

Figure 3.1[6]: Power dissipation in mobile computers over time. Platform power refers to components such as chipsets and graphics controllers, but does not include the display power.

The changing size and shape of notebook computers further complicates this situation. The trend toward thinner and lighter systems makes cooling even more difficult. Currently, most thermal dissipation takes place via air cooling, resulting in heat transfer away from the CPU to the outside of the system. Today, fans that were once vertical, with standard inlets and outlets, are changing to horizontal "flat" fans that take air in from the top and blow air out the side. Unfortunately, current flat fan performance typically lags behind that of vertical fans. Further, the physical confinement of a small computer system makes airflow resistance a serious problem. The net result is that thermal management in mobile computers has become increasingly difficult.

In response to the thermal challenge in mobile computing, Intel has invested substantial resources to enable customers to cool processors. The Mobile Mechanical Applications group is an Intel organization that helps major OEM customers to develop and improve thermal solutions in mobile computers. This group supports the customers' efforts to develop Thermal Solutions.

3.2 Design Goals

To meet the needs of the Configure to Order/Build to Order (CTO/BTO) dealer channel, Intel decided to develop a thermal solution to complement the boxed mobile processor. The thermal solution described is "processor-centric", which means that its purpose is to keep the Intel microprocessor cool with no guarantees about the other platform components. To further satisfy the needs of the channel and the OEMs who will serve the channel, the design team identified several ideal attributes of this thermal solution. Table 3.2 lists and explains these targets:

Design Goal	Purpose
Easy to integrate	Simplicity. Speed of installation and removal.
Standard solution	Reduces need for training. Interchangeable from one system to the next.
High power dissipation	Robustness: Technology forecasts plus some headroom in case installation is imperfect
Shipping model	Flexible enough to allow shipping thermal solution with OEM system or with the processor
Symmetric	Allows installation on either side of notebook. Increased design flexibility.

Table 3.2: List of the ideal characteristics of a thermal solution that would complement the boxed mobile processor, as identified by the boxed mobile processor engineering team.

3.3 Design Constraints

To meet the Design Goals described in Table 3.2, certain system-level constraints were identified. These include the thickness of the base of the notebook computer (the part on which the keyboard sits), for which a minimum thickness of 30mm was specified. This implies a limit on the thermal solution z-height. Second, at least a 2-spindle system design is required. A spindle is defined as a device with a spinning axis such as a hard drive, floppy drive, CD-ROM, or DVD-ROM. This requirement gives a rough x-y size range for the system based on typical sizes of spindle devices. Combined, these constraints give the design team an idea of the form factor of the notebook that the thermal solution will support.

In addition to physical size constraints, certain processor area constraints may need to be defined. For example, air intake and exhaust ports for the thermal solution may need to be specified. Also, ducting for the fan associated with the thermal solution may need to be defined. Further, the processor location may need to be specified and/or the motherboard may need to be cut away. These constraints are additions that would ensure that the thermal solution achieved the performance criteria identified above.

3.4 Principles of Thermal Cooling

Before beginning a discussion of the specific thermal solution design for the boxed processor, it is informative to discuss the general problem of thermal management in notebook computers to see how the design targets were chosen. For the present discussion, the assumption is that the microprocessor is the only device in the computer that needs to be cooled. Then, if one imagines that the microprocessor is the only heat source in the system and no thermal solution is present, some heat dissipation will still occur passively via conduction, convection, and radiation. Figure 3.3 illustrates.

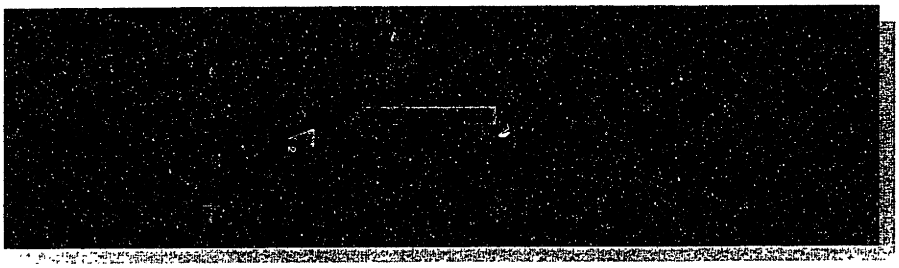

Figure 3.3[7]: Heat transfer mechanisms inside a notebook computer

Q1	Conduction through printed circuit board
Q2	Convection from component
Q3	Convection from board
Q4	Radiation from component
Q5	Radiation from board

By making some assumptions about radiative and convective cooling, one can show that there is a limit to passive heat dissipation (dissipation that would occur without a thermal solution) in a notebook computer of a given size[8]. For this purpose, only the notebook "base" is considered, i.e. the part of the notebook that contains all of the functional components of the computer except the display.

3.4.1 Radiation[9]

Radiation can be defined as the energy emitted by matter at a finite temperature. To simplify the analysis of radiative cooling, one assumes that the notebook computer is a small surface that is surrounded by isothermal air. Further, the notebook base is considered as a surface with five sides radiating heat (Figure 3.4).

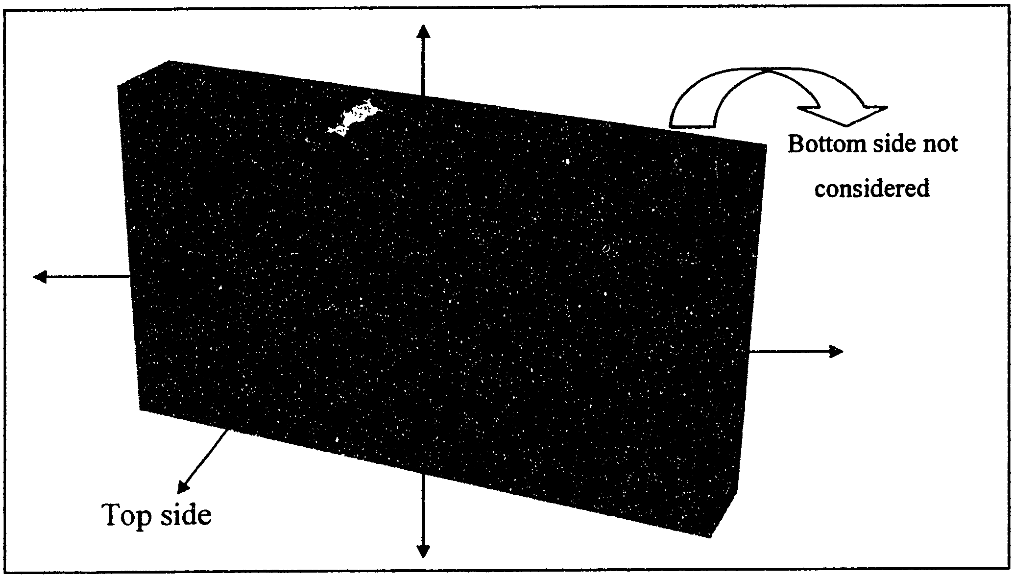

Figure 3.4: The notebook computer base is assumed to be a rectangular block with arrows showing five sides exposed to emit heat. The sixth side is assumed to be the bottom of the computer, flat against a surface such as a desk, hence unable to radiate.

The governing equation for the radiative heat transfer from the box is given by:

$$q_{rad} = \varepsilon\sigma(T_{surf}^4 - T_{surround}^4) \qquad \text{(E1)}$$

where ε is the emissivity of the surface and σ is the Stefan-Boltzmann constant. The rate is given in terms of energy per unit area.

3.4.2 Convection[9]

Convection can be defined as heat transfer between a surface and the fluid medium moving over the surface. Since we are only considering passive cooling of a box, we only consider the effects of free, or natural, convection. Convective heat transfer is governed by an equation of the form:

$$q = h(T_s - T_\infty) \tag{E2}$$

where h is the convection heat transfer coefficient that is influenced by the geometry and the fluid thermodynamics. To describe h in the case of a notebook computer base, we make the following assumptions:

- A five sided rectangular block
- The large exposed side is assumed to be a horizontal surface
- The smaller exposed sides are assumed to be vertical surfaces
- The air is quiescent around the box

Using the above assumptions and some empirical correlations, we can compute a combined h to account for all of the surfaces. The form of these calculations is as follows:

$$h = \frac{Nu_L k}{L} \tag{E3}$$

where k is the thermal conductivity of the medium, L is the characteristic length, and Nu is the Nusselt number, which provides a measure of the convective heat transfer at the surface. The Nusselt number differs depending on the geometry and a slightly different Nusselt number was used for the horizontal and vertical surfaces of the box.

3.4.3 Limits on passive cooling[8]

Making the above assumptions, and combining the results from radiative and convective cooling, the limit to the passive cooling ability of the notebook base can be obtained. Figure 3.5 illustrates.

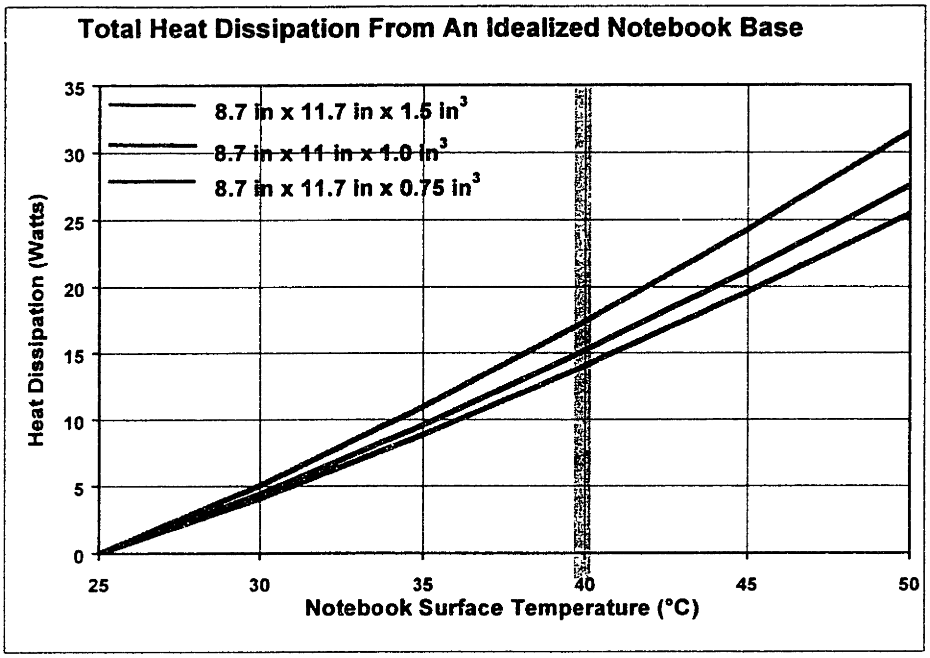

Figure 3.5: Estimated limits on passive cooling by a notebook computer base in a 25°C ambient. The shaded line marks the limit if one assumes that the notebook skin should be no hotter than 15°C above ambient. The limit is approximately 15-17W depending on the size of the base.

Given the platform powers and processor powers described above, this leaves a substantial amount of heat that needs to be dissipated via an active thermal solution (depending on the processor power).

3.4.4 Thermal Management Technology and Metrics

To meet the challenge of dissipating these kinds of powers in a notebook computer environment, the main technology employed is the Remote Heat Exchanger (RHE). Figure 3.6 illustrates the general concept of an RHE.

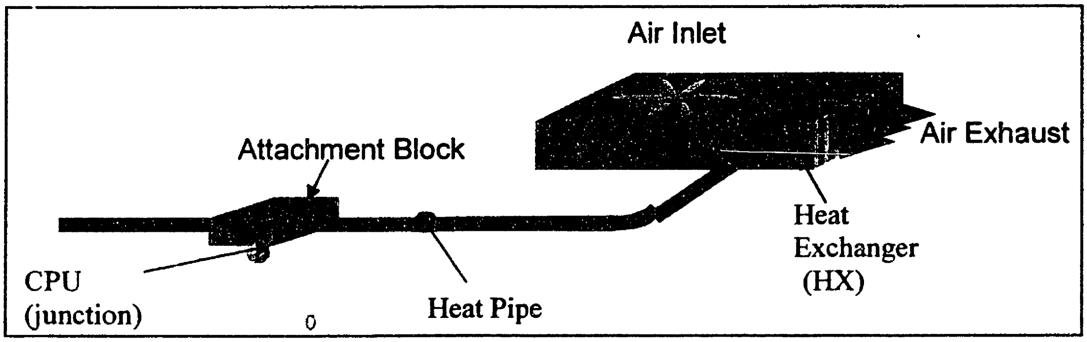

Figure 3.6[4]: A remote heat exchanger in general form consists of: 1) CPU junction block and interface 2) heat pipe 3) heat exchanger and fan.

The purpose of the RHE is to take the heat from the microprocessor to a heat exchanger located relatively far away from the CPU where it can be efficiently removed from the notebook computer via airflow over the heat exchanger. An RHE has potential to be more efficient than simple airflow over the processor because relatively cool air is taken from the outside and passed rapidly through the heat exchanger. If one tried to pass air directly over the processor, the limited surface area and the location of the processor would likely result in poor heat transfer characteristics. By keeping the RHE to die attachment relatively thin, the heat pipe narrow, using a flat fan and a good heat exchanger design, the RHE can be kept fairly small, facilitating its use in mobile computing applications. When designing an RHE, a few general principles should be followed:

- Keep distances short
- Minimize the number of interfaces and loss at interfaces
- Eliminate airflow restrictions

The purpose of these principles is to minimize the thermal resistance between the processor die and the ambient air. Thermal resistance is empirically defines as follows:

$$\Theta_{j-a} = \frac{T_j - T_a}{P} \qquad \text{(E4)}$$

where θ is the thermal resistance, T_j is the temperature of the die, T_a is the ambient temperature, and P is the power being emitted from the die. The units of θ are °C/Watt. Using this kind of definition for thermal resistance, one can represent the die with remote heat exchanger as a system of resistors: the die being the power source and each section of the remote heat exchanger being a resistor. Figure 3.7 illustrates.

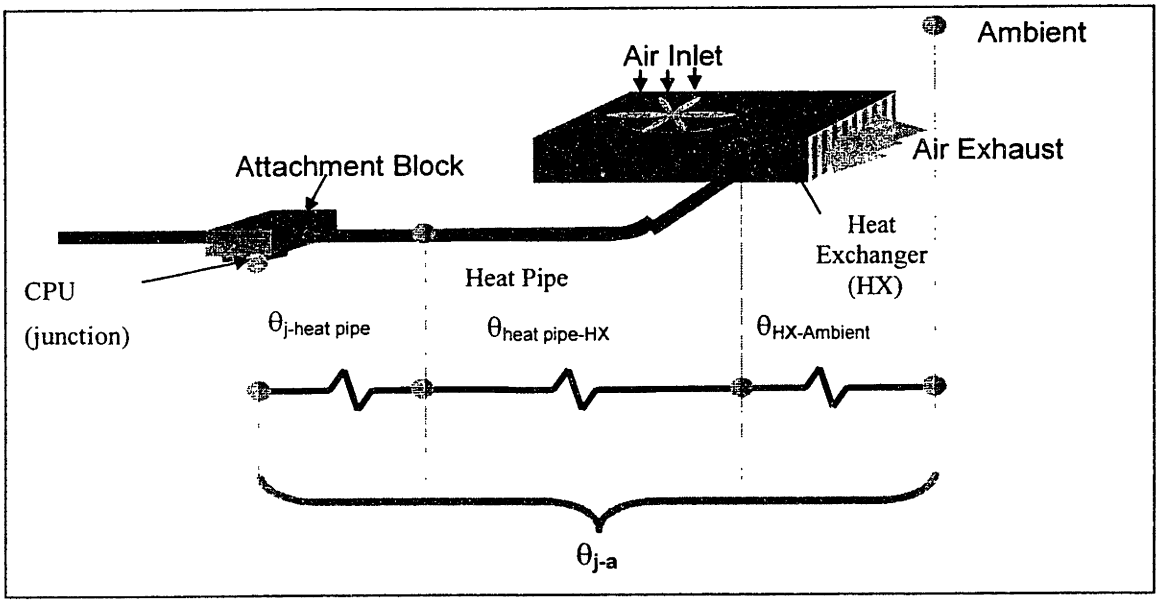

Figure 3.7: RHE with thermal resistance definition mapped onto the diagram.

3.4.5 T_{sys}: Definition and Use

At this point, one additional item must be added to the analysis: the heat contributed by the other elements of the notebook computer. As mentioned previously, the platform power is increasing along with the processor power, and this heat must be removed from the system as well. Platform components such as hard drives, floppy drives, graphics controllers, and batteries will not be actively cooled using an RHE because the box itself is capable of dissipating some heat. However, platform-generated heat will interact with the processor and RHE setup, increasing the internal ambient temperature of the system and limiting the effectiveness of the RHE. The effect of the non-processor components is captured in a variable called T_{sys}. When one incorporates the effect of T_{sys} into the thermal resistance analysis, one finds that the thermal resistance effectively increases. Adding T_{sys} to equation E4 captures this:

$$\Theta_{j-a} = \frac{T_j - T_a - T_{sys}}{P} \qquad \text{(E5)}$$

T_{sys} has been empirically determined to vary between 10-15°C; for a worst case, conservative approach, one would choose a T_{sys} = 15°C[4]. At this point, the challenge of designing a thermal solution for the boxed processor is better defined. For example, a 10W capable thermal solution, assuming a T_{sys} of 15 °C, a maximum die temperature of 100 °C, and an ambient temperature of 35°C will require a total thermal resistance of approximately 5 °C/Watt.

3.5 Components of the Remote Heat Exchanger

As identified earlier, there are three key elements to a remote heat exchanger the die interface, the heat pipe, and the heat exchanger. These three elements create three identifiable thermal resistances: $\theta_{j\text{-}hp}$ associated with the die interface, $\theta_{hp\text{-}hx}$ associated with the heat pipe, and $\theta_{hx\text{-}amb}$ associated with the heat exchanger. These resistances were visually identified in Figure 3.7.

3.5.1 Interface Materials

The interface material makes contact with the backside of the die and the block portion of the RHE. The ideal interface material is one that is highly thermally conductive, yet pliable so that it can absorb small uncertainties in the bond lines. For example, Intel specifies the height of the die above the socket within a certain range (e.g. 7mm +/- 0.01). Thus, a thermal solution that relies on the die height to be exactly 7.0mm would suffer performance hits based on the distribution of the die height. Thermal interface materials are usually able to absorb small variations in height and generally come in categories described in Table 3.8.

Thermal interface material	Physical description	Thermal resistance	PROs	CONs
Grease	Viscous fluid	Low	Best thermal performance	Dispensing problem, bond line control (repeatability)
Phase change	Soft solid that changes to liquid	Low to Medium	Good thermal performance, easy to handle one time	Reworkability
Elastomer	"Rubbery" solid	High	Easy to handle, reworkable	Poor thermal performance

Table 3.8: Comparison of three main types of interface materials

Table 3.8 illustrates a classic engineering tradeoff, an ease of use versus performance question. While grease is the most thermally effective material, it is also the most difficult to use in an assembly process. For the boxed mobile processor application, one initially might expect that the low level of expertise that dealers possess would prohibit use of grease or phase change, but Intel has previously shipped boxed (desktop) processors with a vial of grease and a thermal solution. Another important question is how to handle returns and upgrades. The boxed mobile processor clearly opens the door for upgrading just the CPU; this requires an interface material that is reusable, reworkable, or easily replaceable.

3.5.2 Heat pipes

Heat pipes are literally what their name indicates: they are ideal conduits in the shape of a pipe that transport heat from one area to another. The operation of a heat pipe is illustrated in Figure 3.9:

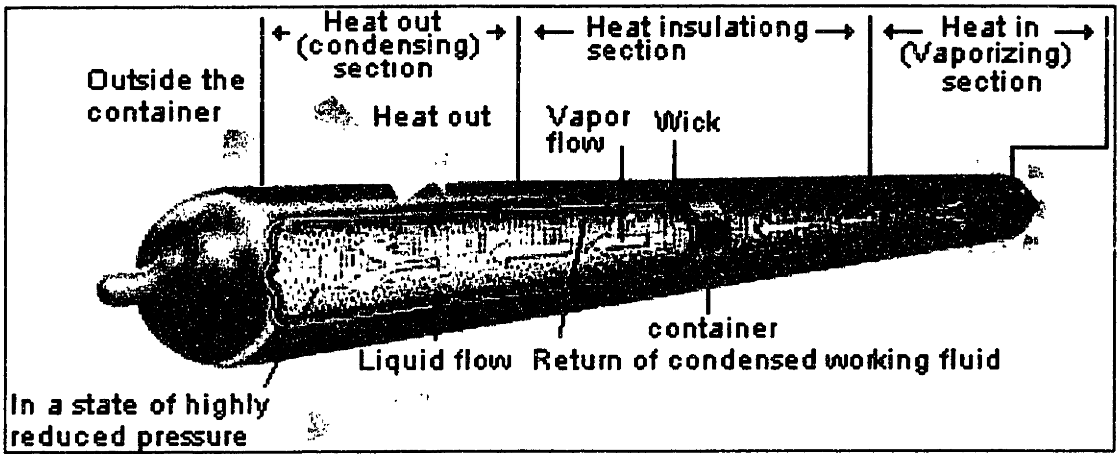

Figure 3.9[10]: Cross section and operation of a heat pipe

The essential operation of a heat pipe is 1) the fluid in the pipe (often water) vaporized 2) vapor moves to the cooler section 3) the vapor condenses 4) the fluid flows back to the hot section via capillary action to begin the process again. While heat pipes are extremely efficient, boasting thermal resistances in the 0.2-0.5 °C/W range, they are also fragile and expensive. Thus, while in some sense an ideal RHE would be a single piece heat pipe in the appropriate shape and size, heat pipe reliability and cost keep this from occurring.

3.5.3 Heat Exchangers

There is a large volume of literature devoted to the design of heat exchangers[9, 10,11]. Rather than go into exhaustive detail about design of a heat exchanger, I will simply mention that the goal of the heat exchanger is to expose as much area as possible to the flow generated by the fan. This goal must be balanced by having the fins on the heat exchanger parallel to the flow directions so as not to impede the performance of the fan. The importance of minimizing flow resistance will be discussed later. Suffice to say that the thermal solution under consideration is expected to use an extruded fin heat sink setup with a fan that will look something like what is pictured in Figure 3.10.

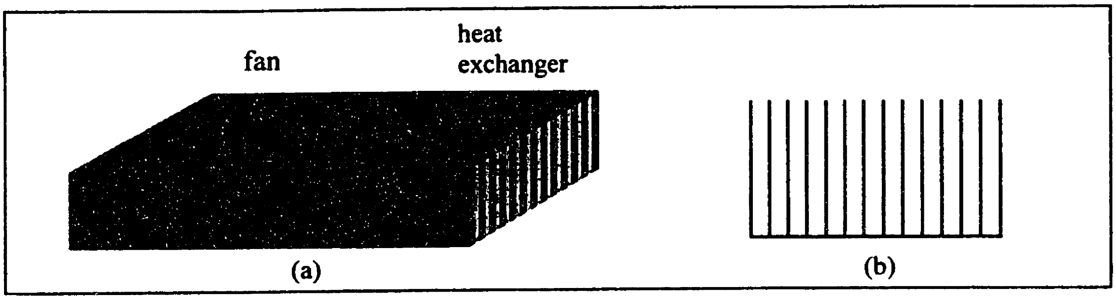

(a) (b)

Figure 3.10: (a) Fan next to heat exchanger (b) Side view of extruded flat fin heat exchanger

3.5.4 Flow Considerations

An additional factor that has been left out up to this point is the flow of air generated by the fan. Active cooling in notebook computers is different from passive cooling (radiation and natural convection) in that active cooling involves using a fan. The performance of the fan is absolutely crucial to the success of the thermal solution. To show just how important fan performance is, consider the following example of how to approximate the cooling capacity of an RHE.

Given the following (see Figure 3.7 above):

Thermal resistances

$\Theta_{\text{j-heat pipe}} = 0.5$ °C/W

$\Theta_{\text{heat pipe-heat exchanger}} = 0.5$ °C/W

$T_j = 100°C$; $T_{air}=35°C$; $T_{sys}=15°C$

Processor power : 15W

1) Calculate the temperatures at the block, heat pipe, and heat exchanger.
2) The processor temperature is "derated" by T_{sys}:

$$T_j = 85°C$$

3) The temperature at the heat pipe is calculated via application of equation E4:

$$\theta_{j-hp} = \frac{T_j - T_{hp}}{P} \Rightarrow 0.5 = \frac{85 - T_{hp}}{15} \Rightarrow T_{hp} = 77.5°C$$

4) Likewise, the temperature at the heat exchanger can be computed to be $T_{hx} = 70$ °C.
5) Then if one assumes that the air flowing through the heat exchanger will heat up approximately half of the difference between the heat exchanger and ambient:

$$dT_{hx-air} = T_{hx} - T_{air} \Rightarrow \frac{1}{2} dT = 17.5°C$$

6) Then, the cooling capacity of the heat exchanger can be estimated by a heat/mass transfer analogy:

$$q = \dot{m} C_p \Delta T = \dot{V} \rho C_p \Delta T$$

7) Making the ideal gas assumption, using C_p of air at the correct temperature, and assuming the fan has a volumetric air flow rate of 2 cubic feet per minute:

$$q = 19.1W$$

Notice that the heat dissipation is directly proportional to the flow rate of the fan.

3.6 First Concept

Based on the concepts described above, the first concept for the Whitebox thermal solution was developed. It is pictured in Figure 3.11:

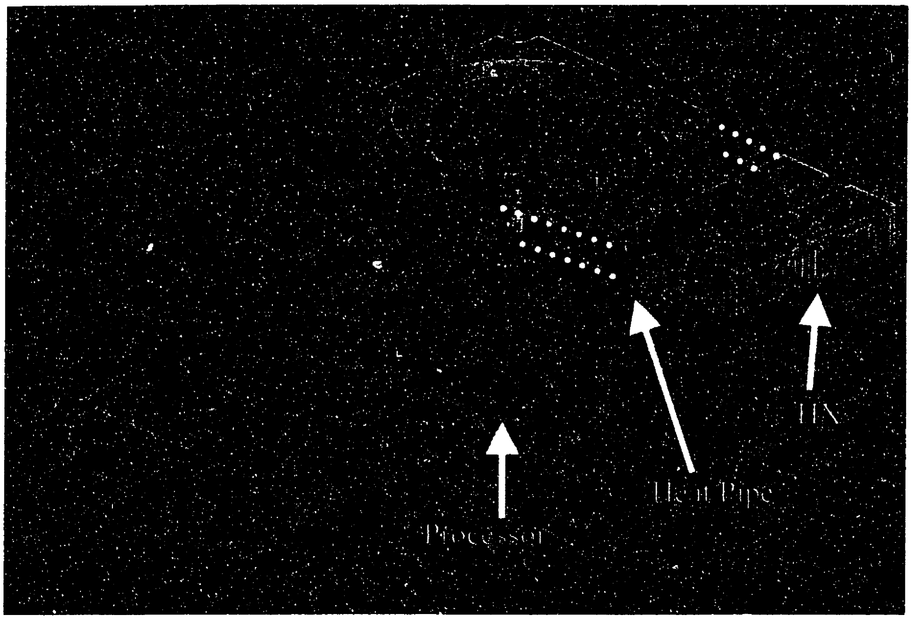

Figure 3.11[13]: First concept of Whitebox Thermal Solution.

As the picture illustrates, the concept is nothing more than a standard RHE that follows some basic design principles. The unique aspects of this design are the method of attaching the block and heat pipe to the processor and the direct heat pipe contact to the die. With regard to the attach method, the use of a spring-loaded clip was proposed. This design feature was suggested as a way to make the thermal interface extremely simple and repeatable, such that the dealer installing the thermal solution would have as simple a task as possible. The direct die contact to the heat pipe is a way of reducing thermal resistance. Some early technical validation on the above design yielded the results discussed in Section 3.6.1. The main task of these early tests was to verify the feasibility of developing a thermal solution capable of meeting the thermal demands of the next generation processor by testing some interface materials and looking at some heat exchanger designs.

3.6.1 *First Concept Data*

Ultimately, the goal of any thermal solution is to keep the die (junction) as cool as possible. However, when designing a thermal solution, data must be collected to determine the thermal bottlenecks in the solution. Thus, the data collected is temperature data at various points along the thermal solution. Using equation E4 above, one can then determine the thermal resistance between points along the thermal solution, which gives an idea of where to try to make improvements.

Prototype testing was conducted in an open-air environment, which does limit its usefulness for overall thermal solution evaluation. The raw data and thermal run charts are given in Appendix 1; key results are presented in Table 3.12:

Thermal interface	**Θj-a (°C/W)**	**Θj-hp (°C/W)**	**Theoretical max power dissipation* (W)**
Grease 1	3.22	1.23	17.1
Grease 2	3.07	0.99	17.9
Phase Change 1	4.21	2.26	13.1
Phase Change 2	3.78	1.55	14.6

Table 3.12: Early testing results of different interface materials. * Assumes a Tsys of 10°C, maximum die temperature of 100°C, and ambient temperature of 35°C.

The results above illustrate two things: 1) the difficulty of cooling high-performance processors 2) performance of phase change material is somewhat worse than grease. For a first pass design, these results were encouraging. With factors to change such as the type of fan and some improvements in the heat exchanger, even higher power dissipation is possible.

3.6.2 *First Concept Modifications*

After some internal consultation among engineering teams, changes were made to the first thermal solution to try to make it even simpler for the dealer channel to work with. This second concept is illustrated in Figure 3.12, along with how such a thermal solution would fit into a working notebook system.

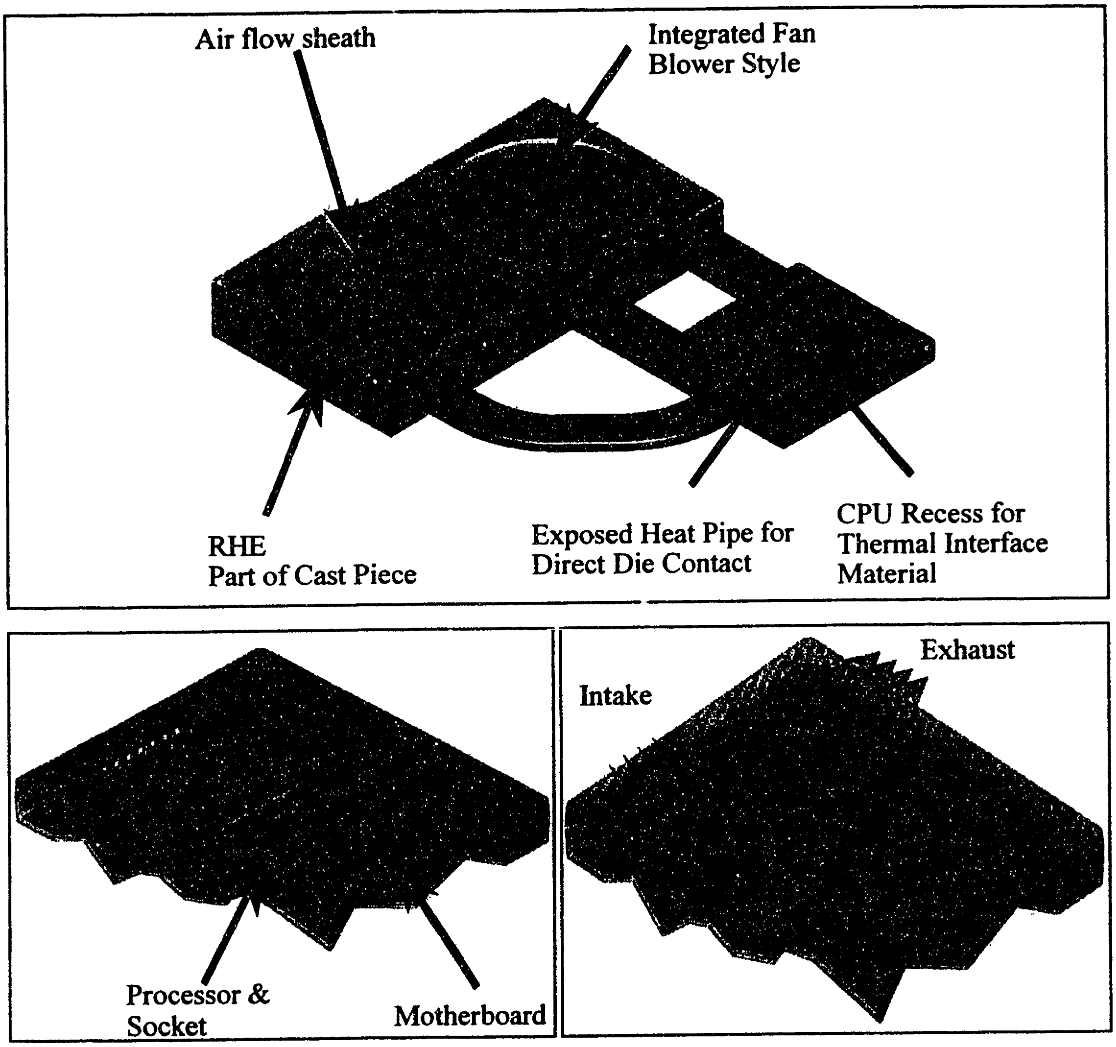

Figure 3.13[13]: Improved first concept of Whitebox thermal solution. Above: bottom view of thermal solution concept. Bottom: how thermal solution would fit into a corner of a notebook.

The theoretical advantages of the modified design include the one-piece solution, the direct die contact to the heat pipe, and its symmetrical nature. The single-piece design is a feature that reflects the practices of Design for Assembly methods. By eliminating multiple pieces, the time and difficulty of installation is reduced. The direct heat pipe contact to the die is a method of reducing the thermal resistance between the heat exchanger and the die. The integrated fan is close to the heat exchanger and is well-ducted, providing an ideal airflow path. Finally, the design is symmetric, which means that it can be installed in any corner of notebook computer, giving some design flexibility to the OEM with regard to motherboard layout.

3.6.3 OEM Feedback

With the new paper design and the promising initial results, a meeting was scheduled with the OEM notebook manufacturers to discuss the thermal solution feasibility. The feedback received from these stakeholders was as follows:

- The design takes up too much space. Further, if located in the back corner of the notebook, it would block the display hinge.
- A detachable fan is desired.
- The direct heat pipe contact to the die is difficult from a manufacturing standpoint. It is hard to get very flat heat pipes. Heat pipes are also very fragile. The direct contact and the curve in the heat pipe are risky from a reliability point of view.
- The clip setup may not put adequate pressure on the die, as there is not enough of a lever arm to generate the appropriate force. This is also a potential problem in that clips are a risk to fail the shock and vibration testing typically done on notebook computers.
- The lack of a passive spreader plate to complement the thermal solution was a concern.

These comments were enlightening because they came from a system designers' perspective. The root causes of these concerns are twofold, design flexibility and reliability. The concern with space and the detachable fan reflect the OEM's desire to be able to do the layout of their motherboard. Further, the detachable fan gives flexibility in terms of where the air inlet and outlet are located. Essentially, the OEMs preferred greater design flexibility than the symmetric RHE could provide. The direct heat pipe attach, while theoretically ideal, is practically difficult for the reasons identified above. The clip attach is a method that OEMs claimed to have looked into and abandoned after some failures.

3.7 Early Adopter and Second Concept

Based on the technical feedback, a second thermal solution was designed and prototyped. It is pictured in Figure 3.13.

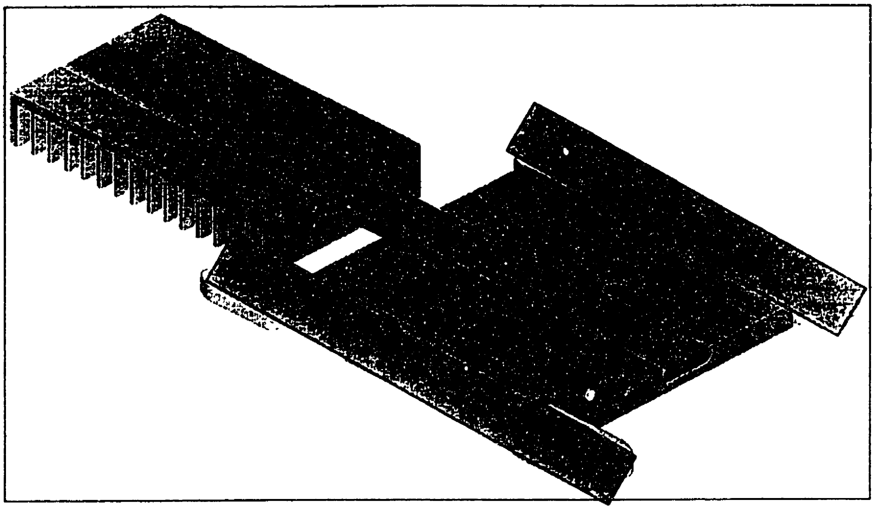

Figure 3.14[13]: Second thermal solution concept. Actual attach hardware is grayed out to protect proprietary information.

The important new characteristics of this design are a new attach mechanism, a separate fan, and having indirect die contact to the heat pipe. The new attach is proprietary; suffice to say that it remains a screwless design. The heat pipe is now no longer directly contacting the die, rather the aluminum block is solid and the heat pipe is mounted into a slot on the block. The fan is clearly separate. In reality, this design is almost a return to the very first concept with a modification to the attach method and a move away from direct heat pipe contact to the die.

Preliminary thermal test results for this setup were very encouraging (raw data in Appendix 1). Using grease as the thermal interface material and testing in an open air environment:

$\Theta_{j\text{-}hp} = 0.63°C/W$	$\Theta_{j\text{-}a} = 2.7°C/W$	Max power est. = 20.4W*

*Assumes a Tsys of 10°C, maximum die temperature of 100°C, and ambient temperature of 35°C.

3.8 Beyond the RHE: System Design Constraints

As described above, the RHE is only one of the specifications that might need to be made to meet the thermal design goals. Depending on the final RHE design, the motherboard layout, the fan, and the air inlet and outlet may need to be specified. The latest incarnation of the RHE design provides a considerable amount of flexibility with regard to motherboard layout. The RHE is symmetric and can be placed at many angles and in many locations. However, for the RHE to work well, the fan and flow path and location must be considered.

The fan choice and specification will be based on experimental data. Given the performance of the above RHE in the notebook, a fan with a minimum flow rate will be specified (see section 3.5.4 for details). Another consideration with regard to fan choice is quality and reliability. Intel will select a fan that meets quality and reliability guidelines that thermal solution parts typically must meet. Finally, where the fan is located and the resulting system resistance must be considered.

In a notebook computer or any such constrained environment, the performance of the fan will be degraded because of the physical resistance imposed by structures in the flow path.

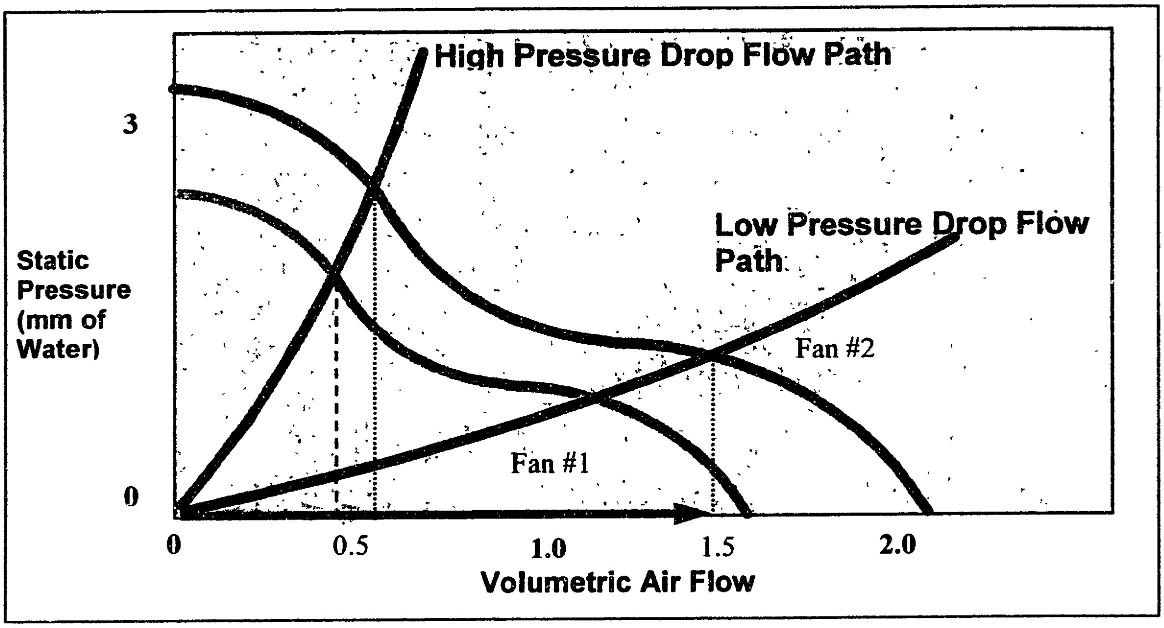

Figure 3.15[8,12]: Fan performance and system resistance curves. The graph shows two fans with different performances in two systems with different resistances.

Figure 3.14 demonstrates a couple of basic fan flow concepts. First, consider the decreasing curves representing Fan #1 and Fan #2. Where these curves intersect the x-axis is the maximum flow rating of the fan as it would operate in a pressure free (open-air) environment (Fan #2 is a higher performance version of Fan #1). If the backpressure is increased, the flow out of the fan decreases until it drops to zero where the curves intersect the y-axis. The increasing curves on the graph represent two different flow paths. The high pressure drop flow path implies a path with a lot of obstacles while the low pressure drop path indicates a well-ducted and clear path for the air. The intersection of a fan performance curve and the system resistance curve gives the fan performance as measured by the volumetric airflow. The figure clearly illustrates that a well-designed flow path is critical to getting the most out of a fan. And, as the sample calculation in section 3.5.4 demonstrates, the airflow rate is directly proportional to the cooling ability of the RHE. Empirical data suggests that a fan in a notebook computer environment operates at approximately 50-60% of its maximum flow rating.

To help clarify the fan issue, the performance of fans has been studied in a number of environments. The data collected gives an indication as to how close to the inlet or outlet an obstacle may be placed without altering the performance of the fan. One example of this is the flow rate of the fan measured as a function of the distance a plate is placed from the inlet. Considering a blower fan (top inlet, side outlet) and the setup pictured in Figure 3.15:

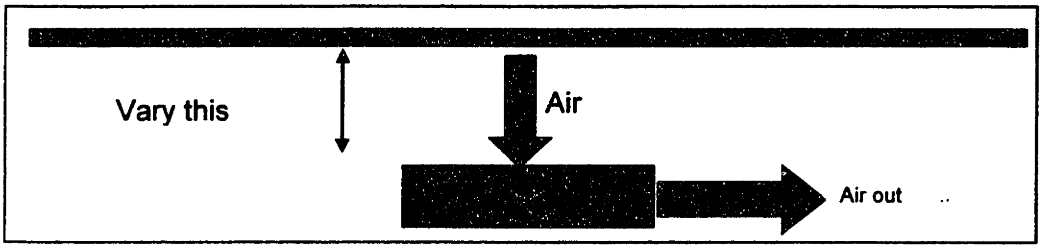

Figure 3.16: Experimental setup for testing fan performance relative to inlet obstructions.

The fan performance results obtained are pictured in Figure 3.16.

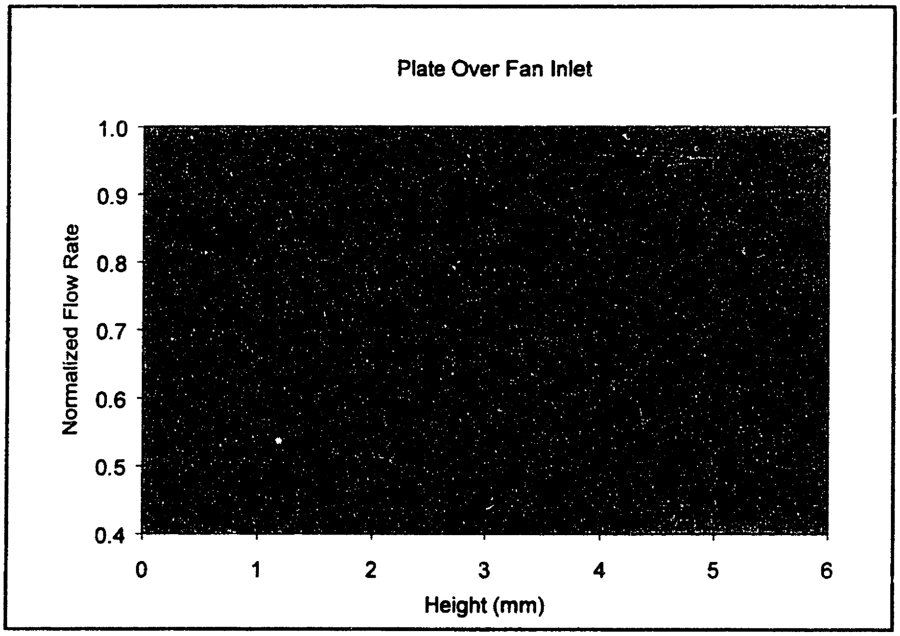

Figure 3.17[8]: Fan performance as a function of plate distance from inlet. Each curve represents a different fan. The y-axis plots the fan performance relative to open-air performance.

The conclusion of this data is that an obstruction at the inlet of a fan should be at least 3mm away; otherwise, serious degradation in fan performance will occur. This is likely to be a specification that will complement the thermal solution design. Note also that different fans react differently to a resistive inlet path. Again, this may influence the fan specification.

3.9 Continuing Challenges

The final thermal solution design for the boxed mobile processor has not been completed. Thus, some issues remain unresolved. These include air inlet and outlet specification. Several types of inlet and outlet shapes are being examined along with the spacing described above. Another key issue is the interface material and reproducibility. Because grease requires some precision in its application, the thermal solution would have to be extremely robust to accept a grease interface. If an alternate material such as elastomer or a reworkable phase change material is feasible, expect this to be the recommended choice.

3.10 Conclusions

Hopefully, the reader has gained an appreciation for the difficulties faced by the design team and the approach that developed. To aid in the development of the thermal solution, Intel engineers proposed an Early Adopter program. This program involved bringing in a specific OEM to help with the design of the thermal solution. The decision to begin this process and the larger enabling context is discussed in Chapter 4. The final specifications will be made jointly between Intel and the OEM partner. The broader implications of the thermal solution and associated specifications are discussed in Chapter 5.

4 Chapter 4: Enabling the Boxed Mobile Processor

4.1 Enabling at Intel

As a component supplier to the computing industry, Intel is in a unique position. In most industries, the component supplier produces a commodity or near-commodity item and is at the mercy of the downstream designer and integrator. For example, in the automobile industry, powers such as Ford and GM have a great deal of control over their suppliers of seats, interiors, and electronics. Typically, the downstream company is the dominant force that influences design and integration of the components. However, in the personal computer industry, it is Intel, the component supplier, that has more ability to influence the downstream product.

Intel makes the most complex and valuable portion of the computer, the microprocessor. Consequently, everything else in the computer must be built around that component. Furthermore, because of its immense success, Intel has the resources and expertise to take an active role in downstream computer design activities. Though its main business is supplying microprocessors, Intel goes so far as to have a business competency in the assembly of entire desktop and server computer systems. In parallel, Intel devotes considerable resources to "customer enabling" activities, which ensure continuous demand for its products.

“Enabling” at Intel encompasses many different activities that help Original Equipment Manufacturer (OEM) customers to design their systems. Key technical areas of enabling include thermal solutions, power delivery mechanisms, voltage regulation technology, and manufacturing enabling, among others. Several organizations within Intel exist with the sole purpose of enabling customers. These include Enabling Enclosure Design (EED), a group that designs thermal solutions for desktop, server, and workstation processors, and Enabling Technology Division (ETD), a group that looks for technologies that will serve as key enablers for future processors. The type of interaction between the various enabling groups at Intel is illustrated in Figure 4.1.

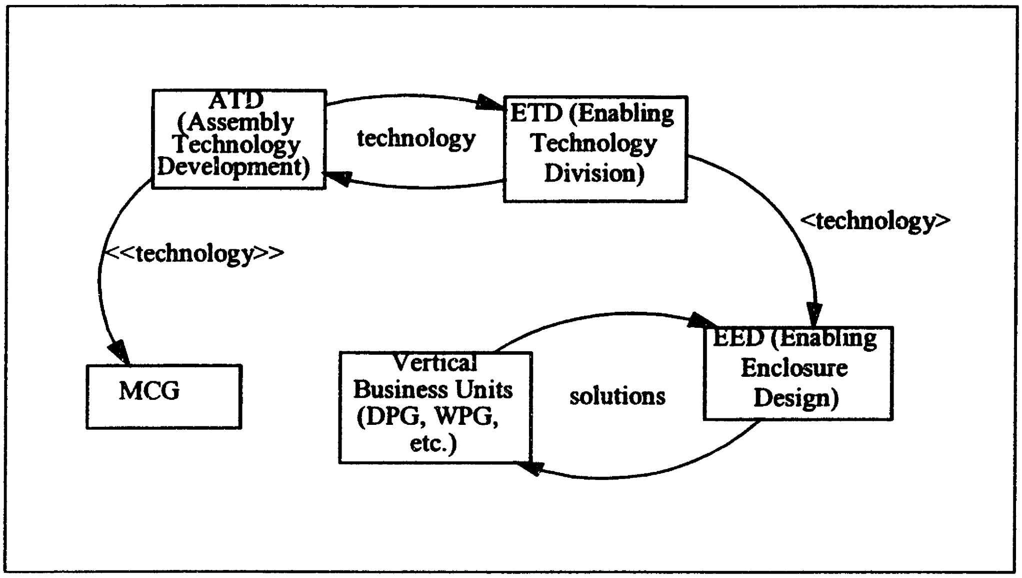

Figure 4.1: Groups at Intel involved in customer enabling activities and their information exchange. Vertical business units represent a particular market, market segment, or product family. An example of information flow would be as follows: ATD and ETD identify and develop a thermal cooling technology. They pass this information to the design group, EED. When a vertical business unit releases a new product (processor), they contract the actual design of the thermal solution to EED, which is now equipped with the latest technology and information.

The logic for customer enabling is fairly simple. From Intel's perspective, the easier it is for a customer to use an Intel product, the more likely they are to use it. Enabling promotes goodwill between Intel and customers and fosters continuous growth of the market. Recently, Advanced Micro Devices (AMD), Intel's largest processor competitor, gained a toehold in the low price market segment. Through enabling, Intel can maintain differentiated products (i.e. higher power, faster processors) that command premiums in the market. Another purpose behind enabling is to get customers ready for new products as soon as they are introduced. This is especially important for Intel, where continuous price pressure on processors means that on average, the highest margins are usually obtained at the earliest stages of product life. If customers are enabled to accept the newest processors as soon as they come out, Intel can generate more profits.

4.2 Enabling the Boxed Mobile Processor

The technical enabling plan for the boxed mobile processor consisted of a number of business and technical issues. The focus here will be on thermal solution enabling. But before discussing the thermal enabling plan for the boxed mobile processor, it is useful to mention Intel's enabling efforts for desktop boxed products. To keep the description simple, consider only the enabling activity surrounding thermal solutions at Intel. In the desktop market, Intel works with outside thermal solution suppliers to design specific thermal solutions for desktop processors. Desktop thermal solutions currently consist of a heatsink and a fan. Specifications for the heatsink and fan are created by Intel and adopted by vendors. This thermal solution design is then manufactured by outside suppliers and can be packaged in the same box as the microprocessor (for boxed processor products). The thermal solution vendor wins by being able to sell its thermal solutions to the OEM computer makers for distribution with their systems. Intel wins because it has a proven design that ensures that its product will function in the desktop environment. However, the desktop chassis is physically large and the motherboard is standardized, resulting in a simplified design process.

In the mobile arena, conditions exist that require a different enabling model. First, the notebook chassis is much smaller and space becomes a critical issue. Furthermore, the average chassis size is shrinking as people gravitate to thinner and lighter notebook designs that offer greater portability (See Figure 4.2). Second, because of the uniqueness of each chassis, motherboard standardization does not exist. These two factors mean that a single thermal solution cannot satisfy all of the existing notebook designs. Thus, instead of a specific design, Intel's Mobile Computing Group (MCG) releases "reference designs" for thermal solutions. Reference designs are paper designs that can be tailored to fit a given notebook system. An OEM can take a reference design and alter it as needed to fit inside the chassis and complement the motherboard. Intel, instead of working with one or two thermal solution vendors, enables as many as it can. Intel validates thermal solutions from various vendors via lab experiments and then publishes the results during industry-wide seminars.

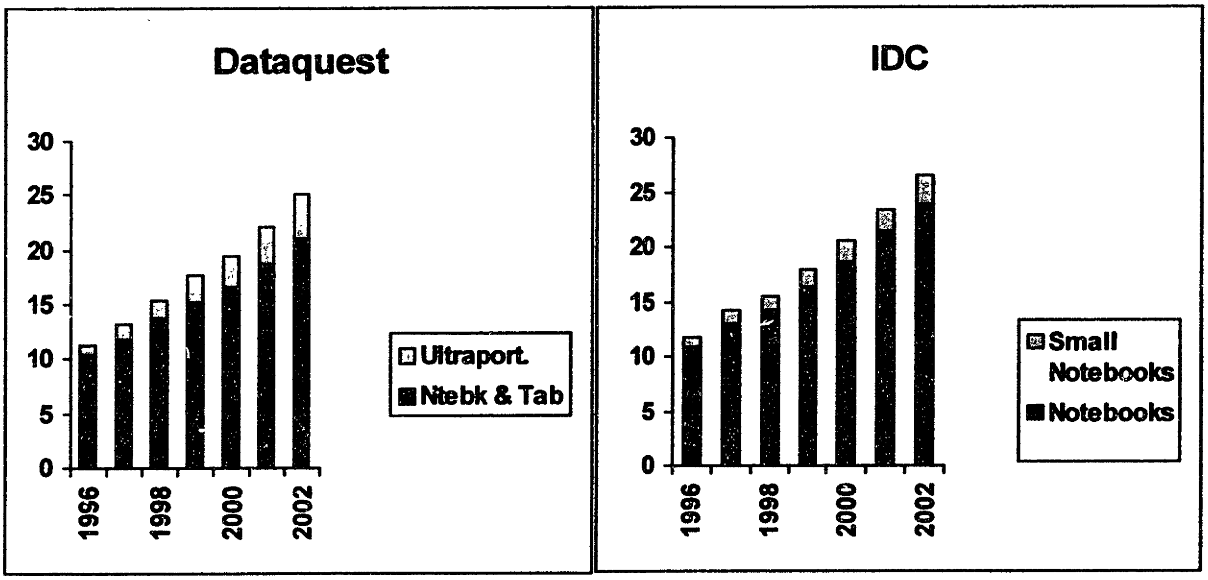

Figure 4.2[14]: Trends in notebook computer sales as predicted by market research sources. Notice that a great deal of growth takes place in the small or ultraportable segment.

An example of a reference design is the Remote Heat Exchanger (RHE), described in Chapter 3. The RHE is general enough that an OEM can make modifications to the mechanical attach mechanism, the heat pipe, or the heat exchanger, but still have a framework for how to cool the microprocessor. Intel validated (and continues to validate) RHE designs by testing interface materials, heat pipes, and prototype designs made by several thermal solution vendors. These results are then presented to OEMs at Intel-sponsored seminars, demonstrating the technical feasibility of cooling the chip. Another service that Intel provides to mobile customers is direct design support. An OEM can create a thermal solution, and if the OEM design has trouble meeting the thermal target of the processor, Intel will help the OEM redesign the thermal solution.

The specific enabling challenge of the boxed processor is the market that it is going to serve, namely the dealer channel. This is a different customer than MCG has been used to supporting, and has certain characteristics that are at odds with the usual MCG customer, the large OEM. Table 4.3 explains.

Characteristic	**Large OEM customer**	**Dealer channel customer**	
		Whitebox OEM	Dealer/integrator
Product design/development	Significant team, investment, and expertise	Small team with fewer resources	N/A
Customer support	Large, experienced organization	Small, though experienced organization	Very small, little experience supporting mobile products
Experience	Mobile business for some time	Mobile business for some time	Limited experience with mobile computing
Volumes	Very large	Large	Small
Communication method	Direct relationship with Intel	Indirect relationship with Intel	Direct relationship with Intel

Table 4.3: Comparison of the dealer channel stakeholders and the traditional MCG customers.

As Table 4.3 demonstrates, the Whitebox OEMs and the dealers do not have the resources or expertise needed to deal with the complexity of notebook computers. After some preliminary discussions with the Whitebox OEMs and the dealer channel, Intel concluded that to ensure success of the boxed mobile processor, a specific thermal solution design was needed. A specific thermal solution design for the mobile arena is clearly a departure from Intel's traditional reference design. A specific thermal solution design gives Intel the option to include the thermal solution in the same box with the processor. However, for cost and flexibility reasons, Intel decided against including a thermal solution in the box with the processor for the initial launch of the processor. The potential implications of a specific mobile thermal solution will be discussed in more detail in Chapter 5.

4.3 Partnering for Design

During the early stages of thermal solution development, a few designs were proposed by Intel and presented to the Whitebox OEMs (see Chapter 3 for details). In response to the feedback obtained from the computer OEMs, two things happened. First, Intel engineers went through several iterations of design, trying to meet the needs of the key stakeholders. The second result of the feedback was the creation of an Early Adopter program that would involve Intel working closely with one particular OEM on the thermal solution design. This OEM would then gain an advantage as the system around which the thermal solution was designed, and would have extra input to its design. Intel would gain the system design experience that the OEM possessed as well as the resources to build prototype notebook chassis and motherboards. The added benefit is a "real" system in which to test thermal solution prototypes. Additionally, other OEMs who saw the thermal solution working in a competitor's system would be more likely to trust the design. Intel also decided to partner with a specific mobile thermal solution supplier. This is the same as the strategy used for enabling boxed desktop processors. In fact, many of the tactics used to enable the boxed mobile processor came directly from models of enabling desktop processors. The idea is to use previous organizational learning and capability for a slightly new product.

4.4 Valuation of "Enabling"

The larger question encompassing this chapter is precisely how to place a value on "Enabling". This question is similar to the question of how to fundamentally value basic research. In both instances, the benefits are difficult to measure, have long-term implications, yet may have a tangible cost in the short term. Nonetheless, I propose two alternate methods for evaluating an enabling investment: one tied to customer satisfaction and speed to market, the other approach involving the use of real options.

4.4.1 Customer Satisfaction and Readiness Approach

First, consider the benefits of enabling:

- Customer satisfaction and Customer loyalty
- Increased processor sales
- Processor sales when prices are high

Now, also consider the costs of enabling:

- Engineering cost
- R&D costs
- Marketing cost
- Marginal cost to produce more processors that are sold as a result of enabling

Then of course, one would try to determine the return on investment of Enabling:

$$\text{Return} = \frac{\text{Revenue from Enabling - Cost of Enabling}}{\text{Cost of Enabling}}$$

The costs of enabling are relatively simple to estimate. Intel invests in people, technology, and marketing activities to drive the enabling effort. These costs are set in the strategic plan for the various business units and are understood to be useful to the customers. Finance groups within Intel track these costs and set budgets based on historical data and technology forecasts. The greater problem is trying to calculate the revenue generated as a result of enabling.

The immediate difficulty is trying to value customer satisfaction and loyalty. Marketing studies have revealed that satisfied customers are more loyal, and that it is less costly to maintain current customers than to obtain new ones[15]. Further studies have related customer satisfaction to market share and profitability. Rust and Zahorik have made the connection between marketing investment, customer satisfaction, and the present value of that satisfaction[16]. This link is made by relating customer satisfaction to retention and then to market share. The analysis proceeds by first identifying the factors that improve customer satisfaction. Then, one must determine the dollar effect that spending has on customer satisfaction. Finally, the link between customer satisfaction changes and customer retention must be made.

There are several methods of estimating the connection between spending and customer satisfaction[16]. For a first pass effort, one can use the judgement of managers of the firm of interest. This technique is quick and has been shown to be remarkably accurate. At this point, the link has been made from dollar invested to customer satisfaction. Another estimate must be made to relate customer satisfaction to retention. Again, the marketing managers at a firm can provide accurate estimates of this data. Thus, the link has been made from customer satisfaction to retention and market share, and thus from dollar invested to bottom line impact.

Another benefit from enabling is more direct, the increased sales of processors or product. Valuing the increase in processor sales as a result of enabling is difficult, especially for a program such as the Mobile Whitebox program. One might assume that Intel could sell very few processors without some enabling activity. However, even without enabling, the market would accept some boxed mobile processors, with the engineering taking place downstream to support the product. Perhaps, a more relevant question is how many processors are sold initially, and how quickly the market is enabled to accept new processors when the prices and margins are the highest for Intel. Because of competitive pressure in the industry, prices of microprocessors are a decreasing function of time. For a company like Intel to extract the maximum profit from the market, it wants to sell while prices are still high. Enabling the market for its products is a key aspect of this strategy. Consider the following purely hypothetical scenario:

- The boxed processor program ideally results in 20k processors sold per year assuming full enabling
- With no enabling, the volume is reduced 10k
- With no enabling, the ramp up in sales is much slower – for example 1k, 2k, 3k, 4k per quarter
- The price of microprocessors decreases approximately 20% over the course of a year (linearly)
- The life cycle of a given processor is 1 year

In this scenario the financial benefit from enabling could be easily calculated. Assuming an initial processor price of $500 per unit, the revenue generated by an "enabled" product launch is approximately double the revenue of that generated by an "un-enabled" product launch. The question then becomes what kind of return one desires from enabling. The limitation of the above analysis has to do with estimating the "un-enabled" market. It is reasonable to assume that there will be some decrease in processors sold if Intel did no enabling. Another factor to consider is competition: AMD could offer a customer-enabled product that could rapidly displace a non-enabled Intel product. However, the exact level of volume decrease is difficult to predict with any accuracy. Again, managerial judgement is the simplest way to create an estimate.

To summarize, one can attempt to value the benefits of the enabling effort as follows:

1. Gather customer satisfaction data - the data should be a quantitative measurement system. Intel already has a system to measure customer satisfaction, which should suffice for this step.
2. Estimate the link between dollars spent and customer satisfaction metric. This requires managerial judgement: a method known as decision calculus can be employed[4].
3. Estimate the link between customer satisfaction and retention. This should be possible using historical data of customer satisfaction and whether customers have switched or adopted new suppliers. Again, managerial judgement must be employed.
4. Market segment share is determined from baseline market segment share and the effect of retention on market segment share. This gives a rough indication of revenue.

5. Add in the estimates of processor sales and high-margin sales to come up with the "value" of enabling.
6. Given the known costs, one should be able to create an optimal NPV graph for the enabling investment analogous to Rust and Zahorik's graph presented in Figure 4.4.

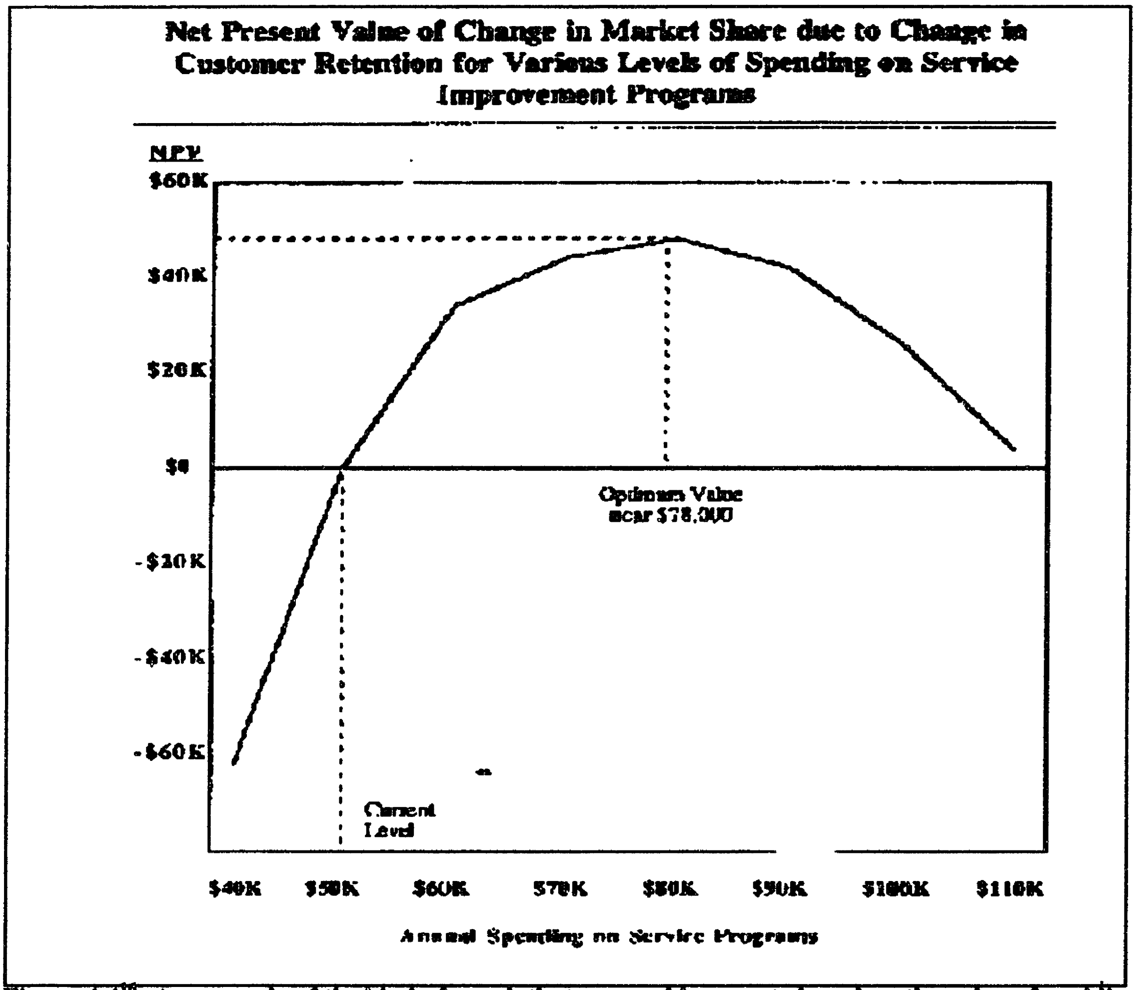

Figure 4.4[16]: An example of the kind of graph that one could generate based on the value of enabling method described above. This graph is shows the optimal spending level for a service program that a bank based on an analogous method.

4.4.2 Real Options Approach

The value of the enabling effort can be estimated as described above; however, enabling can play a larger strategic role in the sense that a well-executed enabling plan can literally make an industry happen. Consider a market in which activities analogous to enabling have taken place, for example the fiber optics industry[17]. Corning Inc. was a leader in the development of optical fibers, the backbone of telecommunication transmissions today. In the 1960's, optical fibers were used, for example, in cars to transmit light from a source to the dashboard. However, these fibers exhibited too much "loss" to transmit light the long distances needed for telecommunications. Fortunately, researchers at Corning made a major breakthrough in 1967 when they demonstrated extremely low-loss fiber that would enable long haul communications.

In the early 1970's a number of technologies related to optical fibers needed to be developed including light sources, detectors, and fiber connectors. Corning did not have the solutions to these problems, but to sell their core technology, optical fiber, they needed to develop a strategy to make the market happen. Thus, Corning put significant resources into enabling the telecommunications industry to use fiber optics. Activities included development of technologies, standards, and marketing these items. Corning invested over $100M during the period from 1967-1982 to enable the market for fiber optics. Yet, in 1982, Corning's optical fiber business was barely breaking even. It was not until the late 80's and the 90's that the optical fiber market exploded and Corning began to reap profits from its longstanding business.

How did Corning make the decision to make these long-term enabling investments? Did Corning "know" how things would turn out? The answer is that Corning probably did not know what was going to happen. The key to the long-term success of the optical fiber group seemed to come more from having a person at the top who "knew" that it was worth making the investment in optical technologies. Thus, while the company was apparently wasting money on a research project with little commercial payoff, it was visionary leadership that drove the development effort. What can Intel learn from this? Part of Intel's enabling effort involves product development and research that will help the industry as a whole advance, enabling the use of Intel's latest processors. Within this context, is there a better way to evaluate an enabling-like effort?

One possible approach to evaluating investments in enabling is the use of Real Options. Real options apply the theory of financial options pricing to capital investment decisions. There are numerous descriptions of real options in the literature. Faulkner presents a method for using "options thinking" and decision trees to value an investment in R&D[18]. The point of using a real options approach is that in the face of uncertainty, the discounted cash flow (DCF) models normally associated with ROI analyses do not properly explain future decisions and contingencies. Consider a comparison between typical DCF methods and a real options approach for a new product development effort.

Assuming an up-front investment of $6M with a 60% chance of a good outcome, 30% chance of an excellent outcome and a 10% chance of failure at the end of Year 1. At the beginning of year 2, another investment of $5M is required to continue the project, which at the end of Year 3 will yield some kind of return on the product. If the life cycle of the product is considered, the total return must be calculated and discounted back to the present to evaluate an ROI.

Now, using a DCF approach, one could estimate the net present value (NPV) of the product development by assuming the most probable outcomes or by taking probability weighted averages of the outcomes. However, the DCF approach does not really model reality, because if the project is deemed a failure, it will be abandoned after Year 1. An options approach explicitly recognizes this fact and can more accurately model the NPV of the development effort. Figure 4.5 illustrates a decision tree process outlining a potential R&D investment.

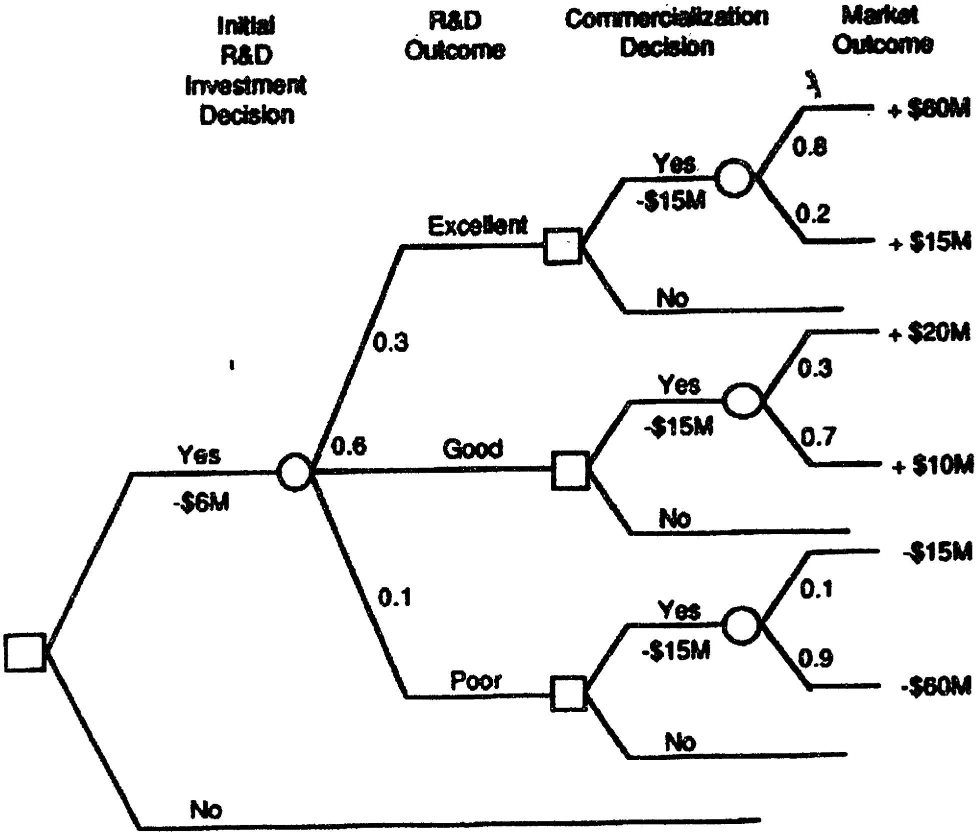

Figure 4.5[18]: Decision tree for an R&D investment. Use of several DCF approaches and an options approach are compared in the table below.

Valuation Method	**Net Present Value**	**Year 1**	**Year 2**	**Year 3**
DCF: Most likely results	-11.4M	-6M	-(15/1.12)M	$+(10/1.12^2)$M
DCF: weighted averages	-5.4M	-6M	-(15/1.12)M	+Weighted averages.
Real Options	+2.2M	-6M	-0.3(15/1.12)M	$+0.3(0.8*60+0.2*15)/1.12^2$

The options approach is the only one that produces a positive NPV for the project. The real options approach considers continuing the project only if the "Excellent" result is obtained, which will only occur with a probability of 30%. The point of this example is not to show that the options approach will always provide a higher NPV than a typical DCF method. Faulkner has concluded that this is not the case. Rather, the example should be considered a simple use of options thinking, and how such thinking might be applied to an investment in customer enabling.

4.5 Summary

This chapter has described the customer enabling effort at Intel and its relation to the Whitebox project. A combination of enabling models from the desktop and mobile markets will be used to enable a thermal solution for the boxed mobile processor. This decision led to a discussion of how to value the enabling effort as a whole. The use of a customer satisfaction framework and a real options approach were suggested. As the enabling process for the boxed mobile processor is further defined, strategic issues for the industry must also be considered. Chapter 5 discusses these issues and introduces the potential strategic value of enabling.

5 Chapter 5: Corporate Strategy and the Boxed Mobile Processor

5.1 The Mobile Computing Market

The mobile computing market as a whole is growing rapidly, and it is in Intel's best interests to see its continued growth. With more devices that require microprocessors, Intel has the potential to serve a larger market. However, like the desktop computer, the notebook is seeing steep price declines, resulting in price pressure on the microprocessor. Figure 5.1 illustrates these trends.

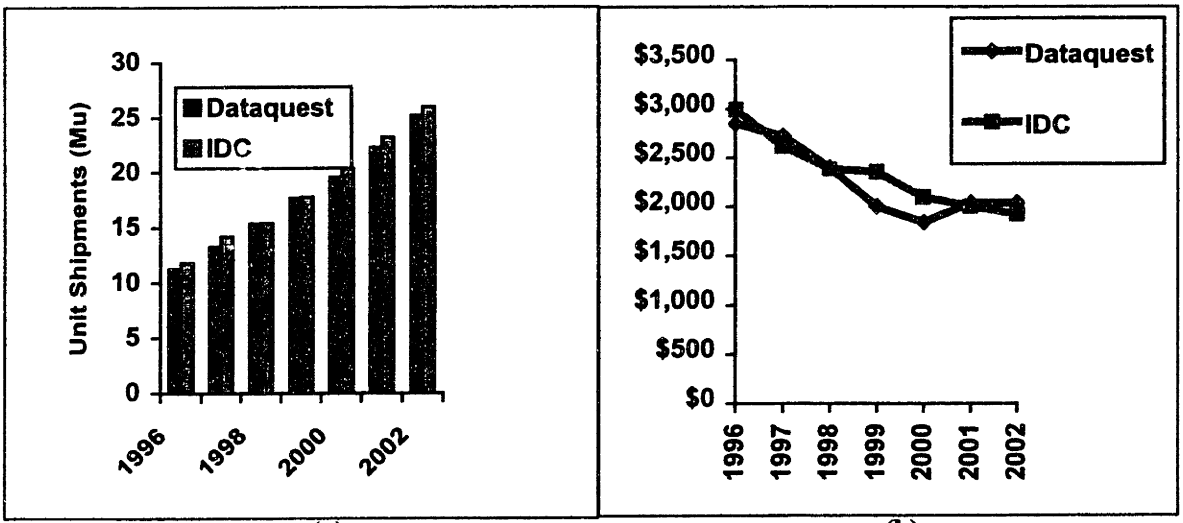

(a) (b)

Figure 5.1[19]: (a) Worldwide market for mobile computers. (b) Average Selling Price (ASP) of notebook computers worldwide.

The aim of the Mobile Whitebox program is for the boxed mobile processor to increase the Total Available Market (TAM) for mobile computing. The worldwide small to medium business is a new adopter of the notebook computer, and Intel can sell more mobile processors to this market. There are a few key assumptions that go along with the claim of a TAM expansion:

- Large OEMs do not lose business to the dealers now selling both OEM systems and their own branded mobile whitebox systems
- Market share of mobile processors is taken from Intel's competitors (i.e. AMD)
- There is no cannibalization of business of desktop processors; the dealers will continue to sell as many desktops as before, and notebooks are simply added on

Given these assumptions, one could estimate the return on investment of the boxed mobile processor program. Three potential scenarios of boxed processor volume are presented in the Figure 5.2, resulting in three different ROIs.

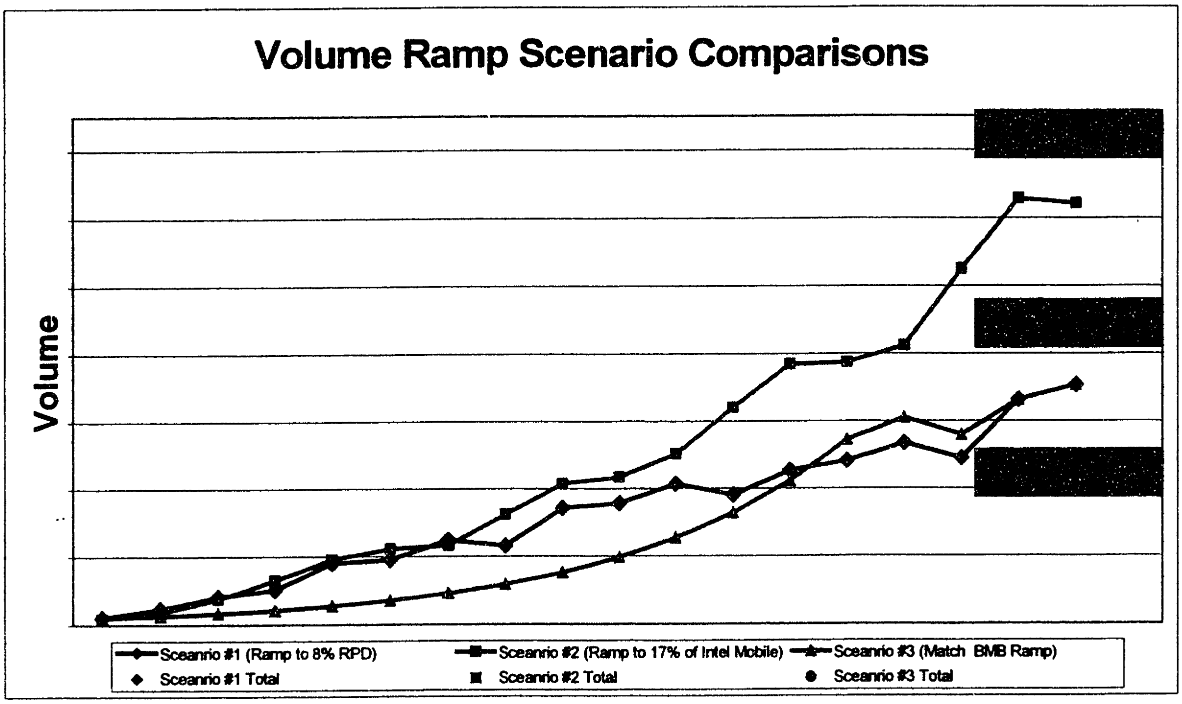

Figure 5.2[20]: ROI analysis of the boxed mobile processor. Scenario 1: The mobile processor becomes ~8% of RPD's business. Scenario 2: The boxed processor becomes ~17% of Intel's mobile processor business. Scenario 3: The boxed processor sales approximately match Intel's motherboard business.

Keeping the assumptions in mind, the boxed mobile processor appears to be a sound investment.

5.2 *Standardization Effects*

A side effect of the boxed mobile processor project might be further standardization of the notebook computer. The Whitebox thermal solution is significant for more than just its functionality. Through Intel specifying a thermal solution for a notebook computer, and perhaps some of the area around the microprocessor and thermal solution, a significant portion of the system is specified. If the processor location and air inlet and outlet are specified, even more system design freedom is removed. The logical conclusion of these specifications is the creation of a fixed notebook layout. If the thermal solution design is exceptional and adopted by a significant portion of the industry, a standard notebook design may result. The relevant issue is whether standardization has a positive or negative effect from Intel's perspective. The other issue is whether Intel's present strategy with regard to influencing standardization in the industry has been the correct one.

To help understand the effects of standardization, I will discuss three topics

- How standardization typically evolves and affects an industry
- Current standards in the notebook computer industry and projections for the future
- How Intel has treated the issue of standardization and what it might do going forward

5.2.1 Standardization in an Industry

There are two major categories of standards in an industry[21]:

- Quality standards
- Compatibility standards

For the computer industry, compatibility standards are of most interest. Compatibility is relevant for notebook computers because the modularity of the components of the computer relies on certain standard interfaces within the system. To examine how standards evolve in an industry, consider Utterback's dominant design concept[22]. A dominant design is created when some kind of innovation results in a new product or feature set in a given product. The dominant design is the one that competitors in an industry must adopt or follow if they wish to win a share of the market. An example of a dominant design is the QWERTY keyboard for typewriters: once it was adopted, typewriter firms either built their machines around the design or went out of business. If a dominant design has developed, one would expect the pace of innovation to slow, standards to be enforced, and the number of firms in an industry to peak and decrease. Utterback has demonstrated this effect in numerous industries including automobiles, televisions, and typewriters. Dominant designs foster the creation and adoption of standards, and shift the basis of competition to process improvements (manufacturing competency) rather than product improvements. Once a dominant design is in place in an industry, firms must either try to change the standard (innovate) or to compete within the standard.

Of course, a given firm in an industry is always trying to maximize its own profits, and generally a very open industry standard makes this more difficult. An example of the difficulty that standards impose on a firm is depicted in Figure 5.3.

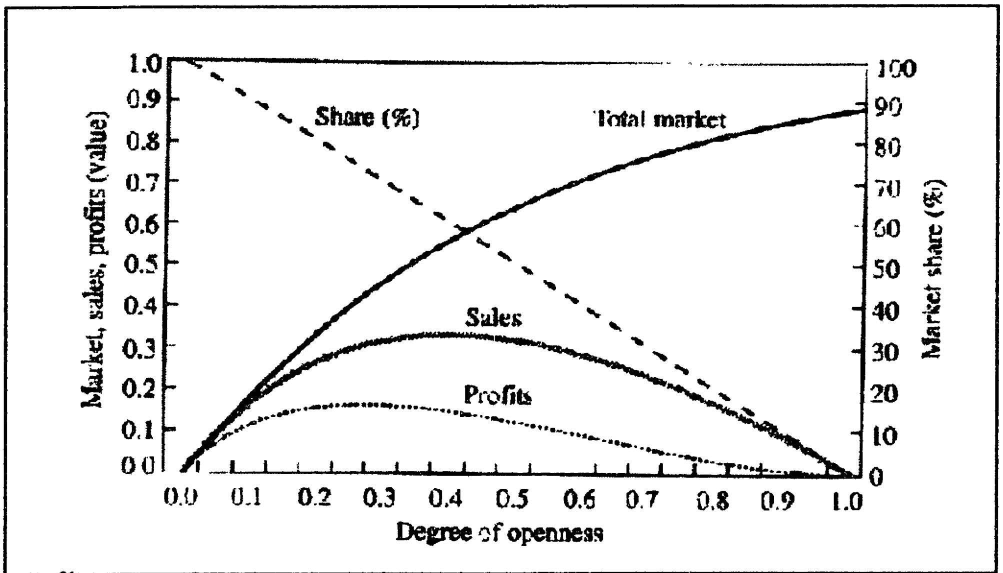

Figure 5.3[21]: An individual firm seeks an optimal amount of openness to profit from standardization.

The tradeoff illustrated in Figure 5.3 is that the total market for a product is likely to increase as the standard becomes more open and universal. And, as the market increases, the share of a given firm will decrease while sales and profits will peak somewhere in the middle. Some companies, such as Sun Microsystems, embrace the policy of open standards, thinking that any technological advance they gain is only temporary. Other firms pursue proprietary technology as the means to generate competitive advantage.

Charles Fine argues that an industry itself fluctuates back and forth between an open, modular product and an integral, proprietary product[23]. His concept of industry flow is illustrated in Figure 5.4, the "double helix" concept.

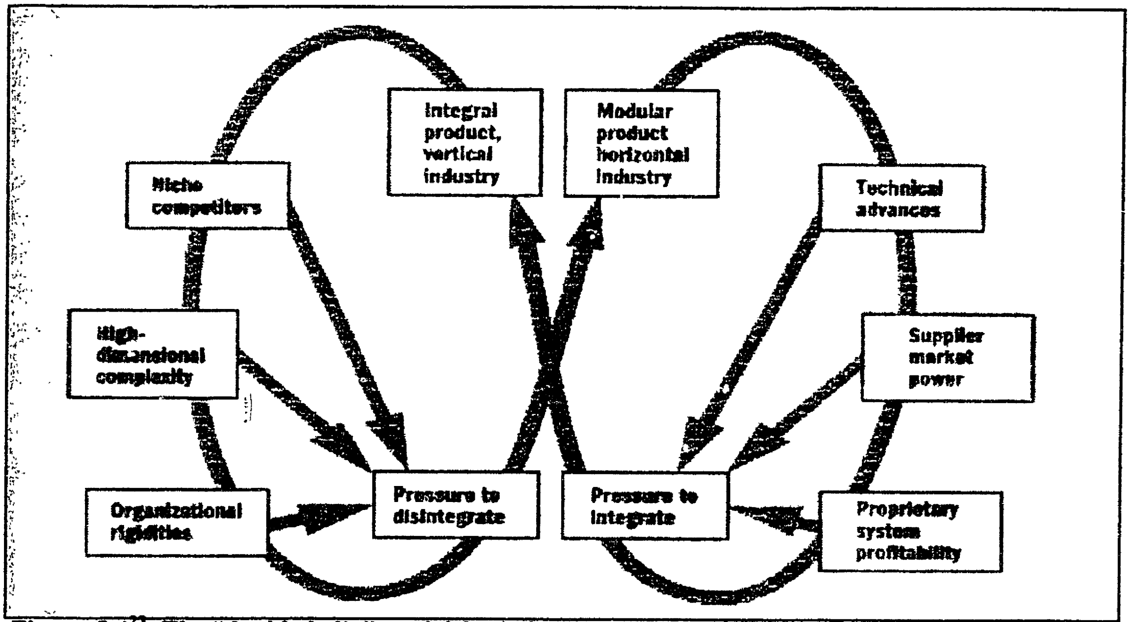

Figure 5.4[23]: The "double helix" model for industry and product structure. Over time, the model predicts that an industry will shift back and forth between a vertically integrated, proprietary product and a horizontally separate, modular product.

The implication of figure 5.4 is that when an industry is characterized by vertically integrated competitors, factors such as complexity, niche competition, and organizational rigidities push that industry toward disintegration and a horizontal structure. Likewise, if an industry is characterized by a very horizontal structure with modular products, pressures such as technology, supplier power, and the desire for margins push companies toward more vertical forms. A horizontal industry with modular components implies a high level of compatibility standardization. Note that someone will profit in either scenario, but how the profit is divided changes depending on the industry structure.

5.2.2 The Notebook Computer Industry

Where does the mobile computing industry fit in the realm of standardization? One could argue that the notebook computer has already become largely standard and modular. Most every notebook computer has certain features: the keyboard is nearly standard, the displays are similar, most hard and floppy drives are removable and modular with the battery, and the PCMCIA card has become standard. Overall, performance of mobile computers has approached that of desktop PCs, and the slogan for many mobile computers is that they are a "desktop replacement". As the market and the products mature, competition in the mobile arena is shifting to price, and an even greater level of standardization is probable. Figure 5.5 illustrates a model that describes some of the drivers of the trend toward standardization in the industry.

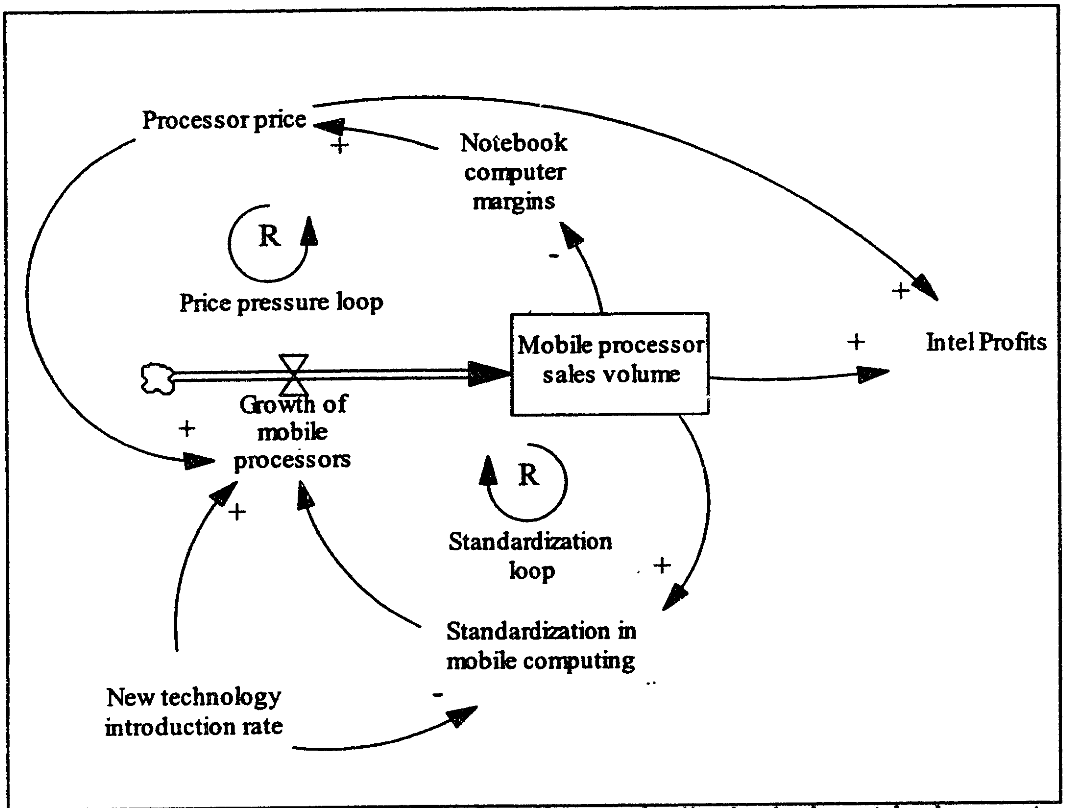

Figure 5.5: An illustration of the forces that affect standardization in the notebook computer industry today. Two "reinforcing" loops are at work that drive the prices of notebooks and processors down and encourage standardization in the industry. The loops interact in the following way: the Price pressure loop means that lower prices increase growth and volumes, promoting the Standardization loop, causing more growth, margins to fall, prices to fall, and growth to increase.

In Figure 5.5, "new technology" includes new hard drives, displays, processors, or peripherals that might alter the typical design of a notebook. While such new technology would promote growth of the industry, it is likely to decrease standardization. On the other hand, as the volumes of notebooks rise, standards are favored because firms gain economies of scale in manufacturing. And, as standards are favored, a horizontal industry structure is likely to evolve, where a different firm can produce each modular component.

To examine the current state of evolution of the industry, one might examine the number of firms that manufacture notebook computers over time. By manufacturer, I mean the final integrator of the components of the system. Intel experts have estimated that the count of notebook manufacturers has peaked and begun decreasing in the past ten years, indicating that a dominant design *has emerged*. Validation for this conclusion is given by a recent announcement that one of the top notebook manufacturers is exiting the market. This manufacturer is one of the larger, more vertically integrated companies currently in the industry. The current industry structure fits the dominant design model, as it is extremely horizontally divided, with small players occupying each segment in the value chain. Independent companies make the plastics for the chassis, the motherboard, the chips on the motherboard, disk drives, displays, etc. Recent data indicates that the top ten notebook suppliers control 70% of the market, demonstrating the importance of economies of scale, manufacturing competency, and brand. The implication of such an industry structure is that subsystem components are very modular and compatibility standards must already exist to a large extent.

Into this environment, the boxed mobile processor and the Mobile Whitebox enter. Figure 5.6 illustrates how the Whitebox project might impact the notebook computer industry and standardization therein.

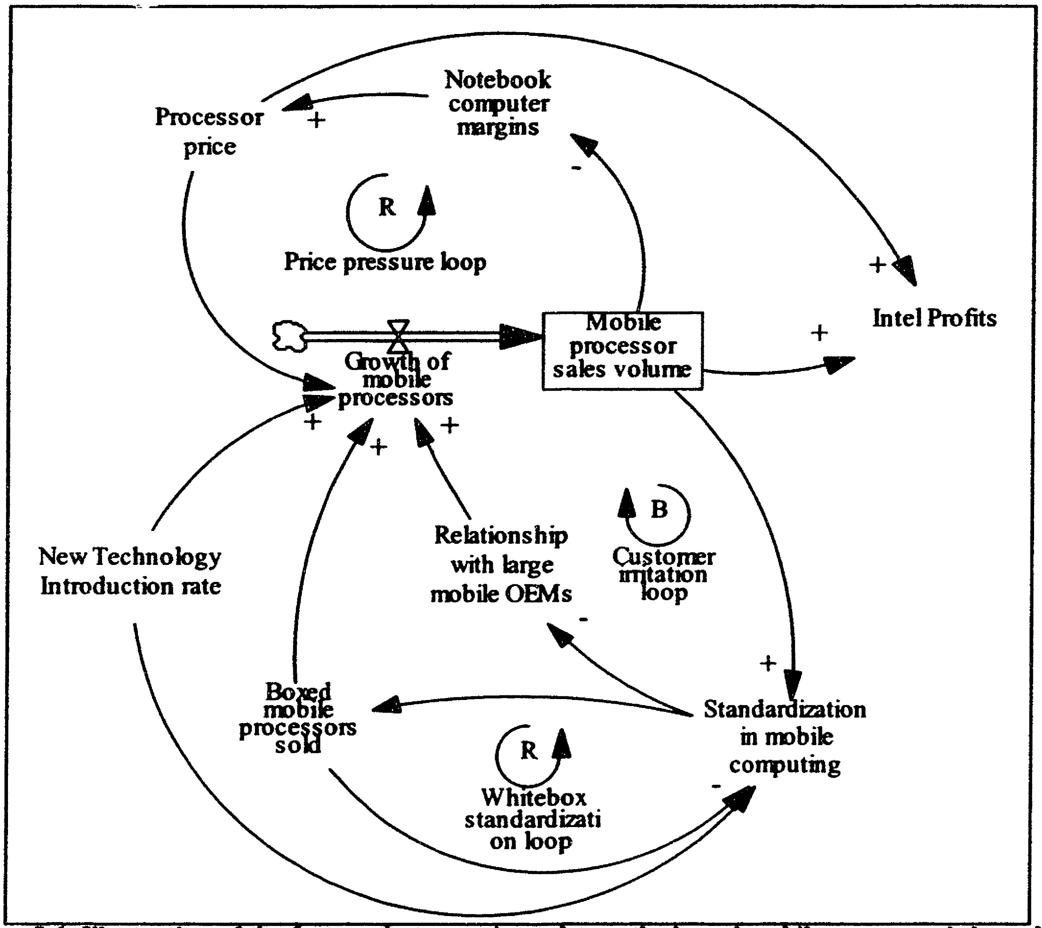

Figure 5.6: Illustration of the factors that come into play as the boxed mobile processor is launched. Assuming that the boxed processor promotes total growth of the market, standardization is also promoted (Whitebox standardization loop). This in turn causes large OEM customers to lose margins and market share and get irritated with Intel selling directly to the channel, thus slowing Intel's mobile processor growth (Customer irritation loop).

Two effects that the Whitebox project is likely to have are as follows:

- Further standardization by increasing volumes and decreasing design freedom
- Because Intel is trying to extract further profits from the channel, its direct OEM customers could retaliate by choosing competitor's microprocessors for notebooks

Thus, while standardization might benefit the growth of the notebook computing industry as a whole, the concern is that it will do so at the expense of Intel's market share. Also recall that an increase in standardization is likely to increase price pressure on Intel microprocessors. A second important point to consider is the market impact of the boxed mobile processor. This is a relatively low volume product from Intel, comprising an estimated maximum of 20% of Intel's mobile processor business (and this kind of volume is still years away at the earliest). Whether this volume of product has the potential to influence the mass market remains to be seen.

The current question of interest is whether the notebook computer market as a whole will ever evolve into a true Configure-to-Order (CTO) market similar to that of the desktop computer. As an example of what I mean by a "true" CTO market, Dell will build a desktop computer to the specifications of the customer including components and peripherals. If the notebook platform is standardized, the door is open for an identical model to be used, whereby the customer does not care what the "box" looks like, only what is inside of it.

A critical variable that would impact the CTO evolution in notebooks is the new technology introduction rate (see Figure 5.6). Two major types of technology that have the potential to alter the dominant design of notebook computers are networking technologies and newer, faster microprocessors. On the network technology side, it is possible that the functionality of handheld computing devices such as personal digital assistants and cellular phones could find its way into the notebook computer. Alternately, these devices might absorb some of the processing power now reserved for the mobile computer. The second new technology driver is the microprocessor. As has been documented, with improving processor performance comes a significant thermal challenge. If new thermal technology must be introduced to physically alter the design of a system, the current dominant design may change. Alternatively, a new design for the microprocessor might emerge to support the dominant design of the notebook computer. One could view the notebook and the microprocessor as two complementary yet competing dominant designs, and it currently appears that the system level design is the more powerful of two. This could prove to be a very interesting challenge for Intel.

Overall, the notebook industry appears to be nearing a fully mature stage. The business models of some of the computer OEMs have reached the point where they are purely managers of the supply chain rather than manufacturers of systems. Some prominent brand name OEMs use a direct-ship model to send notebooks that they buy from Taiwanese manufacturers straight to the customer. As Intel moves forward within the industry, it must continue to be aware of how standards are evolving and influencing the notebook computer.

Another viewpoint is that while the "notebook computer" has experienced significant standardization, the "mobile computer" is still in its infancy. The emergence of the PDA, cellular phone, and Internet appliance may give a whole new meaning to mobile computer. The overall increase in networked devices may present an opportunity for Intel to influence the next dominant design in mobile computing. The final shape, size, and capability of the mobile computing/communications device have yet to be decided. Enabling can be a key method for Intel to influence its design.

5.2.3 Intel's strategy toward standardization

Before discussing Intel's current strategy in the mobile computing environment, consider what Intel has faced in the desktop computer market. The desktop computer platform is now relatively standard and its history could give some insight into what the mobile platform may experience in the future. The desktop platform began with IBM's open standards strategy for the PC. However, the physical size and shape of the motherboard was actually dictated largely by Intel, who introduced the ATX form factor in motherboards. Intel's introduction of this standard was largely made to spur demand for its latest and greatest microprocessors, and in this way, the standard was a success. Using the same models presented in Section 5.2.2, Intel increased the volume of its processors and reinforced the ATX standard. Given an Intel-created standard, releasing a boxed processor in the desktop market was simple.

By contrast, Intel has chosen not to compete in the market for notebook motherboards. With no powerful organization leading a platform standardization charge, the industry has slowly evolved into the form described previously. Given the trends and previous analysis, the question is what Intel should do going forward. Simply put, the issue for Intel appears to come down to a question of overall profit in the channel. If the notebook computer is becoming a commodity, then prices, margins, and profits will fall and the microprocessor will come under price pressure. In this case, the total market expansion implied by standardization benefits Intel, the idea being to capitalize on economies of scale and Intel's manufacturing competency. On the other hand, if, by delaying standardization and commoditization, Intel can extract extra profits from the industry, it should do so. In this second case, allowing standards to evolve on their own is a viable strategy.

Another thing to consider is if Intel's OEM customers have the power to take profitability away from Intel. What I am suggesting is that because Intel has such a strong market presence, a market size increase only helps increase Intel's revenues. However, Intel dictating or influencing a standard might irritate large OEM customers, perhaps pushing them to adopt competitive microprocessors. In a market with open standards, the microprocessor is just another modular component. With AMD's recent re-emergence, especially in the low-cost market segment, competition could easily extend into the notebook arena.

Then, given the current state of the industry and the current capabilities at Intel, an overt push for standardization does not make sense. There is already an acceptable level of standardization and growth in the mobile computing market segment. Furthermore, Grindley concludes that standards are highly dependent upon what happens in the very early stages of an industry. Notebook computing, while a relatively young market, is rapidly approaching maturity. It is not obvious that Intel can even sway standardization in the market at this late stage. The fact that the microprocessor may have to change to adapt to the dominant notebook design supports this assertion.

What does this mean for the boxed mobile processor program? This author concludes that the Whitebox Program will not significantly affect Intel's position in the mobile processor market. With the notebook computer already far along the path to standardization, it is not clear that the boxed mobile processor will truly accelerate the process to a fixed platform. Rather, there are other concerns that the Intel must address.

For example, Intel may have to work even harder to enable its mobile processors. Given the thermal challenge and trends outlined in Chapter 3, it is not immediately clear how to make the newest and fastest processors work in a mobile computer. Unless the thermal solution developed for the boxed processor is thermally scaleable without significant changes to its size and shape, it may be impossible to standardize the mechanical layout. This implies continued opportunities for new technology introductions, proprietary designs, and continued higher margins in the mobile segment. In this sense, Intel must continue to lead thermal solution development activities, enabling the entire industry. Because, the opportunity does exist for a proprietary technology to disrupt the equilibrium that currently exists.

Another concern has to do with the evolution of the mobile computer. The emergence of cellular phones, PDA's, and Internet appliances will likely transform mobile computing as we know it today. As wireless connectivity increases, the trend toward the smaller computer, perhaps even palmtop computer could continue. It is impossible to predict what the dominant design for this new mobile computer will look like; however, Intel has the opportunity to influence this design through its enabling program. While it might be too late to change the "notebook computer", it is certainly not too late to use enabling to influence the next "mobile computer". Enabling customers to use Intel products in the next generation of mobile computing devices can thus provide Intel with a long-term competitive advantage.

5.3 Conclusions

Trying to predict the direction that an industry will go is extremely difficult at best. While the notebook computer industry fits very well into some of the models described previously, it is best to remember that no model is perfect. Instead, I hope that the reader has gained an appreciation for the dynamics that are facing the industry and the issues that Intel could face going forward.

6 Chapter 6: Organizational Structure

6.1 Overview

As has been previously described, the boxed mobile processor program involves complex technical and business issues. Thus, it is no surprise that a cross-functional, cross-geography team was assembled to execute the project. This chapter provides a discussion of the groups involved in the project, their interaction, and an evaluation of the effectiveness of this organizational structure.

6.2 Groups involved in the boxed processor program

Table 6.1 defines some acronyms and explains the key Intel groups that are involved in the boxed processor program.

Group	Purpose
Reseller Division	The Reseller Products Division (RPD) owns the interface with the dealer channel. RPD is the organization that sells all of Intel's boxed processors.
Mobile Marketing	The Mobile Computing Group's Marketing organization owns the interface with the OEMs participating in the boxed mobile processor program.
Intel Taiwan Support	The Mobile Taiwan Support Center (MTSC) is the group that works to enable OEMs (especially in Taiwan) to use Intel products.
Mobile Engineering	The Mobile Platform Technologies (MPT) group is a platform (board level) architecture group.
Mobile Enabling	The Mobile Mechanical Applications (MMA) group is the US counterpart of MTSC, working to enable OEMs to use Intel products.

Table 6.1: Key organizations internal to Intel that are supporting the boxed mobile processor program, especially with regard to customer enabling issues.

Having been involved with the boxed mobile processor project nearly from its inception, I observed several potential problems with the organizational structure of the groups. As a framework for this analysis, I will use the Three Lenses approach described by Nadler and Tushman[24]. The focus of this analysis will be on the organizations involved in thermal issues related to the boxed mobile processor.

Strategic Lens

From a strategic perspective, one looks at the organization from afar, assumes the organizational chart to be perfectly accurate, and assigns appropriate roles and responsibilities for a project.

Using this lens, I observed that for the boxed mobile processor project, too many groups are involved in thermal-related issues, and the groups involved are not necessarily the best ones for the job. First of all, consider the best-case scenario, where the organizations with the expertise perform the functions required to enable the thermal solution. In this situation, the Mobile Enabling group is already the group responsible for thermal enabling of OEMs, and has experience dealing with both the technology and the customers. Intel Taiwan can serve as the local (from a geographic point of view) contact for the OEM, though it need not be very involved from a technical standpoint. The Reseller Division can serve as the expert on the dealer channel, but in essence should not really be involved in the technical enabling of the thermal solution.

Unfortunately, the real organizational interaction looks something like what is pictured in Figure 6.2.

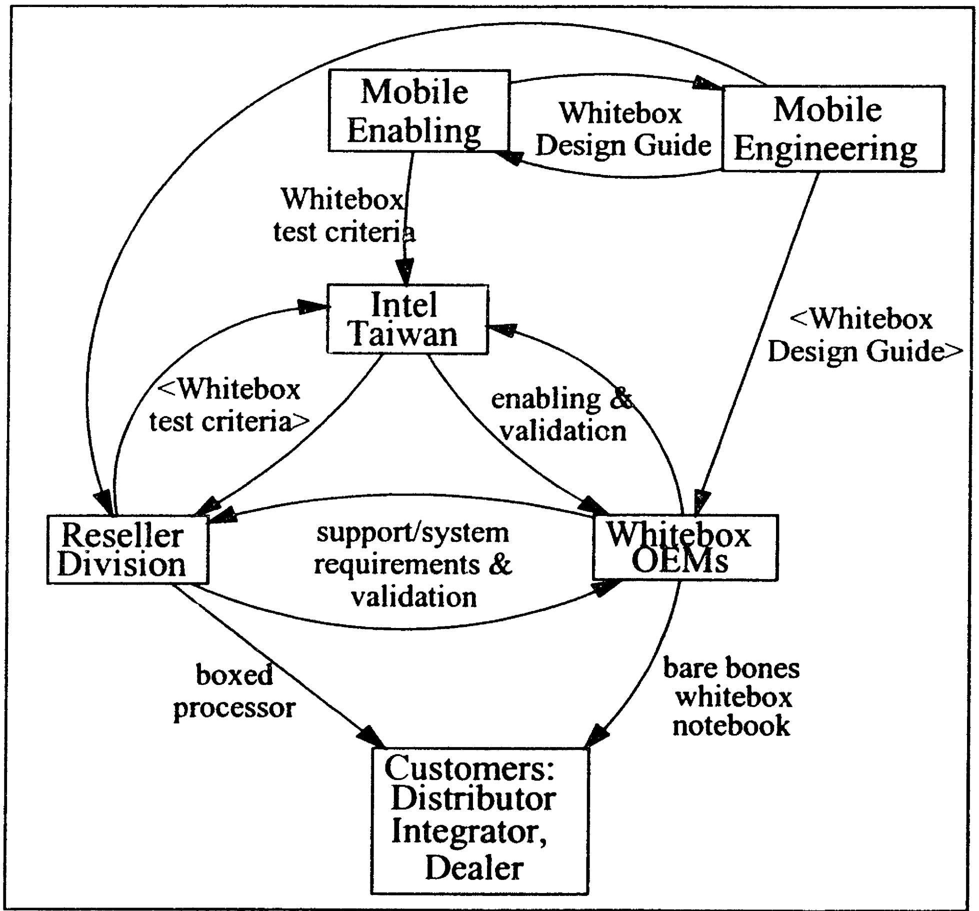

Figure 6.2: Representation of the organizational interaction between groups internal to Intel. Boxes indicate the organizational groups while the text along the arrows indicates the informational material that is being exchanged between groups.

Nothing is inherently *wrong* with the picture in Figure 6.2, but its complexity is alarming, even for a large organization such as Intel. Three things are striking:

- Duplication of effort
- Single voice challenge
- "Passthrough" communication

The duplication of effort is a common occurrence at Intel, and it helps create a system of checks and balances that ensures exceptional product quality. For example, Mobile Enabling and Intel Taiwan essentially have the same charters, but Mobile Enabling has more experience. Of course, Intel Taiwan is geographically closer to customers (most mobile computer OEMs are located in Taiwan), thus is a key customer point of contact at Intel. Mobile Engineering is a complementary group to Enabling, but a group whose expertise grew out of a single Intel product, the Intel Mobile Module. Mobile Engineering developing the Design Guide for the boxed mobile processor is odd given of the wealth of experience that Mobile Enabling has with design guides for notebook computers. The complicating factor is that each of the groups identified in the diagram reports to different direct management, creating potential alignment and prioritization problems.

The single voice challenge can be explained by noting that a large number of arrows in Figure 6.2 are pointing at the Whitebox OEMs. In the chart below, the groups within Intel and their communication with the Whitebox OEMs is described.

Organization	**Location**	**Topic/Level of communication**
Mobile Enabling	Santa Clara, CA	Product based and technical
Intel Taiwan	Taiwan	Technical
Mobile Engineering	Oregon	Technical/design guide
Reseller Division	Folsom, CA	Market level

Here, we notice two things. First, there are many levels of communication with the Whitebox OEMs and second, Intel Taiwan is sometimes serving as a passthrough for communication from Mobile Enabling to the Whitebox OEMs.

Political Lens

When looking at how the organizational structure of the project developed, one must also consider political factors. A political analysis identifies the key stakeholders, their goals, and their basis of power. Mapping a political lens analysis onto a strategic lens analysis enables an improved organizational structure to be developed. Table 6.3 presents the groups involved as described above, with their goals for the Whitebox project and their sources of power identified.

Group	**Goals**	**Sources of Power**
Reseller Division	Expand the dealer market for boxed products. Serve that customer well.	Expertise in the market. Previous relationships with the market. Experience with boxed processor products.
Mobile Marketing	Expand the market for mobile processor products.	Experience dealing with major OEMs.
Intel Taiwan	Recognition/validation? An entry into China?	Direct and close relationships with OEMs in Taiwan.
Mobile Engineering	New applications	Platform experience. Previous attempts at a boxed mobile processor.
Mobile Enabling	Enabling activities for the major OEMs.	Expertise with technical issues. Experience and contact with large OEMs.

Table 6.3: The "political" analysis of the groups involved in the Mobile Whitebox project.

Outwardly, there are no major goal conflicts between internal organizations. One might argue that by enabling smaller Whitebox OEMs, the larger OEM effort might be compromised somewhat. For example, Mobile Enabling's specific charter is to support large OEMs, and now may be pulled into supporting smaller Whitebox OEMs. However, one could easily argue that by enabling a larger part of the industry, the larger goal of enabling is fulfilled. One group that deserves special mention is the Mobile Engineering group. This group was responsible for supporting the Intel Mobile Module product, a product that is being phased out of production. Thus, the people within this group are trying to find new projects to work on that relies on their skills and expertise. The Mobile Whitebox project provides just such an opportunity.

In the context of the boxed mobile processor project, each group has certain sources of power. First, several individuals within Mobile Engineering had previously proposed a boxed mobile processor. Until the most recent incarnation described in this document, all efforts had failed. Thus, these people were regarded as the experts within the organization on issues related to boxed mobile processors.

Cultural Lens

Culture at Intel is extremely strong. This is reflected in artifacts, values, and assumptions. It is a source of unity within an organization that is geographically and functionally diverse. Some of the key features of this culture are:

- Equality: everyone is an equal, everyone has a cubicle, and input from all is accepted and appreciated
- Belief in corporate values: people go so far as to attach a statement of core values to their identity badges
- Paranoia: everyone is encouraged to challenge assumptions, recommendations, and ideas

The relevant question is whether there are subcultures specific to the groups involved in the Mobile Whitebox project and whether these subcultures conflict. The only subculture that I found that was as strong as the Intel culture was one related to geography. Namely, the group in Taiwan was much more focused on the organizational chart than the other groups. This was reflected in a reluctance to make substantial decisions and accept responsibility. However, this was mitigated by other groups' desire for more control. Further, the sense of equality at Intel fostered positive interaction of the many groups involved in the project. Overall, Intel culture supported the way the project was run.

6.3 *Further Analysis*

Still, the way that the Mobile Whitebox project was structured and executed appeared to be sub-optimal. An ideal staffing process might have begun with the strategic marketing and planning group in MCG allocating specific resources within MCG for the project. The staffing process could have been reinforced by vocal messages from top level management stressing the importance of the project. Finally, coordination between the two parent organizations, RPD and MCG, could have been outlined at the beginning of the process rather than emerging in an ad hoc fashion. The nature of this project appeared to be best suited a top-down method of project structuring and management. These things said, the Mobile Whitebox project still managed to get the product launched, which leads to a question of "efficiency versus effectiveness".

The way that a project can be managed has been studied extensively in the literature. Certain project management structures have been identified, ranging from a functional organization to a pure project team with some variants in between[25]. Selecting the right organizational structure is important for the efficiency and effectiveness of the project. The idea is to see whether the approach that was used (an emergent, decentralized one) was reasonable, or whether a more structured and efficient plan might have been better for the organization as a whole.

Duncan identifies several steps that one should take in defining the structure of an organization[26]:

- Environmental analysis – determine complexity
- Identify the pace of change in the environment – dynamic or static
- Determine if market segmentation is possible

Crawford translates these basic principles into a questionnaire (Figure 6.4) to help with the selection of the organizational structure[27].

______ How difficult is it, in general to get new products through this firm?

______ How critical is it for the firm to have new products revenue at this time?

______ How much risk to personnel is involved in this particular new products project?

______ How important is speed of development on this particular project?

______ Will the products require new procedures in their manufacture?

______ Will the products be using new procedures in their marketing?

______ Relative to other projects, how high is the dollar contribution expected from the product in this project?

______ How unskilled in new products work are the people who will be assigned to this team?

______ Will the politics this project faces be more difficult than usual?

______ Are we as a firm experienced in using the more highly projectized forms of organization?

Figure 6.4: A survey of questions to ask to determine the best organizational structure to use to support a new product launch. Each question is scored on the basis of difficulty with a higher score representing more difficulty. A lower score (<25) implies that a more functional structure would work while a higher (>35) score implies a project team/matrix or even a venture.

The Mobile Whitebox project was characterized by the following characteristics:

- New product - actually an "old" product for a new market
- Established customers already exist
- Internal expertise exists to support the customers and the product

- The marketing was "owned" by one division but the product was "owned" by another- this led to a matrixed team executing the project

Given the context of the project, the environmental and the internal analysis, a more functional organizational structure may have been suggested. The boxed mobile processor faced a mostly known customer, the organization was experienced with new product launches, and the overall program offered only incremental revenue to Intel. Thus, when lower level project managers tried to pull in resources from functional groups, there was some slowdown and resistance.

With regard to organizational expertise, Intel is broadly divided into groups with specific knowledge. The idea here is that some kind of strategic planning or marketing group (or any higher level management group) would formally build a team to incorporate the appropriate expertise from the organization. Also, because the divisions involved in the Whitebox project did not have a high degree of familiarity with each other, a higher level organization would have been useful as a mediation tool.

In general, the Whitebox project is symbolic of what can happen to a complex organization when two different functional or divisional groups must interact. The general lesson for Intel is to use the tools of organizational analysis when structuring for a new product launch. The tools described above should be simple and general enough for any organization to use in staffing a product development/launch project. When an idea is generated in one group, the challenge is how to coordinate efforts to support the idea at all levels. Intel is currently diversifying its product line, and entering a number of new businesses will entail a number of new product launches. As these projects are staffed, it makes sense for the company to employ the kind of thinking outlined in this chapter.

7 Chapter 7: Summary and Conclusions

The Mobile Whitebox project could be viewed essentially as an incremental revenue opportunity for Intel's Mobile Computing Group. However, such a simple view leaves out the strategic implications and lessons that can be learned. From a technical perspective, while the work that was done was not groundbreaking in terms of thermal management technology, the development effort provided a look into what might happen to the notebook computer platform in the long term.

First, we saw that a "standard" thermal solution capable of meeting the thermal demands of the next processor platform is technically feasible. Still, it is currently unclear within Intel what the final decision will be with regard to enabling a specific Whitebox thermal solution. Feedback from notebook OEMs indicated that a fully specified thermal solution design might drive too close to a standard notebook platform, thus reducing the value added by systems integration. This has led to the question of how Intel can successfully enable the boxed mobile processor and how Intel might value the enabling effort as a whole. Two valuation methods were presented: the first ties enabling to customer satisfaction and time-to market, the second relies on the use of real options. One can relate enabling to customer satisfaction metrics, then to market share, and thus to the bottom line. Enabling must also be valued in terms of maximizing product sales immediately after introduction, when margins are at their highest. An alternate valuation method is to use Real Options, considering the enabling effort similarly to an R&D project. This approach is useful because some part of enabling involves research and product development. Real options provide a decision tool for managers to use in the face of uncertainty in product and technology development.

Organizationally, the enabling effort for the boxed mobile processor showed that when differently motivated groups within an organization must interact, conflict and inefficiency are likely to arise. Customer enabling is especially susceptible to this because the functional work that goes on is both design and marketing-oriented. As Intel goes forward to enable its newer, more diverse product line, a coordinated understanding of the product and market characteristics will be critical to structuring the organization correctly.

Strategically, enabling can be viewed as Intel's opportunity to influence the dominant design of the downstream product. Through enabling, Intel could have potentially set a notebook computer standard that conformed uniquely to the strengths of Intel's mobile processor. While it is now probably too late in the life cycle of the notebook computer to try to implement standards, Intel must be aware that the current dominant design of the notebook could challenge the dominant design of the microprocessor. As Intel continues to enable the mobile computing industry for its newest products, it must consider the strength of the current dominant design of a notebook computer. If the processor must change to accommodate the notebook, this could prove costly to Intel. However, while the "notebook computer" dominant design may be strong, the opportunity to use enabling to influence the next "mobile computer" still exists. The emergence of wireless devices such as PDA's, cellular phones, and Internet appliances could mean that the notebook computer as we know it today might not exist in 10 years. By enabling this emerging market to use Intel processors, Intel can play a large role in setting standards for the next generation of mobile computing devices. The lesson for components suppliers in a rapidly evolving product design is to consider using an enabling strategy that biases the long-term system design to the strengths of component-supplying firm.

8 Appendix 1: Raw Data from Thermal Tests

The raw data is presented in the form of thermal run charts, which indicate the temperature at locations along the thermal path of the thermal solution. Thermal solution prototypes were attached to the thermal die. The thermal die is form factor replica of a microprocessor; it is a resistor that is powered by an external power supply to an appropriate power level. Temperatures along the thermal solution were measured by physically placing thermocouples at the various locations. Temperatures were tracked over time as the power of the thermal die was increased. Typical abbreviations used in the raw data charts include:

- HP - heat pipe
- RHE - finned heat exchanger
- HSC - block attachment to die
- Core - die (processor) temperature as measured by an on-die thermocouple

8.1 First concept data

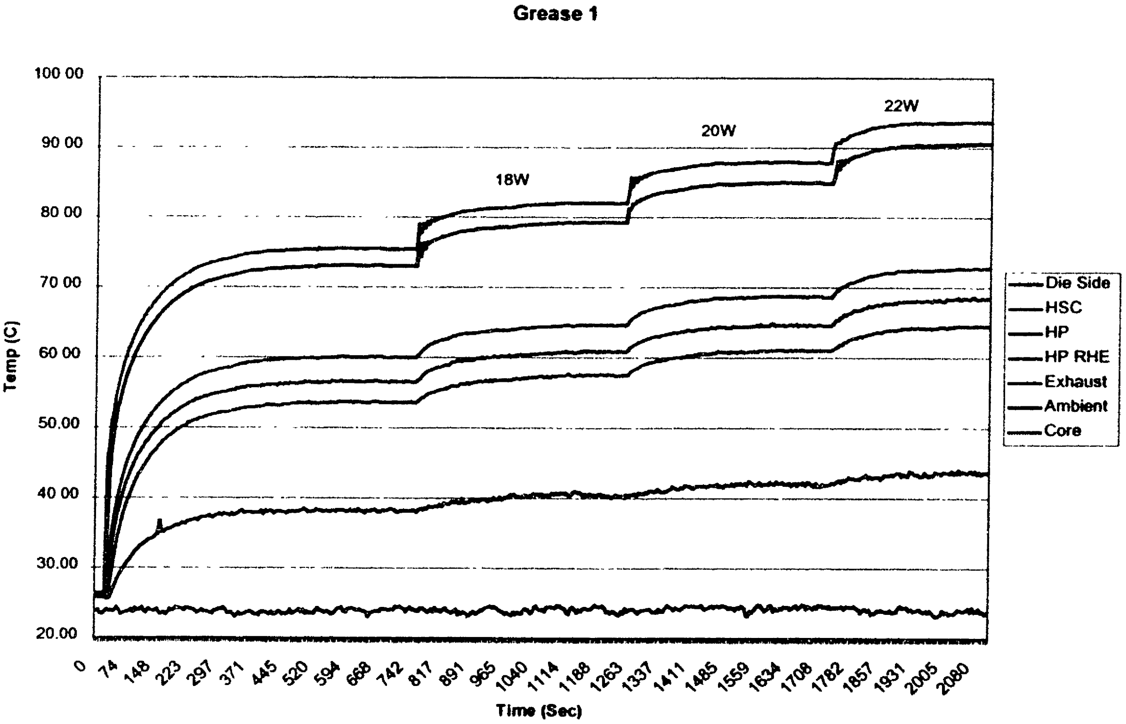

Grease 2

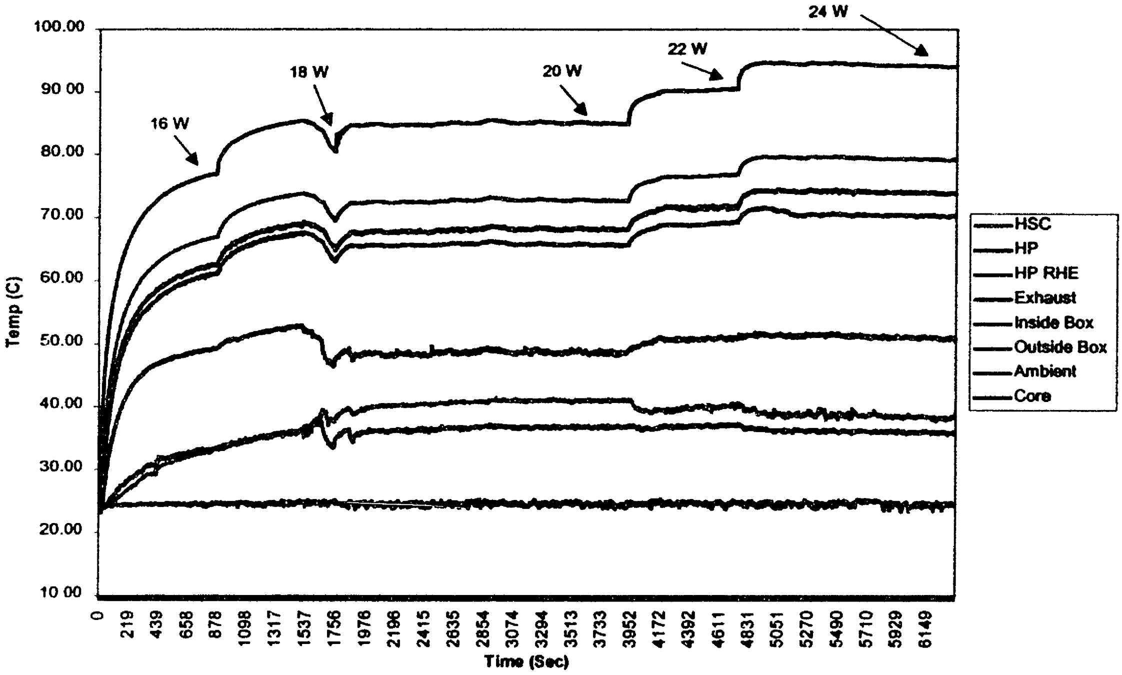

Phase Change 1

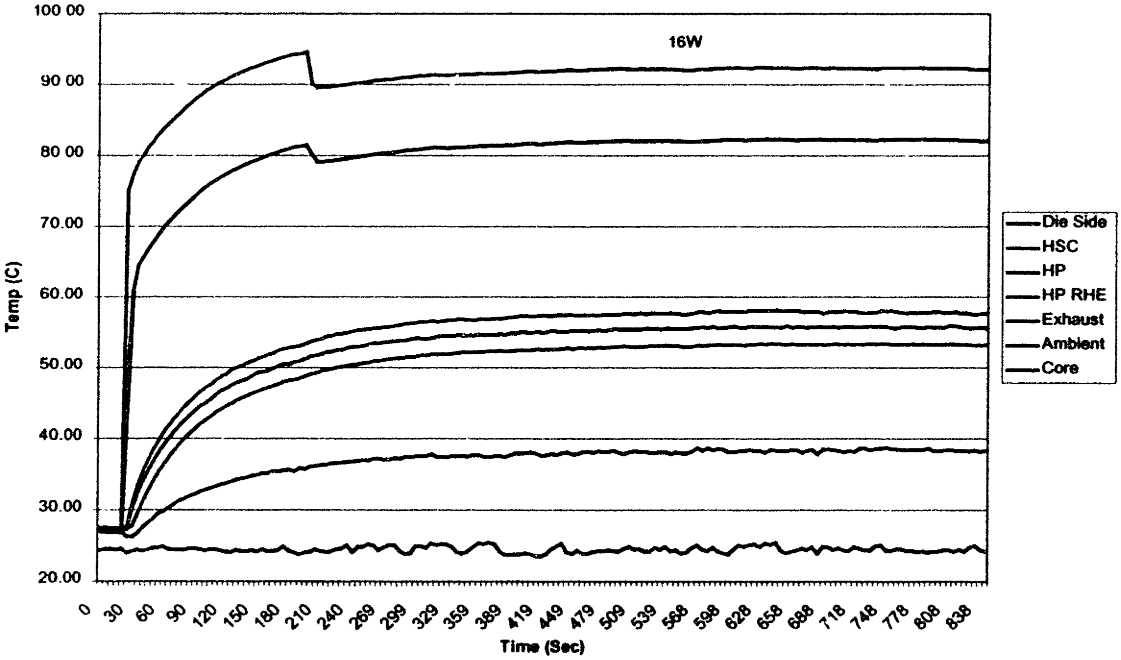

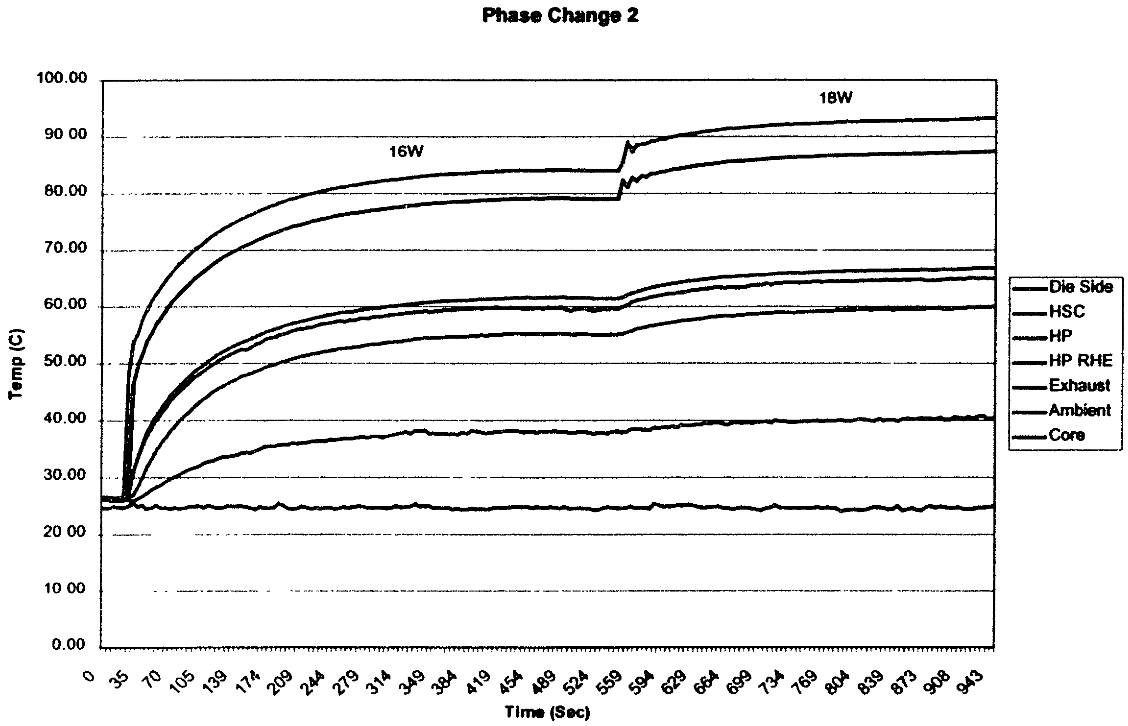

Phase Change 2

8.2 *Second concept data*

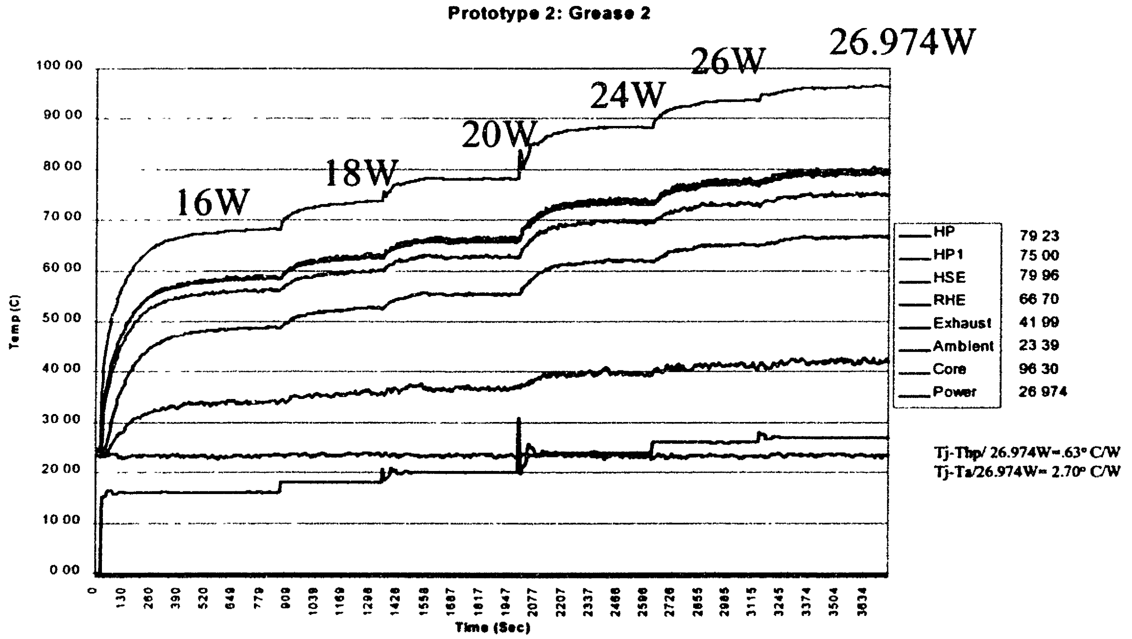

Prototype 2: Grease 2

Tj-Thp/ 26.974W=.63° C/W
Tj-Ta/26.974W= 2.70° C/W

9 Bibliography

1. http://www.intel.com
2. http://mhpg.intel.com
3. http://www-rpd.intel.com
4. Intel Corporation, MCG and RPD marketing.
5. Intel Corporation, conversations with R. Castelino
6. Intel Corporation, Architecture Labs
7. Intel Corporation: Andre Ali, Notebook Thermal Modeling and Simulation, IDF '98 presentation.
8. Intel Corporation: Mobile Mechanical Applications
9. Incropera, F.P. and DeWitt, D.P., *Introduction to Heat Transfer*, John Wiley & Sons, New York, 1996.
10. http://www.furukawa-usa.com/pro/heat/index.htm
11. Bejan, A., "Geometric Optimization of Cooling Techniques". In textbook: Kim, J.S. and Lee, S.W., *Air Cooling Technology for Electronic Equipment*, CRC Press, Boca Raton, 1996.
12. Kordyban, T., *Hot Air Rises and Heat Sinks*, ASME Press, New York, 1998.
13. Intel Corporation: Shawn McEuen, Mobile Platform Technologies
14. Dataquest Personal Computers, 11/98; IDC Personal Systems, 2/99
15. Peters, T. Thriving on Chaos, Alfred A. Knopf, New York: 1988.
16. Rust, R.T. and Zahorik, A. J. "Customer satisfaction, customer retention, and market share", Journal of Retailing. 69(2) 1993.
17. Monroe, J. Winning in High-Tech Markets, Harvard Business School Press, Boston: p125-195, 1993.
18. Faulkner, T.W. "Applying 'options thinking' to R&D", Research-Technology Management. 39(3) 1996.
19. Dataquest Personal Computers, 11/98; IDC Personal Systems, 2/99
20. Intel Corporation: MCG Finance
21. Grindley, P. Standards, Strategy, and Policy, Oxford University Press, (Oxford: 1995)
22. Utterback, J. Mastering the Dynamics of Innovation, Harvard Business School Press, (Boston: 1994)
23. Fine, C. Clockspeed, Perseus Books, (Reading: 1998).
24. Nadler, D. and Tuschmanm M. "Competing By Design: The Power of Organizational Architecture" Oxford University Press: 1997.
25. Larson, E.W. and Gobeli, D.H. "Organizing for Product Development Projects" *Journal of Product Innovation Management*. 5(3): 180-90. 1988.
26. Duncan, R. "What is the right organization structure?" *Organizational Dynamics* 7(3):59-80 (Winter 1979).
27. Crawford, C. M. New Products Management. New York: Irwin, 1986.

THESIS PROCESSING SLIP

FIXED FIELD: ill. ______ name ______

index ______ biblio ______

COPIES: [✓ Archives] Aero Dewey Eng Hum

Lindgren Music Rotch [✓ Science]

TITLE VARIES: ☐ ______

NAME VARIES: ☑ Mohan

IMPRINT: (COPYRIGHT) ______

COLLATION: ______

ADD: DEGREE: S.M. **DEPT.:** Mgt

SUPERVISORS: ______

NOTES: Dewey's copy not available

cat'r ______ date ______

page J13b

DEPT: Mgt. and E

YEAR: 2000 **DEGREE:** S.M.

NAME: PHANSE Vivek