Balanced Technology Extended (BTX) System Design Guide

Version 1.0

IMPORTANT INFORMATION AND DISCLAIMERS

INTEL CORPORATION MAKES NO WARRANTIES WITH REGARD TO THIS BALANCED TECHNOLOGY EXTENDED (BTX) SPECIFICATION ("SPECIFICATION"), AND IN PARTICULAR DOES NOT WARRANT OR REPRESENT THAT THIS SPECIFICATION OR ANY PRODUCTS MADE IN CONFORMANCE WITH IT WILL WORK IN THE INTENDED MANNER. NOR DOES INTEL ASSUME RESPONSIBILITY FOR ANY ERRORS THAT THE SPECIFICATION MAY CONTAIN OR HAVE ANY LIABILITIES OR OBLIGATIONS FOR DAMAGES INCLUDING, BUT NOT LIMITED TO, SPECIAL, INCIDENTAL, INDIRECT, PUNITIVE, OR CONSEQUENTIAL DAMAGES WHETHER ARISING FROM OR IN CONNECTION WITH THE USE OF THIS SPECIFICATION IN ANY WAY.

NO REPRESENTATIONS OR WARRANTIES ARE MADE THAT ANY PRODUCT BASED IN WHOLE OR IN PART ON THE ABOVE SPECIFICATION WILL BE FREE FROM DEFECTS OR SAFE FOR USE FOR ITS INTENDED PURPOSE. ANY PERSON MAKING, USING OR SELLING SUCH PRODUCT DOES SO AT HIS OR HER OWN RISK.

THE USER OF THIS SPECIFICATION HEREBY EXPRESSLY ACKNOWLEDGES THAT THE SPECIFICATION IS PROVIDED AS IS, AND THAT INTEL CORPORATION MAKES NO REPRESENTATIONS, EXTENDS NO WARRANTIES OF ANY KIND, EITHER EXPRESS OR IMPLIED, ORAL OR WRITTEN, INCLUDING ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, OR WARRANTY OR REPRESENTATION THAT THE SPECIFICATION OR ANY PRODUCT OR TECHNOLOGY UTILIZING THE SPECIFICATION OR ANY SUBSET OF THE SPECIFICATION WILL BE FREE FROM ANY CLAIMS OF INFRINGEMENT OF ANY INTELLECTUAL PROPERTY, INCLUDING PATENTS, COPYRIGHT AND TRADE SECRETS NOR DOES INTEL ASSUME ANY OTHER RESPONSIBILITIES WHATSOEVER WITH RESPECT TO THE SPECIFICATION OR SUCH PRODUCTS.

A NON-EXCLUSIVE COPYRIGHT LICENSE IS HEREBY GRANTED TO REPRODUCE THIS SPECIFICATION FOR ANY PURPOSE PROVIDED THIS "IMPORTANT INFORMATION AND DISCLAIMERS" SECTION (PARAGRAPHS 1-4) IS PROVIDED IN WHOLE. NO OTHER LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY OTHER INTELLECTUAL PROPERTY RIGHTS IS GRANTED HEREIN.

Version 1.0, 5 August 2004

Revision History

Version	Description	Date
0.1	Outline	November 2003
0.5	Initial draft	January 2004
0.9	Second draft	April 2004
1.0	Public Release	July 26, 2004

This page is intentionally left blank.

Contents

Equations

Tables

Figures

1. Introduction

The Balanced Technology Extended (BTX) form factor is a scalable form factor that allows for a wide range of system sizes and profiles. There are several benefits associated with a properly designed BTX system over designs from previous form factor specification generations. These benefits include scalability in system design, improvements to system power delivery and power dissipation, acoustics, board layout and routing, high volume manufacturing costs, and structural integrity. Combined, these benefits represent the value propositions of a BTX design. This document will focus on delivering technical content that will highlight design challenges and the associated engineering solutions for a wide range of BTX system sizes, which demonstrate the proposed value propositions.

1.1 Objective

The Balanced Technology Extended (BTX) System Design Guide is intended to be technical reference material that encompasses non-product specific engineering principles as they apply to the BTX form factor. The format in which the material is presented is designed to identify the design challenge, engineering principle, design intent, design trade-offs, and offer engineering solutions. Much of the content is based on learning's from the development and design of the two Intel© BTX system reference designs. However, this Design Guide is not limited to the design of systems similar to the Intel© reference designs; instead it can and should be used to engineer systems of many different sizes, with many different feature combinations, and usage models.

This material should provide a detailed understanding of both individual component or ingredient function and their interdependencies in the overall system design. The content of this System Design Guide will aid in the design and development of systems compliant with the *BTX Interface Specification* that demonstrate the proposed value propositions. Value propositions included in this design guide will include scalability in system design, system power delivery and dissipation, acoustics, HVM cost, and structural integrity.

The BTX System Design Guide is not intended to be an Intel© product or platform based document; however, certain boundary conditions derived from Intel product performance requirements have been assumed in the development of the Intel reference designs and will be discussed. These boundary conditions do not necessarily represent limitations of a BTX system or of its subsystem components or ingredients.

1.2 Reference Documents

Document	Comment
Balance Technology Extended (BTX) Interface Specification	http://www.formfactors.org
Roark's Formulas for Stress and Strain	Young, Warren C., Budynas, Richard G., 7th Ed., 2002
Handbook of Hydraulic Resistances,	I.E. Edelchik, 3rd edition

This page is intentionally left blank.

2. Fluid Dynamic Management

2.1 Airflow Engineering

The intent of this section is to provide guidance on airflow management engineering strategies. It is assumed that the reader will have a reasonable understanding of fluid dynamic engineering principles.

Specifically, this section will address issues and strategies related to impedance and effective management of airflow in a Desktop System.

2.2 Definition of Airflow and Venting Terms

Table 1: Airflow Terminology

Term	Symbol	Units	Definition
Sound Power		BA	Sound energy averaged in all directions from the noise source
Sound Pressure		dBA	Sound energy in one direction from the noise source
Coefficient of Lift	C_L		An indication of the Lift force generated on an airfoil (e.g. fan blade) as it passes through a fluid (air)
Coefficient of Drag	C_D		An indication of the Drag force generated on an airfoil (e.g., fan blade) as it passes through a fluid (air)
Fan			The entire radial fan assembly including the fan motor, fan motor Printed Circuit Board (PCB), bearing tower, impeller, housing, thermistor, and wired connector
Impeller			Fan blades or airfoils that are connected to the fan motor
Stator			Blades or airfoils at he exit of a fan, whose intent is to change the angle of airflow
Duct			A channel created by a housing that directs airflow to a desired location or constrains its path
Volumetric Airflow	W	CFM	The volume flow rate of airflow
Flow Speed	V	LFM	The magnitude of linear airflow velocity
Flow Axial Component	V_{axial}	LFM and $\angle$	The vector component of the airflow velocity that is normal to the face of the fan
Flow Swirl Component	V_{swirl}	LFM and $\angle$	The radial vector component of the airflow velocity that is radial to the face of the fan
Exit Angle	$\angle_{exit}$	degrees	The angle of the airflow that leaves the trailing edge of a fan impeller blade, normal to the fan rotating plane

Term	Symbol	Units	Definition
Angle of Attack	$\angle_{attack}$	degrees	The angle of the airflow that enters a blade or airfoil measured relative to the blade chord
Blade Angle	$\angle_{blade}$	degrees	The angle of the blade or airfoil relative to the rotating axes of the fan
Sweep Angle	$\angle_{sweep}$	degrees	The angle of the blade relative to the normal of the rotating axes of the fan
Chord	L	m	The length of a blade or airfoil from its leading edge (tip) to its trailing edge (tail)
Impedance	dP	"H_2O	Resistance to airflow movement
Near Field Impedance	dP	"H_2O	A source of impedance close to the fan's inlet or exit plane
Far Field Impedance	dP	"H_2O	A source of impedance distant from the fan's inlet or exit plane
Blade Pass Frequency	BPF	Hz	The frequency at which the blades of a rotating fan pass near to another solid object
Heatsink Thermal Resistance	θ_{ca}	°C/W	The resistance to the transfer of heat flux through a heatsink dissipation device, when a uniform heat power is applied
Heatsink Thermal Resistance	ψ_{ca}	°C/W	The resistance to the transfer of heat flux through a heatsink dissipation device, when a non-uniform heat power is applied
Heat Power	Q	W	Power in the form of heat typically generated by efficiency losses from electrical power
Convective Coefficient	h	W/m°K	A description of the heat transfer from the surface of a material
Thermal Conductivity	k	W/m°K	A description of the heat transport within a material
Temperature Difference	ΔT	°C	A description of the difference in temperature at two different locations
Reynolds Number	Re		A dimensionless indication of the airflow regime, which is the ratio between the fluid's internal viscous to inertial forces
Prandtl Number	Pr		The ratio of fluid kinematic viscosity to thermal diffusivity
Thermal Diffusivity	α	m^2/s	A fluid heat transport property
Kinematic Viscosity	ν	m^2/s	A fluid momentum transfer property
Density	ρ	Kg/m^3	Fluid density.
Specific heat	Cp	J/kg °K	Specific heat of the fluid.
Case Temperature	T_{case}	°C	The temperature of the surface of a component
Ambient Temperature	$T_{ambient}$	°C	The temperature of surrounding air
Junction Temperature	$T_{junction}$	°C	The temperature of the silicon of a micro-electronic device
Fan Speed	ω	RPM	The rotational speed of a fan

2.3 Concurrent Design Engineering

The BTX form factor was concurrently developed and engineered by Intel$^{®}$ motherboard and system engineering groups. Motherboard component placement, therefore, takes into account not only the routing and electrical performance requirements, but also the heat transfer and airflow management required to ensure a proper operating environment.

Intel understands the migration from an existing to a new form factor standard is one that will create transition challenges, in addition to new opportunities for innovation. The new constraints imposed by the BTX form factor are ones that Intel anticipates will, in fact, lead to the use of the Desktop Personal Computer in more environments and usage models than is allowed by the existing form factors.

Intel has engineered BTX with the anticipation that all components will continue to evolve and that new components will need to be integrated. For instance, the motherboard core area – processor, chipset, and memory – has been increased to not only decrease component and routing density on today's platform but also for the anticipated growth in component power, I/O count, I/O density, and component physical size. By providing this additional board real estate, it is anticipated that standard Desktop systems will to continue to use four-layer motherboard technology and cost effective component packaging technologies.

2.4 BTX Airflow Management Strategy

2.4.1 Provide the High Power Components with Low Temperature Air at High Velocity

The benefits of providing low temperature air at high velocity to the high power components include lower heatsink costs, improvement in acoustic performance, and the potential for improved component performance.

For a heatsink attached to a component with Heat Flux (Q) and a Package Case Temperature (T_{case}) specified, the heatsink resistance requirement is defined at the Design Ambient Temperature ($T_{ambient}$) as:

Equation 1: Thermal Characterization Parameter (Case-to-Ambient)

$$\psi_{ca} = [T_{case} - T_{ambient}] / Q$$

Similarly, the component may have a silicon junction temperature requirement instead of a case temperature requirement. In that case the heatsink resistance requirement is defined as:

Equation 2: Thermal Characterization Parameter (Junction-to-Ambient)

$$\psi_{ja} = [T_{junction} - T_{ambient}] / Q$$

As can be seen from the governing equation, for a given power into a heatsink and a given component temperature requirement (either T_{case} or $T_{junction}$), lower ambient temperature ($T_{ambient}$) increases the required heatsink resistance, as shown in Figure 1.

Figure 1: Requirement versus Ambient Temperature
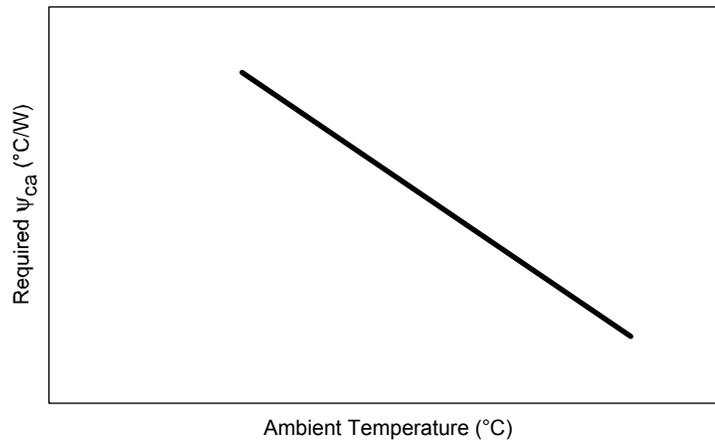
OM16706

Higher resistance heatsinks are generally easier and cheaper to manufacture. The relationship between $T_{ambient}$ and heatsink cost is shown in Figure 2 and illustrates why BTX was designed to provide lower temperature air to all high power Desktop system components.

Figure 2: Heatsink Cost versus Ambient Temperature
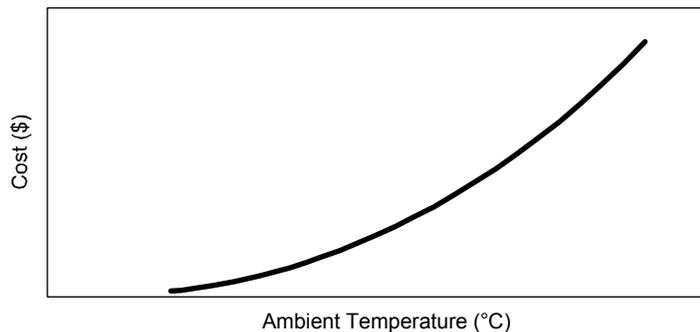
OM16707

Intel provides the required junction or case temperature specification for its processor and chipset components in component datasheets. Intel also provides guidance on designing and measuring the performance of heatsinks for these components in the component design requirements documents and thermal design guides.

Heatsink performance can be defined in terms of its conduction and convection performance. Conduction performance is primarily determined by the selection of heatsink materials (thermal conductivity, k) and design features such as its cross-sectional conduction area (A_c) and the length of the conduction path (L), as illustrated in Equation 3 by the governing equations for one-direction conduction.

Equation 3: Heat Conductive Performance

$$Q = k \, A_c \, \Delta T / \Delta L$$

The governing equations for convection will be used to illustrate the strong relationship between the velocity of the air that moves through the heatsink fin channels and the convection performance.

First, the amount of heat that can be dissipated by any convective surface is a function of the surface area from which it is being conducted (A), the temperature difference between the surface and the surrounding air (ΔT), and the convective coefficient (h).

Equation 4: Heat Convective Performance

$$Q = h \, A \, \Delta T = h \, A \, (T_{surface} - T_{ambient})$$

This equation illustrates the impact of lowering the ambient temperature. It is desirable to decrease $T_{ambient}$ because it increases the amount of heat (Q) the surface can dissipate while maintaining a target surface temperature. Alternately, decreasing $T_{ambient}$ can be used to lower the surface temperature for the targeted heat dissipation (Figure 3) while using the same thermal ingredient design.

Figure 3: Effect of Decreasing Ambient Temperature
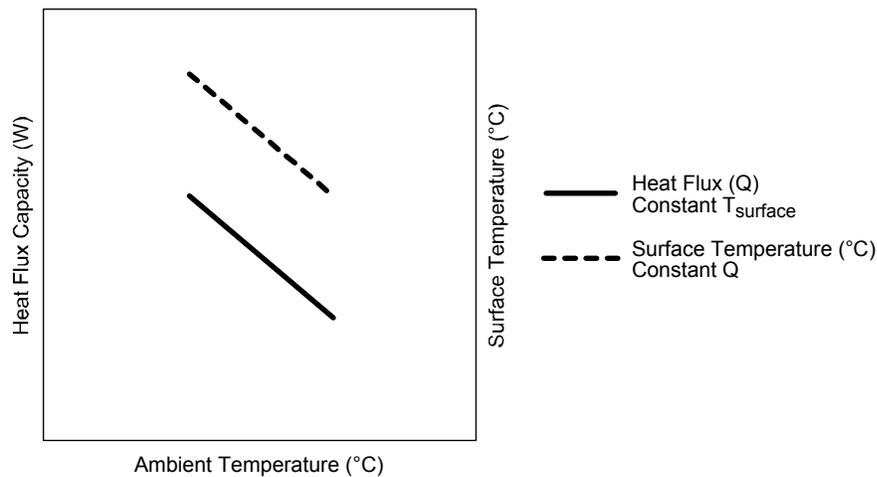
OM16708

Equation 4 also illustrates the impact of increasing the convective coefficient, h. It is desirable to increase h because it increases the amount of heat (Q) the surface can dissipate while maintaining a target surface temperature. Alternately, increasing h can be used to lower the surface temperature for the targeted heat dissipation (Figure 4).

Figure 4: Effect of Increasing Convective Coefficient, *h*
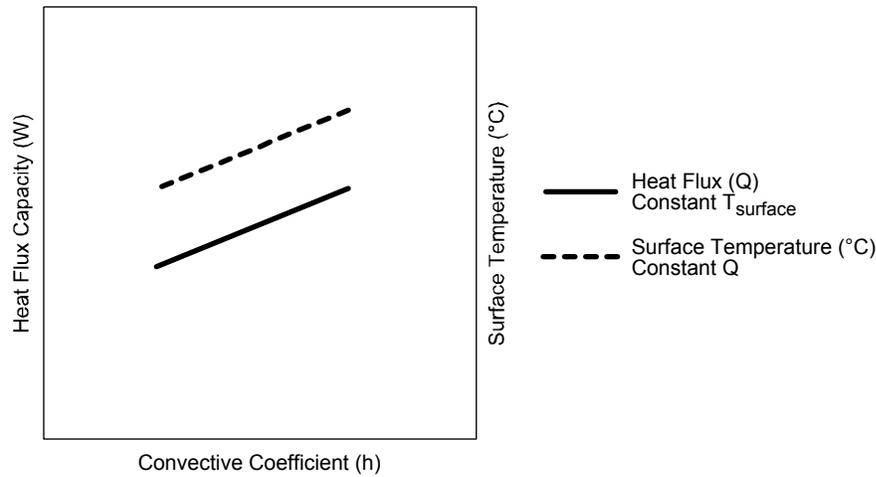
The convective coefficient, *h*, is a function of the surface geometry, material conductivity (*k*), Reynolds Number (Re), and Prandtl Number (Pr):

Equation 5: Convective Coefficient

$$h = 0.664 \, k \, \mathrm{Re}^{1/2} \, \mathrm{Pr}^{1/3} / L$$

Note: The constant, 0.664, in Equation 5 is a function of the flow regime; this value increases when the flow transitions from the laminar to the turbulent regime. The flow regime in most areas of a Desktop system will be laminar, so the use of this equation is appropriate.

This equation is limited in application to the case where the air is flowing over a flat plate (for instance, over the surface of the motherboard, add-in card, or drive bay). For flow within a narrow channel, the boundary layer development will limit the improvement in convection coefficient due to increasing velocity.

For reference, the Prandtl Number is a function of the air viscosity (ν) and the thermal diffusivity of air (α):

Equation 6: Prandtl Number

$$\mathrm{Pr} = \nu / \alpha$$

The Reynolds Number is a function of airflow velocity (V), the length of the surface (L), and air viscosity (ν). The Reynolds Number for flow over a horizontal surface (not the flow within a channel) can be calculated using the following equation:

Equation 7: Reynolds Number

$$\mathrm{Re} = V \, L / \nu$$

From Equation 5 and Equation 7, it can be demonstrated that the convective coefficient, *h*, is proportional to airflow velocity. Therefore, the convective heat transfer capability of a surface increases with an increase in velocity. Generally, it can be stated that heatsink resistance, ψ, decreases with increasing velocity.

At a target surface temperature, the heat power that can be dissipated increases with increasing airflow velocity through a heatsink. Alternately, at a target heat power, the surface temperature decreases with increasing airflow speed.

Figure 5: Effect of Increasing Airflow Velocity
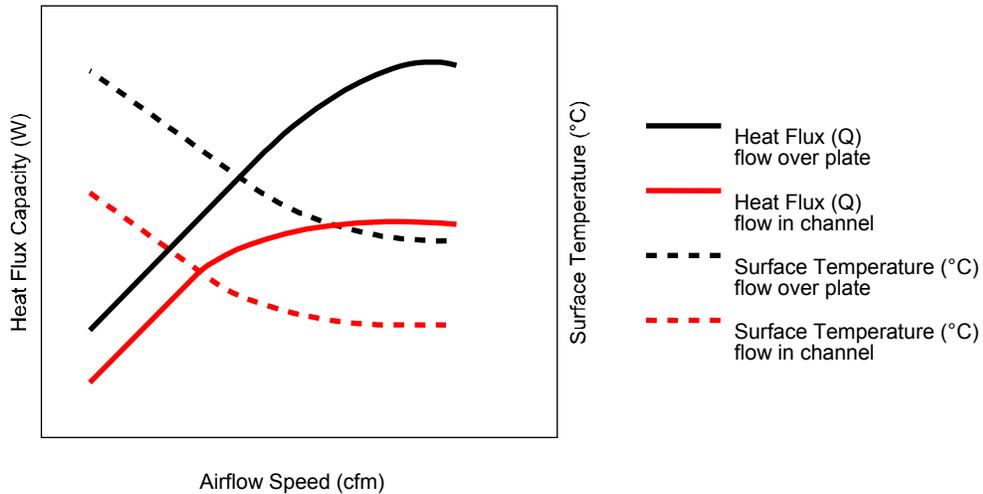
OM16710

In summary, the performance of every convective surface is a function of the ambient airflow temperature and velocity. Lower airflow temperature increases ΔT and higher airflow velocity increases the convection coefficient. The BTX has been engineered to reduce the ambient temperature and increase the airflow velocity in the region of all high power components.

2.4.2 Minimize the Total System Impedance

Impedance is the resistance to movement that a moving fluid encounters. Generally, low impedance is more desirable than high impedance because it allows the air movers (e.g., fans) to provide higher airflow at each operating speed.

In Desktop systems, there are impedance sources that are more preferable than others. For instance, when air is flowing through a channel, the walls of that channel represent impedance to the airflow, but if the fins of a heatsink create that channel then the air is also working to remove heat. In other words, impedance introduced to remove heat or improve heat transfer performance is "good" impedance. Examples of undesirable impedance include airflow turning or forcing air through narrow channels and expansions or constrictions that do not remove heat from components.

A network resistance diagram should be constructed for every Desktop system design, illustrating the parallel and series airflow impedances from the airflow entrance to its exit. The impedance (resistance) characteristics of each subsystem can be estimated using hand calculations, or loss coefficient estimates from fluid flow handbooks, measured in a wind tunnel, or predicted using fluid dynamic numerical tools. A typical, simple resistance network diagram for a Desktop system is shown in Figure 6.

Figure 6: Typical Network Resistance Diagram

OM16711
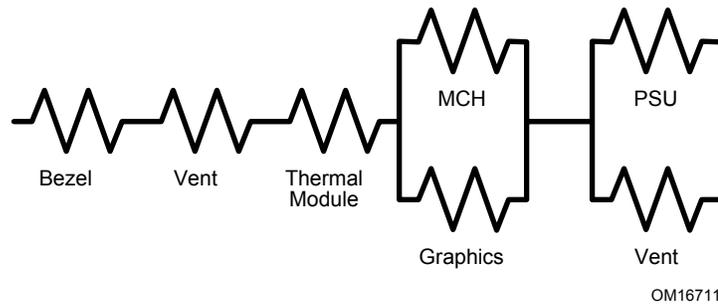
The total resistance to airflow in a Desktop system determines that amount of airflow that the various fans can deliver. Lower impedance systems allow more airflow to pass through them and, as seen from the previous section, this will improve the system's heat transfer performance.

Acoustic performance is strongly correlated to the system impedance. That is, for any required system airflow, higher system impedance will require higher fan speeds and, therefore, create more acoustic noise. The relationship between impedance and the number and operating speeds of the fans is described in Section 1.

2.4.3 Provide Above-board and Under-board Airflow

There are several sources of heat within and directly attached to a motherboard. Metal traces and planes in the motherboard generate heat during the conduction of electrical current. Surface Mount components directly attached to the motherboard also generate and transfer heat into the motherboard. Motherboard reliability is compromised when its temperature exceeds operating specifications; therefore, removing heat from the motherboard is an important part of a Desktop system's thermal design.

By providing airflow above and below the motherboard, two convective heat transfer paths from the motherboard are established (effectively doubling the available convective heat transfer surface area relative to that in an ATX system), as shown in Figure 7. The improvement in total heat transfer capability from the motherboard allows higher power components to be attached to the motherboard without compromising motherboard reliability.

Figure 7: Two Heat Transfer Paths from Motherboard

OM16712
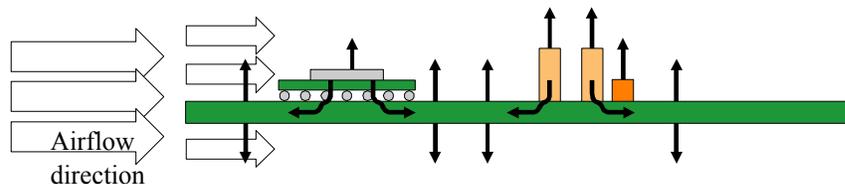
2.4.4 System Airflow Direction

The three important airflow management objectives identified in Sections 2.4.1, 2.4.2, and 2.4.3 were used by Intel system and motherboard engineers in the placement of high power components in a single airflow stream generated and managed by a minimum number of fans. With this airflow management and component placement strategy outlined, the next challenge was to determine the most appropriate direction for that airflow.

Creating a single airflow stream direction implies that the Desktop system must be designed with the inlet and exhaust at opposite sides of the chassis.

The path from the chassis pan to top cover was not considered feasible because airflow would impinge on the motherboard, as shown in Figure 8. The associated turning impedance is not "good" impedance, so this flow direction was not selected.

Figure 8: Top-to-Bottom Airflow Illustration – Impinging Flow at Motherboard
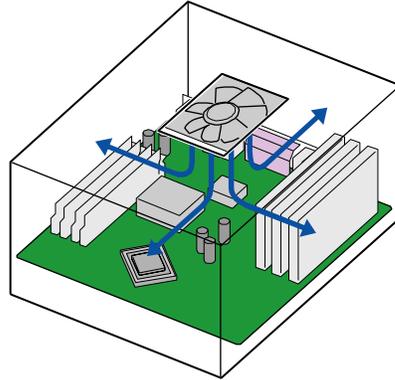
OM16713

Since it is common in Desktop system design to have a single chassis design operate in desktop and tower orientations, inlet and exhaust from side-to-side or top-to-bottom would have the inlet or exhaust blocked by the surface on which the chassis rests (Figure 9); therefore, these airflow directions were not selected.

Figure 9: Side-to-Side Airflow Illustration – Exit Flow Blocked When System Is Rotated
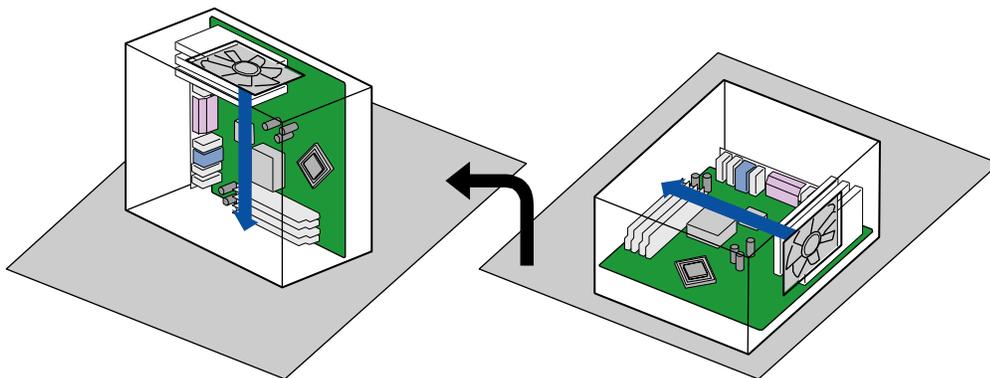
OM16714

Finally, exhausting airflow out of the front panel of the chassis directs warm air exhausted from the system toward the user, as shown in Figure 10.

Figure 10: Back-to-Front Airflow Illustration – Exit Airflow Directed at User
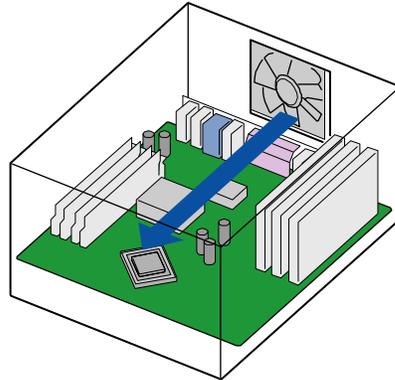
OM16715

Therefore, Intel designed BTX to use front-to-back airflow with the high power components aligned so that a single airflow stream services them all. Placing a fan at the front of the system creates unique acoustic concerns because the fan is now a noise source that is directly in front of the system user. Acoustic management strategies for a noise source at the front of the system are discussed in the Acoustic Engineering section of this design guide.

2.4.4.1 CPU Location

The highest power component with the most stringent temperature specifications is the Central Processing Unit (CPU), so it is located in the front of the motherboard, near the chassis front panel (typical component locations on a BTX motherboard are illustrated in Figure 11). In this position, the CPU will receive the lowest ambient temperature air moving at its highest speed, as described in the front-to-back airflow management strategy outlined above.

Figure 11: Typical Component Locations on a BTX Motherboard
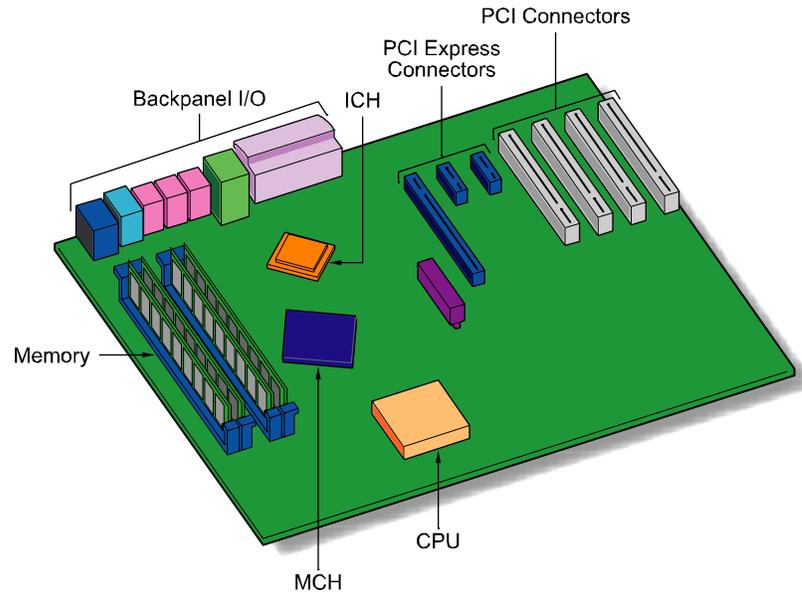
OM16716

The fan in the BTX Thermal Module provides the airflow for the CPU heatsink and the remainder of the system. Figure 12 illustrates the position of the Thermal Module relative to the CPU.

Figure 12: Thermal Module Location Relative to the CPU
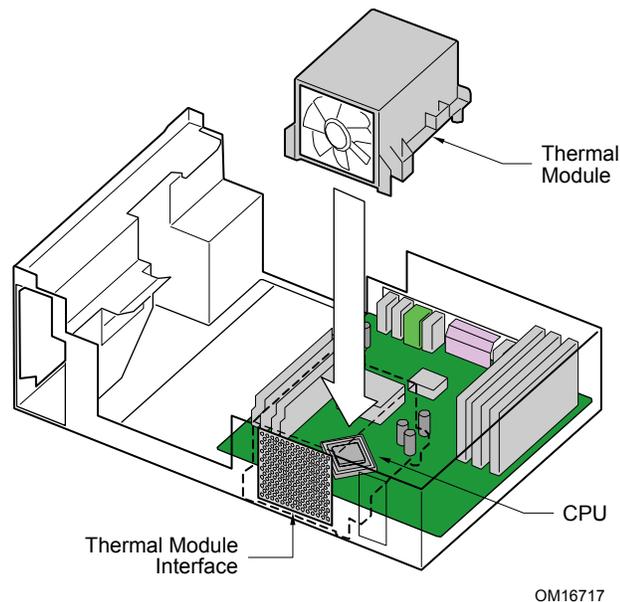
OM16717

The *BTX Interface Specification* requires that the Thermal Module Interface be designed in such a way as to eliminate the potential for air that has entered the system to re-enter the Thermal Module inlet. Eliminating recirculation ensures that the CPU heatsink receives the lowest temperature air, which comes from outside the chassis. It also ensures that the temperature of Thermal Module exit airflow (which

provides the airflow for the remainder of the system) is not increased by recirculation. Figure 12 also illustrates a Thermal Module Interface provided by a vented portion of the chassis sheet metal front panel.

The Thermal Module supplier is responsible for providing a fan, in addition to the CPU heatsink and duct assembly. The Thermal Module exhaust air – which exits the CPU heatsink and duct assembly – provides the airflow for the remainder of the Desktop system components. Figure 13 illustrates the primary airflow path generated by and exiting from the Thermal Module fan.

Figure 13: Thermal Module Airflow Path
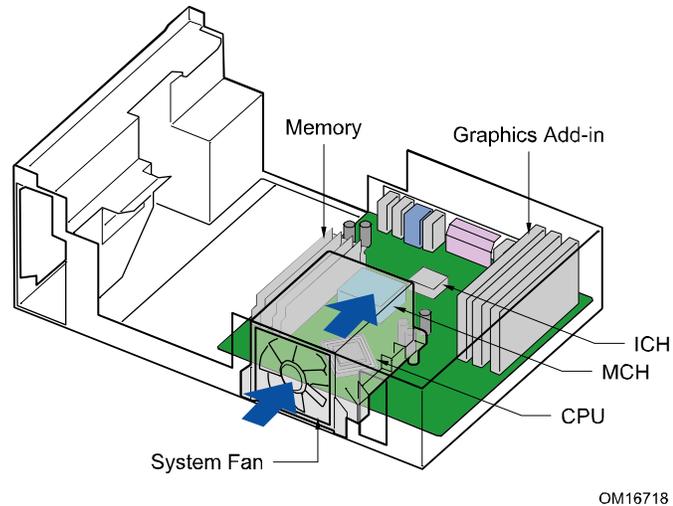
With a properly selected Thermal Module, the exhaust airflow will provide low temperature air at high velocity to the remainder of the system components. The importance of selecting a Thermal Module with an appropriate effective fan curve is discussed in Section 2.4.5.1.

The location of the CPU at the front of the system is beneficial not only to the CPU heatsink performance, but also to the CPU socket performance. The CPU socket generates heat when it is delivering the operating current required by the CPU. In BTX, the temperature of the socket and motherboard are reduced because there is low temperature, high velocity airflow directed above the motherboard at the socket and because the bottom of the motherboard is a convective surface due to the under-board airflow (refer to Figure 7).

Comparison to ATX: ATX CPU thermal solutions typically have a dedicated fan directing airflow through a heatsink. This impinging flow configuration forces airflow to turn at the motherboard, which introduces substantial undesirable impedance to the CPU heatsink fan. This increase in impedance decreases the velocity of the air through and exiting the ATX CPU heatsink. The exhaust airflow path is constrained by the location of memory and the first expansion slot, forcing the exhaust air to recirculate into the CPU heatsink fan. The increase in CPU heatsink inlet temperature decreases its heat transfer capability. The increase in exhaust airflow temperature also decreases the heat transfer efficiency around the socket and motherboard. A BTX Thermal Module will not have turning impedance or recirculation; therefore, the CPU heatsink and socket-motherboard heat transfer efficiency will be higher.

Refer to Section 2.4.5.1 for more information on the Thermal Module.

2.4.4.2 CPU Voltage Regulation Location

CPU Voltage Regulation component performance is typically constrained by the Voltage Regulation component or motherboard temperature specifications. The electrical efficiency losses that result from the management of the CPU operating current create heat in the Voltage Regulation components and in the motherboard traces and power planes.

CPU Voltage Regulation is typically designed using multiple phases; each phase is comprised of Voltage Regulation components (drivers, capacitors, inductors, etc.) and is responsible for managing a portion of the CPU operating current. When the Voltage Regulation component or motherboard temperature specifications are exceeded because the heat associated with the phase's operating current, the following design options are typically considered:

a. Attach heatsinks to the Voltage Regulation components to improve their convective heat transfer performance and thereby lower their operating temperature.
b. Provide more airflow to Voltage Regulation components to improve their convective heat transfer performance and thereby lower their operating temperature.
c. Add a phase to the Voltage Regulation design to reduce the operating current per phase and thereby reduce the amount of heat generated per phase.
d. Redistribute the current going to each phase such that it is not equally distributed.

The BTX Thermal Module provides above and below motherboard airflow at low temperature. This airflow pattern creates the heat transfer illustrated in Figure 6. CPU Voltage Regulation components are typically placed near the CPU socket. The greatest heat transfer benefit will be realized if these components are placed in front of the CPU socket, near the front of the BTX motherboard because they will receive low temperature air directly from the Thermal Module fan.

Comparison to ATX: The exhaust airflow from an ATX CPU thermal solution typically provides the airflow for cooling CPU Voltage Regulation and the motherboard. As discussed in Section 2.4.4.1, this impinging flow configuration decreases the velocity and increases the temperature of the heatsink's exhaust airflow. This low velocity, high temperature airflow on only one side of the motherboard does not create an efficient heat transfer path for CPU Voltage Regulation or the motherboard. A BTX Thermal Module provides the same high velocity, low temperature airflow to the CPU Voltage Regulation, socket, and motherboard as it does to the CPU heatsink. It provides this airflow above and below the motherboard, creating a very efficient heat transfer path.

For more information on engineering CPU Voltage Regulation airflow, refer to Sections 2.4.5.2 and 2.4.5.3.

2.4.4.3 Graphics Location

Integrated graphics is often included on Intel® Memory Controller Hub (MCH) chipsets. Refer to Section 2.4.6 for a description of the MCH airflow environment.

A graphics add-in card will typically be located in the first add-in card slot position. In BTX, the graphics add-in card slot position was shifted to the right side of the core so that the powered components of the card would be exposed to the exhaust airflow of the BTX Thermal Module. The card is aligned with the exhaust airflow direction so that it does not create unnecessary system impedance.

Whether the card is placed directly in the first slot (perpendicular to the motherboard) or on a riser (parallel to the motherboard), the powered components of the add-in card – the Graphics Processing Unit, Voltage Regulation, and memory – are exposed to the Thermal Module exhaust airflow.

Comparison to ATX: In ATX, the powered components on a graphics add-in card are on the side opposite the CPU, and are facing into a very narrow channel created by the adjacent card. There typically is very little system-generated airflow on this side of the card and it is not uncommon for a performance ATX graphics add-in card to include a thermal solution with a heatsink and dedicated fan. The high velocity, low temperature exhaust airflow provided by the BTX Thermal Module moves past the powered side of a graphics add-in card, significantly improving the heat transfer performance of the card's thermal solution. The improvement allows lower graphics card heatsink cost, improved operating temperature, or higher power.

For more information on engineering graphics add-in card airflow, refer to Section 2.4.9.

2.4.4.4 Memory Controller Hub (MCH) Location

In BTX, the MCH is located immediately behind the CPU on the motherboard. In this position it will receive the high velocity, low temperature Thermal Module exhaust airflow. That airflow direction will be consistent in every BTX system so the MCH heatsink fins can be continuous instead of crosscut.

In addition, a conductive heat transfer path typically exists between the MCH and motherboard - the MCH solder joints conduct heat from the MCH silicon and package into the motherboard. This heat flow into the board is eventually transferred by convection from the topside of the board but generally increases the board temperature. The above-board and under-board airflow provided in a BTX system provides two convective heat transfer paths for this heat to leave the motherboard.

Comparison to ATX: The exhaust airflow from an ATX CPU thermal solution typically provides a portion of the airflow for the MCH heatsink. This impinging flow configuration decreases the velocity and increases the temperature of the CPU heatsink's exhaust airflow. An ATX system fan may also provide a portion of the airflow for the MCH heatsink, but the direction of that airflow is not consistent in all ATX system designs; therefore, an MCH heatsink is typically crosscut. The high velocity, low temperature exhaust airflow provided by the BTX Thermal Module directly to the MCH heatsink significantly improves its heat transfer performance. And the below board airflow offers another MCH heat transfer path. These improvements in heat transfer allow lower cost MCH heatsinks, improved operating temperature, or higher power.

For more information on engineering MCH airflow, refer to Section 2.4.6.

2.4.4.5 Memory Location

Much the same as the graphics add-in card slot, memory is aligned with the Thermal Module airflow to minimize the impedance it introduces. Thermal Module exhaust airflow will be drawn into and through the memory channels by the power supply fan, if the PSU is located near the memory fixed motherboard edge. In addition to Thermal Module exhaust airflow, memory airflow can be augmented by the airflow that exhausts below the CPU heatsink from the CPU Voltage Regulation. A BTX Thermal Module should be designed such that the below-heatsink airflow path is not fully ducted from the heatsink base to the motherboard. The below-heatsink airflow that is directed toward memory by the socket is an important system airflow path.

Comparison to ATX: The exhaust airflow from an ATX CPU thermal solution typically provides a portion of the airflow for memory. An ATX system fan may also provide a portion of the memory airflow but the direction of that airflow is not consistent in all ATX system designs. BTX Thermal Module and CPU Voltage Regulation exhaust airflow will provide consistent inlet flow to the memory airflow; however, managing memory airflow bypass and exit is important to ensuring an appropriate thermal environment.

For more information on engineering memory airflow, refer to Section 2.4.9.

2.4.5 Subsystem and Component Airflow Management Strategies

2.4.5.1 Thermal Module Role

The Thermal Module fan is the primary air mover in a BTX system. Not only does the Thermal Module fan supply airflow directly to the CPU heatsink, CPU socket, and CPU Voltage Regulation components, its exhaust airflow is used as a medium to provide the heat transfer for the remaining subsystems and components.

In fact, a typical BTX chassis will not ship with any fans installed since the system airflow will be provided by the Thermal Module and PSU fans. The Thermal Module fan will be provided by the Thermal Module supplier and the PSU fan will be provided by the PSU supplier.

2.4.5.1.1   Thermal Module Effective Fan Curve

The Thermal Module exit airflow characteristics are described by its effective Thermal Module fan curve, which is a function of the stand-alone fan curve, the CPU heatsink impedance, CPU Voltage Regulation impedance, CPU socket impedance, and any near-field impedance effects.

Measuring the volume of airflow that a fan can generate against increasing impedance generates a stand-alone fan curve. This measurement is typically conducted in a wind tunnel. The stand-alone fan curve varies with its operating speed, as shown in Figure 14. All fan curves reflect the following performance characteristic: at a particular fan operating speed, as the impedance in a fan's airflow path increases, the amount of airflow that the fan can generate decreases. A stand-alone fan curve at the maximum operating speed is typically in the fan performance data sheet provided by the fan supplier.

Figure 14: Fan Performance Curve versus Fan Speed (RPM)
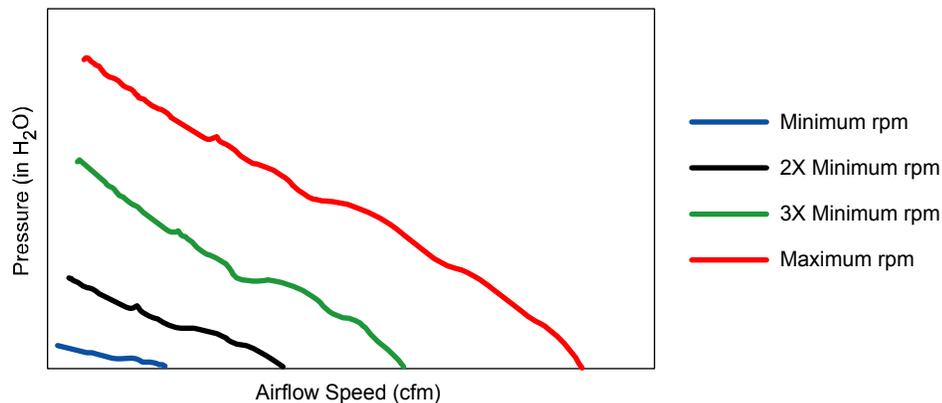
Fans typically include a closed-loop feedback design that monitors and maintains the requested speed. The impact of the closed-loop feedback circuit on the fan curve is best described using the following sequence of events:

a.   A fan converts electrical power into the momentum of the air it is attempting to move. The more difficult the air is to move, the more electrical power is required to impart that momentum.

b.   Impedance at the exit of a fan increases the difficulty of imparting momentum to the air, because the increase in exit impedance increases the pressure in the fan. The increased pressure in the fan makes it more difficult for the impeller to move through and impart momentum to the air.

c. If the electrical power input to a fan is fixed, then increasing the exit impedance causes the fan to slow down. That is, since it finds the air more difficult to move, the rotation of the fan is slowed.

d. The fan's closed loop feedback circuit has a tachometer that monitors fan speed. When higher impedance slows a fan down, the circuit will attempt to return the fan speed to the requested speed. It attempts to do this by increasing the electrical current supplied to the fan motor windings.

e. If the fan feedback circuit request for more electrical current can be satisfied, the fan's operating speed will be maintained, even as impedance increases. In an operating regime where the higher current request can be met, the fan will operate at constant speed. The fan curve in this regime is termed an Iso-RPM fan curve.

   Note: Even if the fan is able to maintain its operating speed against increasing impedance, the amount of airflow that can be generated decreases with increasing impedance. That is, the impedance affects the volume rate of air that the fan can move even if the fan speed can be maintained. This is the performance characteristic reflected in the fan performance curve (see Figure 14).

f. The current required to operate a fan increases as the requested fan speed increases and as the impedance increases.

g. Therefore, at higher operating speed and higher impedance, the current requested by the fan motor may eventually exceed the current that can be provided.

For instance, when the fan is operating at its maximum speed against low impedance it will draw electrical current that may be near the current limit of the electrical circuit. As the impedance increases, the fan will request more electrical current in an attempt to maintain the operating speed. However, once the current limit of the circuit is reached, the fan will no longer be able to maintain the operating speed. If the impedance continues to increase, the fan performance curve will reflect the decrease in airflow volume rate (CFM – cubic feet per minute) from both the impedance increase and the reduction in fan speed.

Figure 15 shows the impact of the reduction in fan speed from this constrained electrical current operating condition. One curve shows the fan curve if electrical current were unconstrained; the other shows the reduction in CFM after the current limit is reached. The fan curve in this current limited operating condition is called an iso-power fan curve.

Figure 15: Iso-RPM versus Iso-Power Fan Performance Curve
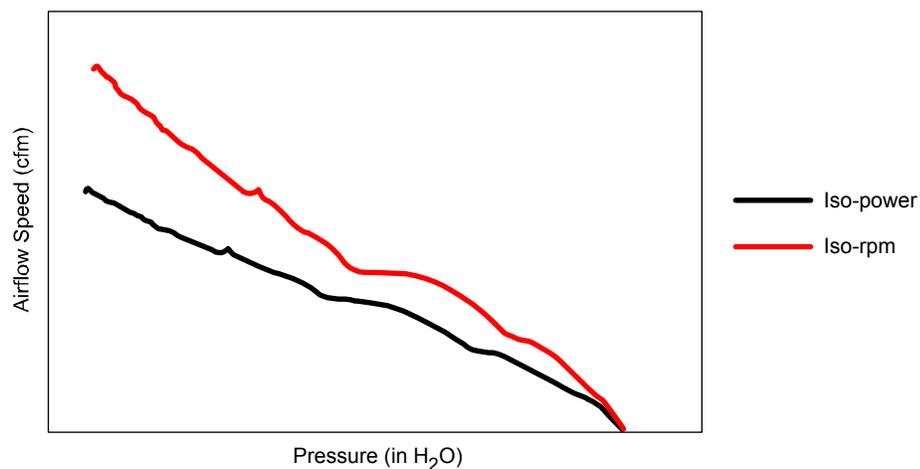
OM16720

The impedance of any particular airflow path changes with the amount of airflow being forced into it. This characteristic of the airflow path is called its impedance curve. An impedance curve can be determined through fluid dynamic numerical modeling applications or wind tunnel measurement. A typical impedance curve is illustrated in Figure 16.

Figure 16: Typical Impedance Curve
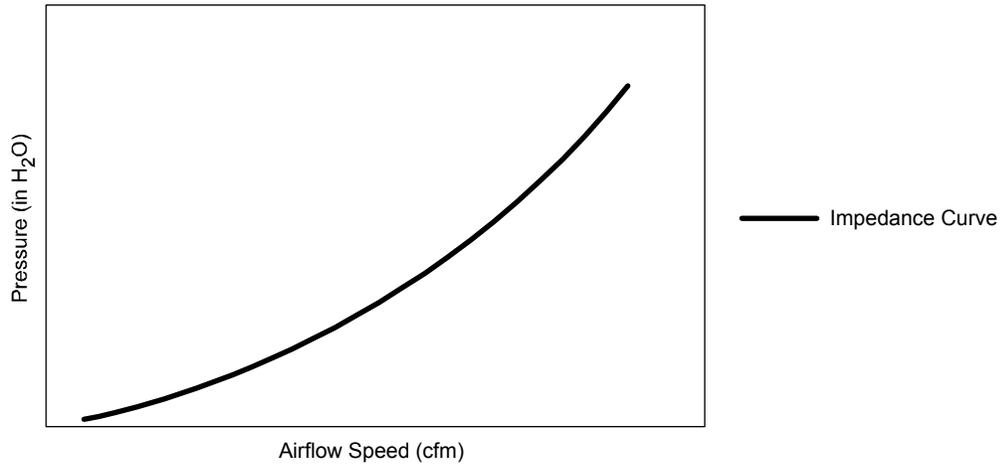
OM16721

The Thermal Module effective fan curve is an indication of the airflow that the Thermal Module can provide a system after the airflow has made its way through the CPU heatsink, Voltage Regulation, and socket. This effective fan curve reflects the impact that the CPU heatsink and Thermal Module duct have on the stand-alone fan curve. Theoretically, if the impedance of the CPU heatsink and duct were measured in a wind tunnel then subtracted from the measured or supplier-provided stand-alone fan curve, the result would be the Thermal Module effective fan curve. This is illustrated in Figure 17.

Figure 17: Theoretical Thermal Module Effective Fan Curve Calculated From Stand-alone Fan Curve and Heatsink Impedance Curve
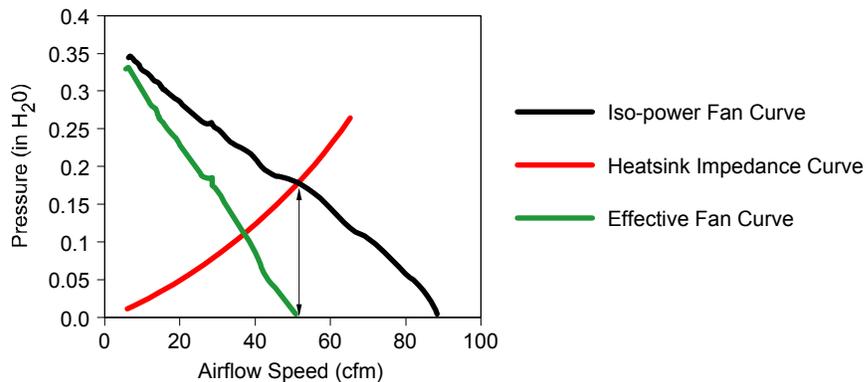
OM16722

However, the near-field impact that the CPU heatsink has on the stand-alone fan curve is difficult to accurately predict from the stand-alone fan curve and the heatsink impedance curve. This is because a wind tunnel can measure only the axial flow characteristics; whereas a rotating radial fan creates airflow

with axial and swirl components. When the swirl component of the fan airflow enters the heatsink, it creates additional impedance that is not present in the axial wind tunnel measurement of the fan curve or the heatsink impedance. It is best to measure the Thermal Module effective fan curve by inserting the entire Thermal Module (fan, heatsink, and duct on a test board) into the wind tunnel.

Increasing the axial component and minimizing the swirl component of airflow that exits the Thermal Module fan can minimize the impact that the CPU heatsink has on the effective fan curve. The Thermal Module supplier can select or engineer the fan impeller to increase the axial component of airflow, and select or engineer a stator that converts a portion of the swirl component to axial. The improvement in the stand-alone fan curve through this type of fluid dynamic engineering is illustrated in Figure 18.

Figure 18: Fan Curve Impact from Impeller and Stator Fluid Dynamic Engineering
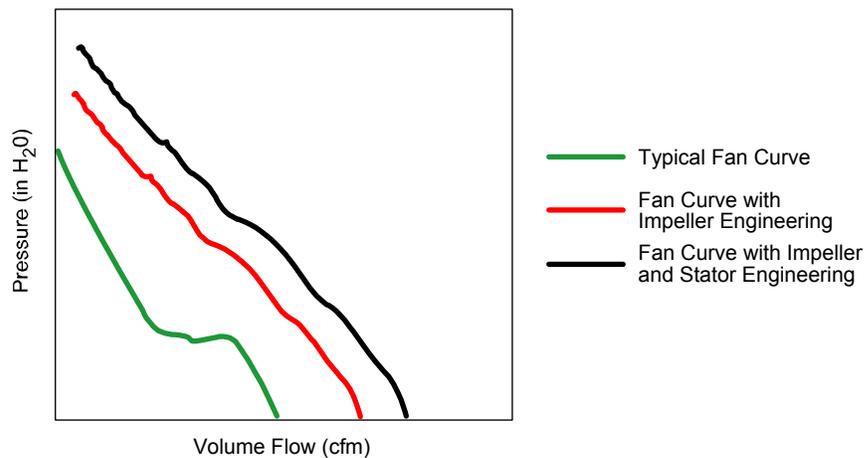
OM16723

2.4.5.1.2   Primary Airflow Paths Within Thermal Module

The Thermal Module is responsible for partitioning the airflow to three distinct airflow paths: through the ducted CPU heatsink, above the motherboard but below the CPU heatsink base, and below the motherboard (see Figure 19).

Figure 19: Thermal Module Airflow Partitioning
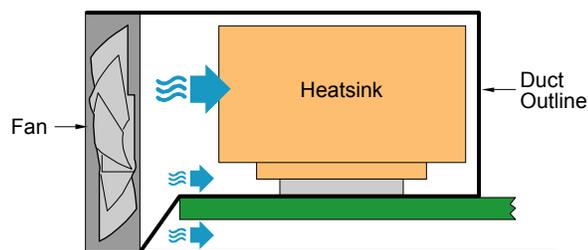
OM16724

The network resistance diagram in Figure 6 can be refined to show the partitioning of these airflow paths in the Thermal Module, as illustrated in Figure 20. This is an excellent method for identifying and managing airflow partitioning and the specific airflow volume through each path.

Figure 20: Thermal Module Network Resistance Diagram
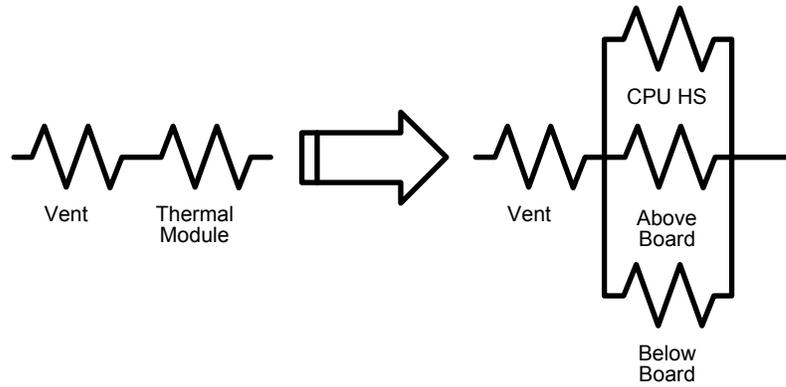
OM16725

The airflow through the CPU heatsink is determined by the interaction between the Thermal Module fan stand-alone fan curve and impedance characteristic of the selected CPU heatsink design (as illustrated in Figure 17). Improvements in CPU heatsink heat transfer performance often involve adding material to the heatsink. Improvements in conduction can be accomplished through increases in the conduction path cross-sectional area, and improvements in convection can be accomplished through increases in surface area and the addition of fins. Both of the engineering approaches typically increase the heatsink impedance because they decrease the channel area available for airflow. The reduction in airflow from an increase in heatsink impedance is illustrated in Figure 21 by the superposition of the heatsink impedance curves and the stand-alone fan curve.

Figure 21: Impact of Heatsink Impedance on Heatsink Airflow
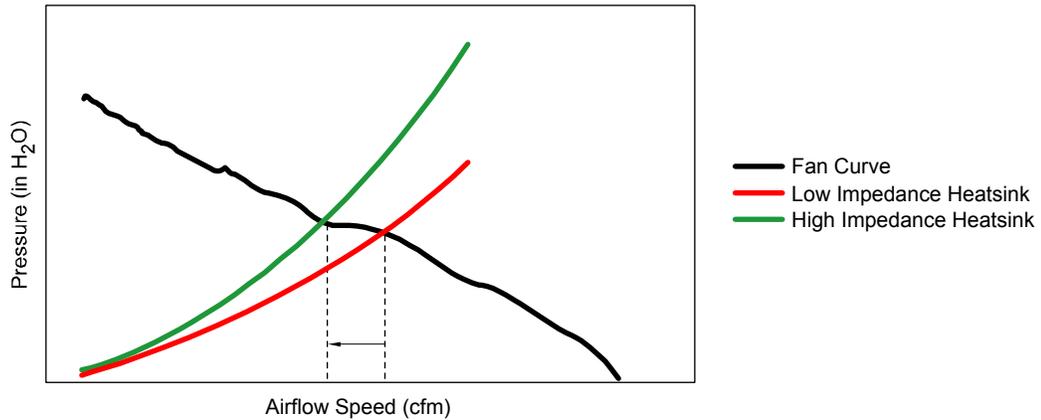
OM16726

The reduction in airflow not only impacts the CPU heatsink heat transfer performance by reducing its convection coefficient (see Equation 4), but it also impacts the Thermal Module effective fan curve (see Figure 17). Recall that the Thermal Module Effective fan curve is an indication of the airflow and pressure capability available for the remainder of the system. It is also important to note that the CPU heatsink impedance also influences the airflow that will go through the remaining paths; a higher impedance heatsink may force more airflow into the above and below motherboard airflow paths.

A Thermal Module supplier will need to carefully select or engineer the fan and heatsink characteristics to achieve the required CPU heatsink heat transfer performance (ψ_{ca}) and an appropriate effective fan curve. Returning to the network resistance diagram in Figure 20 and using the governing equations for airflow as a function of impedance (Equation 8), it is obvious that an increase in the impedance of one of the paths in parallel may increase the airflow to the remaining paths.

Equation 8: Airflow as a Function of Impedance

$$W = [2 \cdot dP/ ((L_1 + L_2) \cdot \rho)]^{1/2} \qquad \text{Flow through channels in series}$$

$$W = [2 \cdot dP/ (\rho \cdot (1/\sqrt{L_1} + 1/\sqrt{L_2})^2)]^{1/2} \qquad \text{Flow through channels in parallel}$$

Where:
W is the volumetric flow rate
dP is the pressure loss through two channels in series
L_1 and L_2 are the length of the two channels
ρ is the density of the fluid (air)

Additional information on selecting or engineering above and below motherboard airflow impedance is discussed in Sections 2.4.5.2 and 2.4.5.3.

A system integrator will need to select the Thermal Module based on the required CPU heatsink performance and the required effective fan curve.

2.4.5.2 Above Motherboard Airflow Path

The airflow for this path enters the channel created by the base of the CPU heatsink and the motherboard. As stated previously, the airflow speed will be a function of the Thermal Module stand-alone fan curve and the impedance of the three primary airflow paths: above and below the motherboard and through the ducted CPU heatsink.

The impedance of the above board airflow path is typically governed by the following physical characteristics:

a. The dimension of the channel created by the CPU heatsink base and motherboard.

   The impedance of the channel decreases as the channel height increases. Since the base of the CPU heatsink typically rests on the top of the CPU package, there are few design options to change this channel height.

   The shape of the bottom surface of the CPU heatsink base will affect the impedance of this channel. For instance, in Figure 22, the heatsink base shown on the right will have slightly lower channel impedance than the one on the left due to the relative reduction in turning loss.

Figure 22: Heatsink Base Impact on Above Motherboard Channel Impedance

OM16727
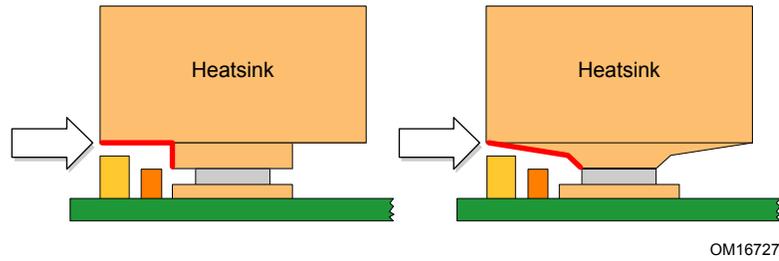
b. The size, quantity, and position of motherboard components near the CPU socket.

The impedance of the channel increases as the total frontal area of these components increases. The frontal area will increase as components are added, as those components become wider and taller, and when the components are closely spaced. For instance, in Figure 23 the Voltage Regulation configuration on the right will have slightly lower channel impedance than the one on the left (in this illustration, the airflow enters the page) because it has more open area available for the airflow that moves through it.

Figure 23: Voltage Regulation Component Impact on Motherboard Channel Impedance

OM16728
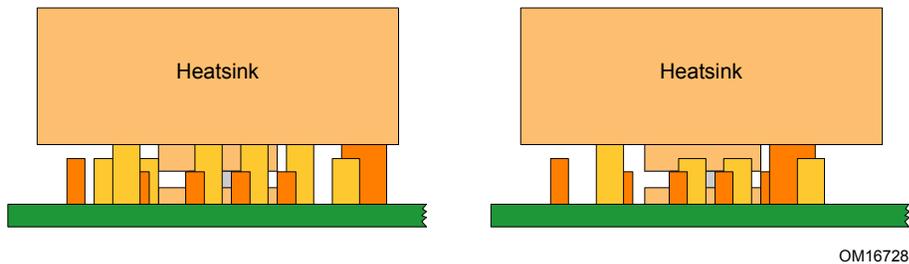
Since low profile Voltage Regulation components are typically more expensive, the system integrator will need balance the cost of the selection or specification of Voltage Regulation components against the airflow impedance that they will introduce.

c. The size and orientation of the CPU socket.

A larger CPU socket will increase the channel impedance; however, the orientation of the socket has a more significant influence. If the socket edge is perpendicular to the airflow direction, its impedance will be at its highest possible value. As the socket is rotated such that its corner is presented to the airflow, its impedance reduces. For instance, in Figure 24 the impedance of the socket on the right will be lower than the one on the left due to the relative reduction in turning loss.

Figure 24: Socket Orientation Impact on Impedance
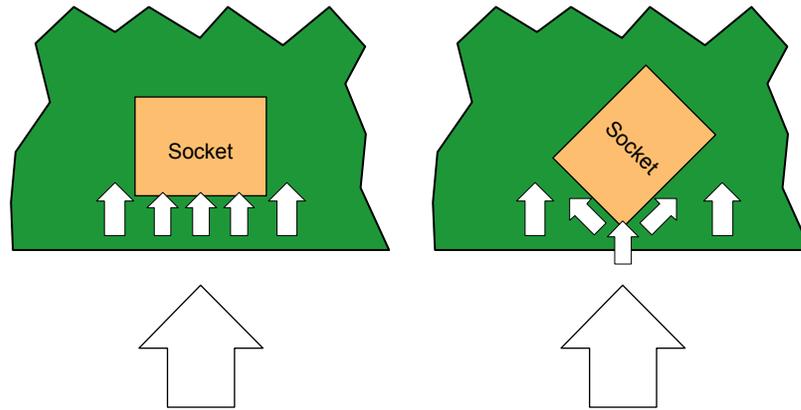
OM16729

Also, if the socket body is rotated, it acts as an effective airflow diverter. This allows the airflow to slipstream around the socket body and be efficiently directed to the remaining portions of the system. Figure 25 is an image extracted from a fluid dynamic numerical model of the airflow near a rotated socket body. In this illustration, the airflow comes from the Thermal Module fan that is mounted to the right. Note how the airflow efficiently begins to move around the socket body early in the airflow stream.

Figure 25: Rotated Socket Acting as a Flow Diverter
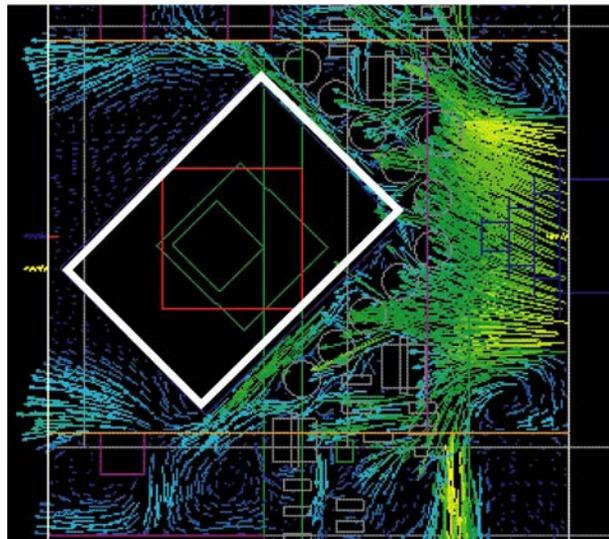
OM17114

The above motherboard airflow that goes around the socket and through the CPU Voltage Regulation components eventually provides inlet airflow to other portions of the system. Memory, located to the left of the CPU socket (in Figure 25, memory is located toward the bottom left of the illustration), receives a portion of the above motherboard airflow from the Voltage Regulation and socket area. This memory airflow path complements the Thermal Module exhaust airflow that exits the CPU heatsink.

2.4.5.3 Below Motherboard Airflow Path

The Thermal Module provides airflow below a BTX motherboard because the Thermal Module fan is located in front of, and its exit face extends below, the motherboard (see Figure 19). This airflow allows the bottom side of the motherboard to become a convective heat transfer surface. The motherboard is, of course, a complex electrical circuit composed of many conduits for managing electrical current (e.g., traces, planes, vias, and pads). As these motherboard features manage electrical current, that current generates heat power and, unless that heat is completely dissipated into the surrounding environment, the temperature of the motherboard will rise. When the motherboard temperature at any given location exceeds the operating temperature specification, the motherboard is subject to performance degradation. Since the above and below motherboard airflow in BTX allows heat to be effectively dissipated through its top and bottom surfaces, a BTX motherboard can manage higher current in critical locations. Especially in the region of the CPU socket and Voltage Regulation components, the bottom side heat transfer characteristic significantly increases the amount of heat, and thereby the amount of electrical current, that can be managed.

Surface mounted motherboard components can conduct a portion of the heat they generate into the motherboard. When the convective heat transfer capability of the motherboard is increased by above and below motherboard airflow, more heat from these surface mount components will conduct into the motherboard, effectively lowering the temperature of those components. This is especially important for CPU Voltage Regulation components. Components mounted to the bottom side of the motherboard will see the direct benefit of below motherboard airflow, since that airflow creates convective heat transfer directly from the components, in addition to removing heat from those components that is conducted into the motherboard. Even for components mounted on the topside of the motherboard, the below motherboard airflow is of significant benefit.

The below motherboard airflow is also a significant benefit for the Memory Controller Hub (MCH) and Input-Output Controller Hub (ICH) chipsets. Again, these components will conduct a portion of their heat into the motherboard and the convective heat transfer on the motherboard bottom side created by the below motherboard airflow effectively reduces the temperature of these components. More detail on MCH heat transfer is provided in Section 2.4.6 and on ICH heat transfer in Section 2.4.7.

The impedance of the below motherboard airflow channel will increase when the channel area is reduced. As described in Chapter 4, a Structural Retention Module is required in BTX system designs. Since this component occupies a portion of the space between the chassis pan and motherboard, it increases the below motherboard impedance and reduces the airflow. Components mounted on the bottom side of the motherboard will also increase the impedance.

Additional guidance on the thermal performance improvement from below motherboard airflow may be provided in subsequent revisions of this guide.

A system designer will need to pay particular attention to the way in which the Thermal Module airflow is partitioned. While the above and below motherboard airflow paths are important for managing CPU Voltage Regulation, CPU socket, and, to some extent, chipset temperatures, the CPU heatsink and the Thermal Module effective fan curve are typically more important. If the impedance of the above and below motherboard channels is particularly low, then more of the Thermal Module fan's airflow will go through these channels – meaning that less of that airflow will go through the CPU heatsink. Not only does this reduce the CPU heatsink heat transfer performance but it can also impact the heat transfer characteristics for the remaining portions of the system, since the Thermal Module effective fan curve essentially determines these important total system flow and flow distribution characteristics.

It may be necessary for the Thermal Module supplier to introduce impedance to the above and below motherboard channels to ensure that the appropriate airflow is directed to the CPU heatsink. System

integrators should select a Thermal Module based on its CPU heatsink ψ_{ca} and effective fan curve, but they should also understand the impact of Thermal Module design options on the above and below motherboard airflow.

2.4.5.4 Airflow into the Thermal Module

The *BTX Interface Specification* defines a requirement for a Thermal Module Interface. Chassis suppliers must always provide a Thermal Module Interface that is sealed and meets the open area, positional, and size requirements defined in the *BTX Interface Specification*. A BTX Form Factor compliant Thermal Module Interface may be provided through the use of a vented or open area in the chassis front panel (Figure 26) or by an airflow duct (Figure 27).

Figure 26: Thermal Module Interface Provided by Vented Front Chassis Panel
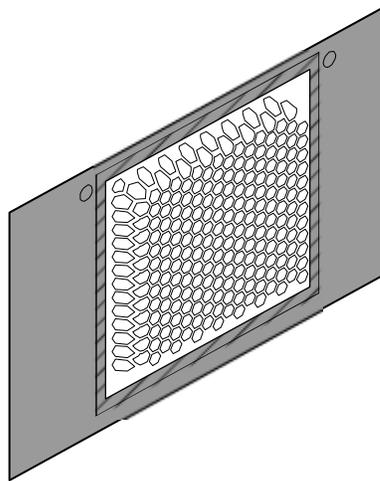
OM16730

Figure 27: Thermal Module Interface Provided by a Front Panel or Side Panel Airflow Duct
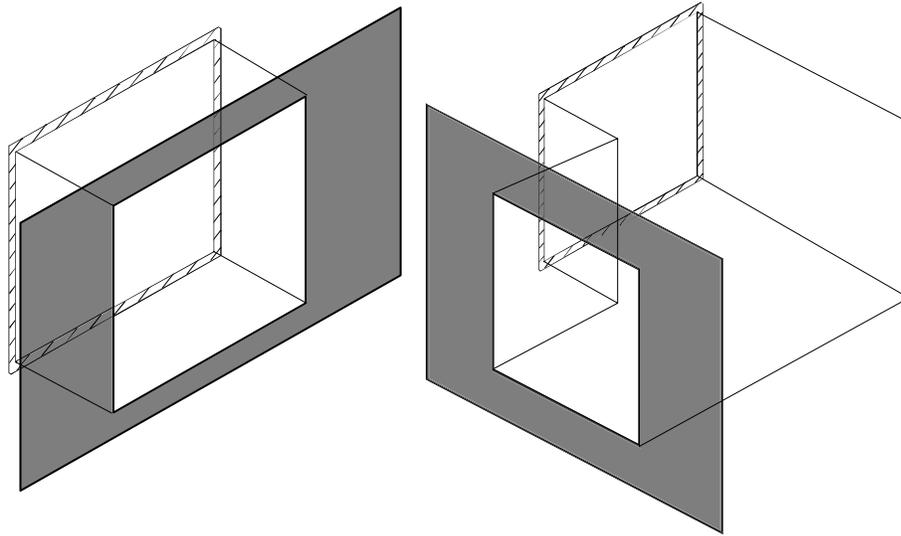
OM16731

The *BTX Interface Specification* also requires that the Thermal Module Interface be sealed in a way that prevents recirculated air from entering it. This assures that the Thermal Module fan will only pull air from outside the chassis through the Thermal Module Interface. A Thermal Module Interface inside the chassis walls (that is, if it is not created by the chassis sheet metal front panel) must be provided by the end of an airflow duct. This airflow duct should start at a chassis panel (Figure 27 illustrates start positions at both the front panel and a side panel) and end at the Thermal Module Interface. It would, of course, be beneficial to design gradual turns into any duct design, as opposed to the sharp turn shown in the rightmost illustration in Figure 27.

Should the chassis sheet metal front panel provide the Thermal Module Interface, the chassis supplier should appropriately engineer the vent open area. The open area is typically constrained by electromagnetic compatibility (EMC) that the chassis must also provide. Recall that the CPU is located near the front of the motherboard and, therefore, near the front panel and that the proximity of this electromagnetic interference (EMI) source to an open area in the front panel may be a concern. Nonetheless, the chassis supplier should optimize the hole size, shape, and pattern to create effective EMC without introducing airflow impedance that will substantially degrade the Thermal Module ψ_{ca} or operating point airflow. In this design situation, it is important for the system integrator to understand the impact that the near field impedance of the front panel vent will have on these Thermal Module characteristics. In particular, it may be beneficial to offset the vent area from the front panel to increase the distance between the vent and the Thermal Module fan.

A Thermal Module Interface is required by the *BTX Interface Specification*. The opening and physical interface requirements described must be provided by the BTX chassis supplier. It is against the Thermal Module Interface physical features that the inlet of the Thermal Module will fit. The system integrator will determine the quality of the seal between the Thermal Module and Thermal Module Interface, although the Thermal Module supplier will typically provide a compliant physical interface (for instance, a compression gasket or seal) to accommodate the expected dimensional tolerance variation.

2.4.6 MCH Airflow and Heat Transfer

Figure 28 illustrates the heat transfer path and associated network thermal resistance diagram for an MCH. Intel typically offers the MCH in a Flip Chip Ball Grid Array (FCBGA) package that does not have a metal lid, molded plastic, or other substance over the bare silicon die. Silicon is not a particularly good heat conductor (from Figure 28, θ_{ca} would be very high in the absence of an MCH heatsink) but the die surface is available for heat transfer augmentation, so a heatsink is typically directly attached to its surface. The heatsink significantly reduces the resistance of this heat transfer path and heat flow from the top surface of the MCH package is the primary heat transfer path.

Figure 28: MCH Heat Transfer Path and Resistance Network Diagram
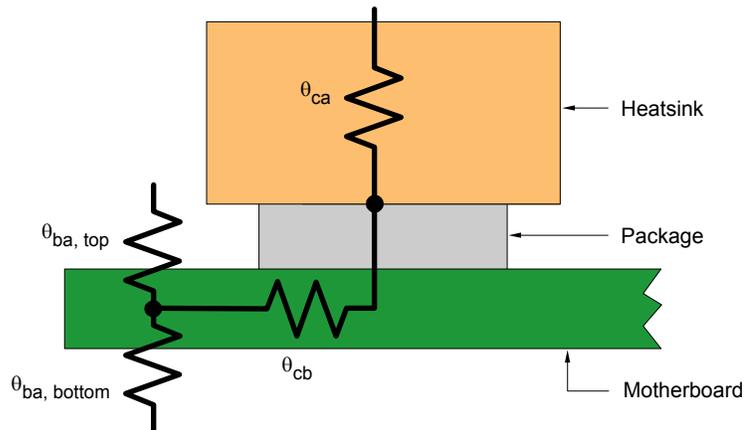
As discussed in Section 2.4.3, the below motherboard airflow augments the heat transfer for the MCH. It is also true that the above motherboard airflow flows over the motherboard area near the MCH and this improves the convective heat transfer from the motherboard topside, as well. Both of these motherboard convective heat transfer paths augment the heat transfer from the MCH and lower its temperature. However, because the resistance of the heat transfer path through the MCH heatsink is so low, the relative importance of heat transfer into and out of the motherboard diminishes.

Nonetheless, the heat transfer capability of the motherboard in the area of the MCH can and should be characterized. Knowledge of these heat transfer paths can be used to optimize the performance and cost of the MCH heatsink. The network resistance diagram in Figure 28 can be used to define the total heat transfer resistance equation (Equation 9) – in this equation, the "up" and "down" terms are in reference to the power source (MCH silicon). The required total resistance can be calculated from the component power and temperature requirements, and the MCH ambient temperature (Equation 10).

Note: MCH power (Q_{MCH}) and case temperature (T_{case}) requirement are specified in the component Datasheet. Ambient temperature ($T_{ambient}$) in the MCH region should be extracted from a system level thermal numerical model or from system temperature measurements. Ambient temperatures characterized in Intel reference design systems are available in Section 2.7.3.

Equation 9: Total Heat Transfer

$\theta_{total} = 1/(1/\theta_{up}+1/\theta_{down})$
$\qquad\qquad = 1/((1/\theta_{ca})+1/(\theta_{cb}+1/(1/\theta_{Ba, top}+1/\theta_{ba, bottom})))$

Equation 10: Required Total Resistance

$\theta_{total} = (T_{case} - T_{ambient}) / Q_{MCH}$

From this equation, the required heatsink resistance (θ_{ca}) can be calculated (Equation 11).

Equation 11: Total Heatsink Resistance

$\theta_{ca} = 1/((Q_{MCH}/(T_{case}-T_{ambient}) - 1/(\theta_{cb} -1/(1/\theta_{Ba, top}+1/\theta_{ba, bottom})))$

Equation 11 shows how the required MCH heatsink performance must improve if the motherboard heat transfer characteristics are unknown. Table 2 offers an example of the relative contribution of the θ_{ba} resistance terms – when the motherboard resistance terms are unknown, they must be assumed to be infinite resistance.

Table 2: Motherboard Heat Transfer Impact to Required MCH ψja

		Case Temperature (°C)	112
		Ambient Temperature (°C)	45
		MCH Power (W)	18
θ_{cb} (°C/W)	θ_{ba}, top (°C/W)	θ_{ba}, bottom (°C/W)	θ_{ca} (°C/W)
15	15	15	3.59
30	30	40	3.66
∞	∞	∞	3.72

The performance and cost for the MCH heatsink is also a function of the airflow through the MCH heatsink. The Thermal Module exhaust airflow is the primary inlet airflow for the MCH heatsink. As discussed in Section 2.4.1, the heat transfer performance of a heatsink is a strong function of the amount of airflow that passes through it. However, unless it is directed through a heatsink, airflow will often choose to go around it (see Figure 29).

Figure 29: Illustration of Airflow through a Heatsink as a Function of Bypass
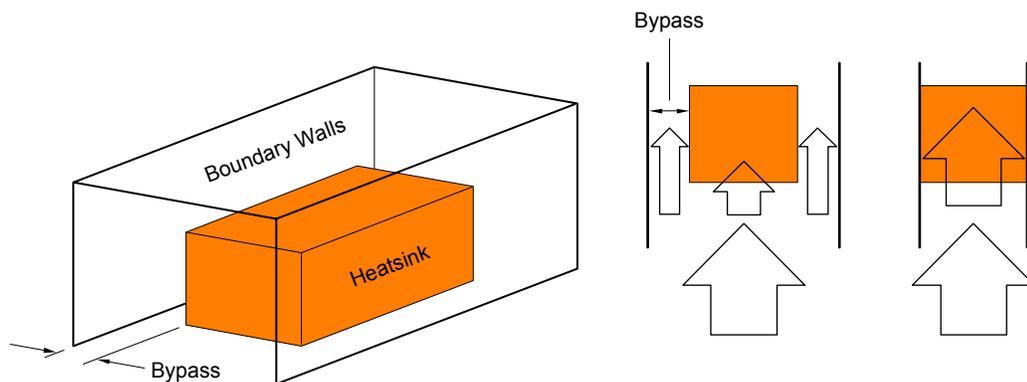
OM16733

The amount of open channel area around a heatsink (termed bypass) determines the amount of airflow that will go through a heatsink. As illustrated in Figure 30, reducing the bypass around a heatsink increases the amount of airflow that will go through it. Although the airflow through a heatsink improves with a reduction in bypass, the impedance of that restricted airflow path also increases (Figure 30).

Figure 30: Impact of Bypass on Heatsink Airflow and Impedance
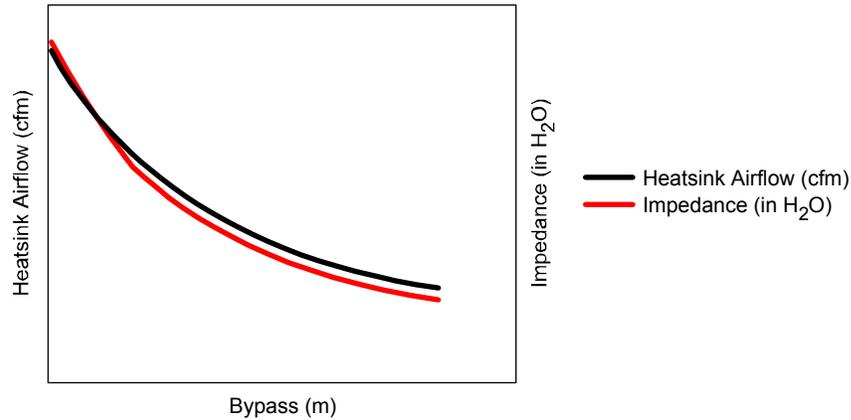
OM16734

A system designer will need to be cognizant of the bypass around the MCH in the design or selection of the MCH heatsink, since there are several system configuration choices that will impact the airflow bypass around the MCH. For instance, an add-in card on a riser may occupy the volume immediately above the MCH (above the Zone B Keepout height required in the *BTX Interface Specification*), as shown in Figure 31 or the area above the MCH may be free all the way to the chassis boundary wall, as shown in Figure 32.

Figure 31: MCH Bypass Illustration – Graphics Add-in Card on a Riser
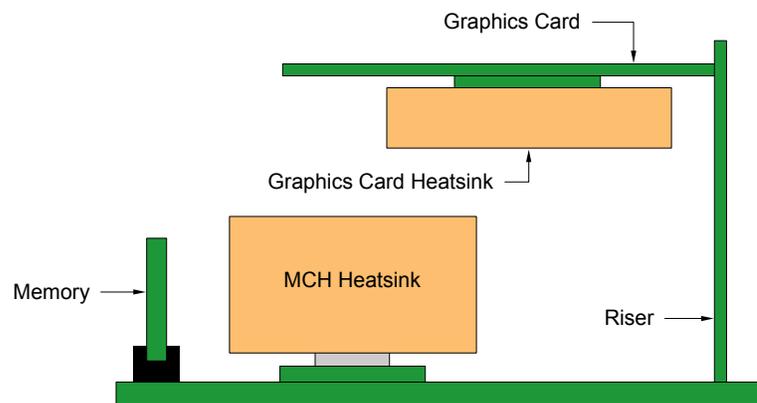
OM16735

Figure 32: MCH Bypass Illustration – Chassis Boundary Wall
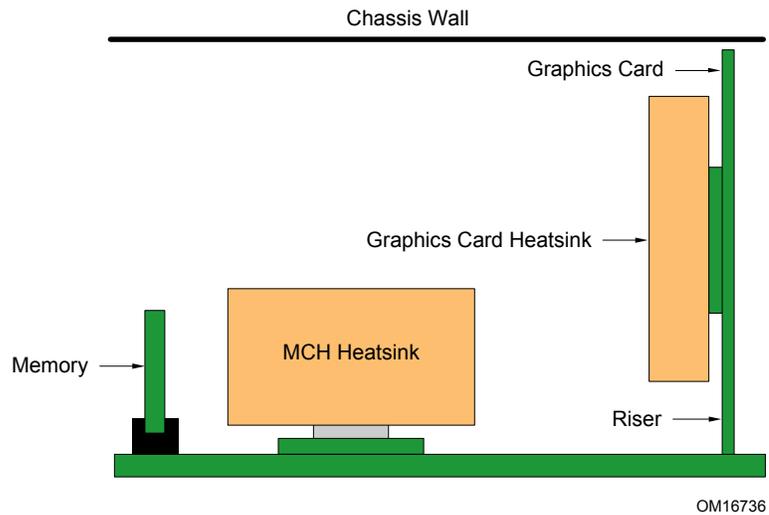
It is not unusual for the motherboard supplier to provide the MCH heatsink as part of the motherboard assembly. Knowledge and characterization of the motherboard heat transfer characteristics in a typical BTX environment will reduce the performance requirement for the MCH heatsink; however, the motherboard supplier is unlikely to have knowledge of the bypass around the MCH heatsink since the system configuration into which the motherboard will be installed may not be known. In this case, the motherboard supplier's thermal engineer should design the MCH heatsink to provide the required ψ_{ja} assuming infinite bypass.

MCH heatsink boundary condition airflow characterized in Intel reference design systems is available in Section 2.7.3.

2.4.7 ICH Airflow and Heat Transfer

Figure 33 illustrates the heat transfer path and associated network thermal resistance diagram for an ICH. Intel typically offers the ICH in a package with a molded plastic body. Neither silicon nor molded plastic are particularly good heat conductors so the relative importance of the motherboard heat transfer characteristics increases.

Figure 33: ICH Heat Transfer Path and Resistance Network Diagram
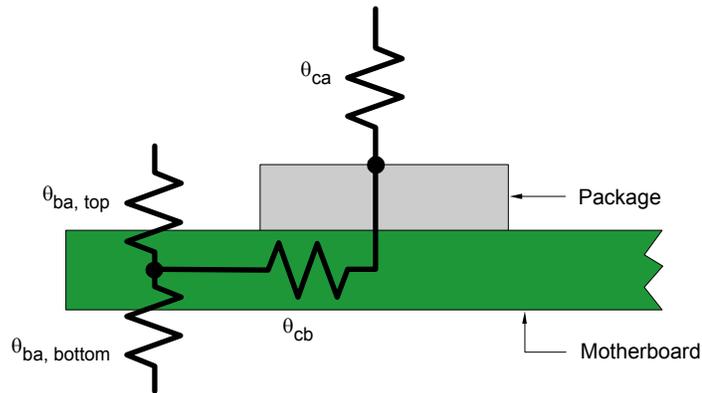
OM16737

The heat transfer capability of the motherboard in the area of the ICH should be characterized. Knowledge of these heat transfer paths can be used to determine if an ICH heatsink is required. The network resistance diagram and the heat flow illustration and the network resistance diagram in Figure 33 can be used to generate Equation 12. The required total resistance can be calculated from the component power and temperature requirements, and the ICH ambient temperature (Equation 13).

Equation 12: Total Resistance

$$\theta_{total} = 1/(1/\theta_{up}+1/\theta_{down})$$
$$= 1/((1/\theta_{ca})+1/(\theta_{cb}+1/(1/\theta_{Ba, top}+1/\theta_{ba, bottom})))$$

Equation 13: Total Required Resistance

$$\theta_{total} = (T_{case} - T_{ambient}) / Q_{ICH}$$

From this equation, the required heatsink resistance (θ_{ca}) can be calculated (Equation 13).

Equation 14: Required Heatsink Resistance

$$\theta_{ca} = 1/((Q_{ICH}/(T_{case}-T_{ambient}) - 1/(\theta_{cb} - 1/(1/\theta_{ba, top}+1/\theta_{ba, bottom})))$$

Equation 13 shows that the ICH heatsink may be eliminated if the motherboard heat transfer characteristics are known. Equation 14 offers an example of the relative contribution of the θ_{ba} resistance terms.

Table 3: Motherboard Heat Transfer Impact to Required ICH $\psi\chi a$

		Case Temperature (°C)	110
		Ambient Temperature (°C)	50
		MCH Power (W)	4
θ_{cb} (°C/W)	θ_{ba}, top (°C/W)	θ_{ba}, bottom (°C/W)	θ_{ca} (°C/W)
15	15	15	14.7
30	30	40	14.9
∞	∞	∞	15.2

Ambient temperature boundary conditions for the ICH characterized in Intel reference design systems are available in Section 0.

2.4.8 Graphics Add-in Card Airflow and Heat Transfer

The PCI Express* specifications approved in 2003 allow high performance x16 PCI Express graphics add-in cards up to 75 watts. This power is distributed amongst the graphics card Graphics Processing Unit (GPU), Voltage Regulation, and memory.

It is not unusual for the graphics card supplier to provide the GPU heatsink as part of the graphics card motherboard assembly. Knowledge of the airflow characteristics around a graphics card in a typical BTX environment will allow the heatsink designer to determine the performance requirement. Ambient temperature and airflow graphics add-in card boundary conditions characterized in Intel reference design systems are available in Section 2.7.3.

However, as is the case with the MCH heatsink, the graphics supplier is unlikely to have knowledge of the graphics card bypass since the system configuration into which the card will be installed may not be known (see Figure 31 and Figure 32 for examples). In this case, the graphics card supplier's thermal engineer should design the GPU heatsink to provide the required ψ assuming infinite bypass. A graphics card supplier should provide a description of the airflow boundary conditions (inlet airflow velocity and temperature) for which their heatsink was designed.

Since a graphics card is typically installed in the first expansion slot (whether directly into that slot or on a riser), it will receive the Thermal Module exhaust airflow. Therefore, the performance of the graphics card heatsink will be determined by the Thermal Module effective fan curve that the system integrator selects. The system integrator should ensure that the selected Thermal Module provides the required graphics card airflow boundary conditions.

The airflow pattern in the graphics add-in card heatsink will be strongly influenced by the rear panel ventilation. More rear panel open area near the graphics add-in card heatsink exhaust will allow the airflow to travel completely through the heatsink. Less open area will force the airflow to leave the heatsink at the top of the fins, reducing heatsink performance.

2.4.9 Memory Airflow and Heat Transfer

Memory is located to the left of the Thermal Module and MCH. Memory will typically receive its inlet airflow from two sources: Thermal Module and CPU Voltage Regulation exhaust.

The Voltage Regulation exhaust airflow will enter memory along the motherboard top surface. The Thermal Module exhaust will cross over memory upon being diverted by the MCH heatsink and, potentially, the graphics add-in card if it is on a riser behind the Thermal Module. The exhaust airflow

from the Thermal Module will have expansion and turning when it exits the Thermal Module, and some of this flow will make its way across the memory channels. However, some of the Thermal Module exhaust airflow will go all the way to the rear panel before entering the memory channel from the rear.

The location of the Power Supply Unit (PSU) and the PSU effective fan curve will have a significant influence over the memory airflow pattern. If the PSU is located next to the left (fixed) motherboard edge, the location on the PSU of its inlet airflow will determine the pressure characteristics near memory, and this will determine how the Thermal Module and Voltage Regulation exhaust airflow streams will travel over and through the memory channels.

Generally, airflow will enter the memory channels from the front (CPU Voltage Regulation exhaust) and the rear (Thermal Module exhaust turned at the rear panel) and will leave the memory channel near the PSU fan inlet. Physical features over the top of the memory channel – drives or drive brackets, cables, etc. – will also influence how the airflow travels through the memory channels.

If a fluid dynamic numerical model of the system is constructed to predict the airflow pattern, the system thermal engineer should consider significantly increasing the model's element mesh density in the Thermal Module exhaust region. This will allow a more accurate prediction of the exhaust airflow expansion and turning characteristics and, therefore, a more accurate prediction of the airflow behavior near memory.

It may also be beneficial for the system thermal engineer to use a detailed fluid dynamic and temperature numerical model to predict memory component temperature. Memory suppliers will often provide detailed numerical representations of their memory card products that include detailed representations of each package's heat power and heat transfer characteristics. The system engineer may extract the memory airflow boundary conditions from their system model and apply them to the detailed memory model to determine if the predicted memory component temperatures meet the memory supplier's specifications. The boundary condition extraction from the system model should include airflow velocity (speed and direction) and temperature on the inlet planes and pressure on the exit planes around the memory, as illustrated in Figure 34.

Figure 34: Memory Airflow Boundary Conditions from the System Model Applied to a Detailed Memory Model (MCH Position Shown to Illustrate Memory Orientation)

Pressure

Pressure

Velocity
Temperature

MCH

Velocity
Temperature

Velocity
Temperature
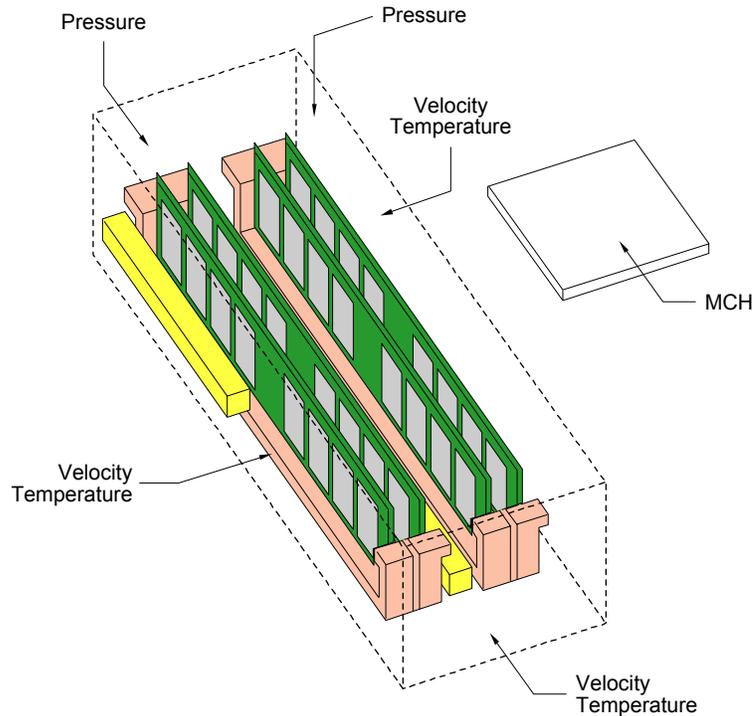
OM16738

2.5 Chassis Vent Design

The selection of chassis venting significantly influences the total system impedance and the system airflow pattern. The in-line airflow pattern selected by BTX form factor engineers is one whose benefits will be significantly reduced by poor vent engineering.

The location and design of vents should be determined based on the desired airflow pattern in the specific system design. Fluid dynamic numerical models are the best tool to use to predict the system airflow pattern. The use of a numerical tool also allows the evaluation of different vent designs and locations without the expense of tooling and evaluating prototype designs.

2.5.1 Rear Panel Venting

Generally, the rear panel should be generously ventilated in the area not consumed by the Rear Panel I/O, Expansion Slots, or PSU interface locations.

Because of the below motherboard airflow generated by the Thermal Module, a BTX rear panel should include ventilation along the edge nearest the chassis pan (as illustrated in Figure 35). Failure to provide that ventilation will increase the below motherboard impedance and decrease the below motherboard airflow.

Figure 35: Below Motherboard Rear Panel Vent Illustration
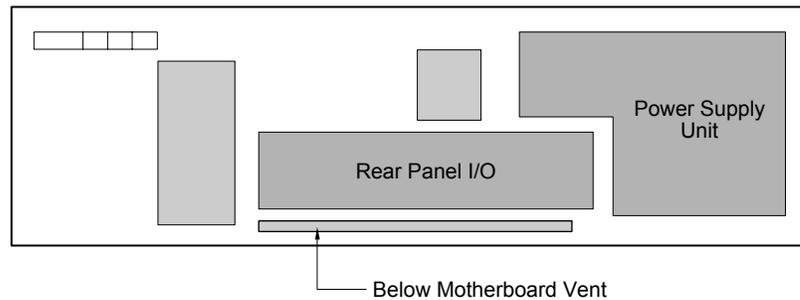
OM16739

2.5.2 Front Panel Venting

Obviously, the chassis inlet to the Thermal Module should be ventilated. Again, a vented portion of the front panel may be used to provide the Thermal Module Interface. If a duct is used to create the Thermal Module Interface, then the inlet that the chassis provides to the duct must be open or ventilated. If ventilated, there may be acoustic benefit from the vented surface away from the Thermal Module fan inlet. The *BTX Form Factor Specification* does not require that the vented area be on a concurrent plane with the Thermal Module Interface, so an offset vent design is completely compliant with the specification. Figure 26 in Section 2.4.5.4 shows an offset vent pattern integrated into the required Thermal Module Interface.

A chassis designer or system integrator should include ventilation around the drive bays. As in most system designs, creating airflow around the drive bays is difficult; nonetheless, airflow can be established in the drive bay area through appropriate inclusion of front panel venting. The system integrator should be aware that these front panel vents might exhaust airflow for some combinations of Thermal Module and PSU fan speed settings. In this case, the exhaust airflow from the vents might be added to the inlet airflow (i.e. re-circulation) for the Thermal Module. It will be beneficial for the front bezel to include feature that direct the front panel vent exhaust airflow away from the Thermal Module inlet. At the very least, the bezel design should now allow immediate recirculation of the drive bay vent exhaust into the inlet for the Thermal Module Interface.

2.5.3 Vent Design

When a vent is used instead of a fully open inlet, the vent will introduce impedance to the airflow that passes through it. This impedance, like all other impedance, will reduce the airflow or force the fan to work harder to maintain the target airflow. Fans that work harder generate more noise.

The impedance introduced by vents can be minimized through the appropriate selection of vent hole size, pitch, and the shape of its perimeter. Large, closely spaced holes will have lower impedance than small, widely spaced holes. However, the EMC offered by larger holes is typically not as good as that offered by small holes. Finally, the tooling cost for closely spaced holes and /or non-circular hole geometries is typically higher.

A chassis designer or system integrator will need to balance the tooling cost, EMC, and impedance of each vent location. The impedance for each vent can be calculated using the method outlined below.

a. Calculate the vent hole hydraulic diameter:

Equation 15: Vent Hole Hydraulic Diameter

$dH = 4 \cdot A_{hole} / P_{hole}$

> where:
> A_{hole} is the area of the vent hole
> P_{hole} is the perimeter of the vent hole

b. Calculate the Free Area Ratio (see Figure 36).

Equation 16: Free Area Ratio; Vent

$FAR = \Sigma\, A_{hole,\, i} / A_{grid}$

> where A_{grid} is the area of the vent hole pattern outline
> (see Figure 36)

Figure 36: Open Area Ratio illustration ($A_{grid} = X \cdot Y$)
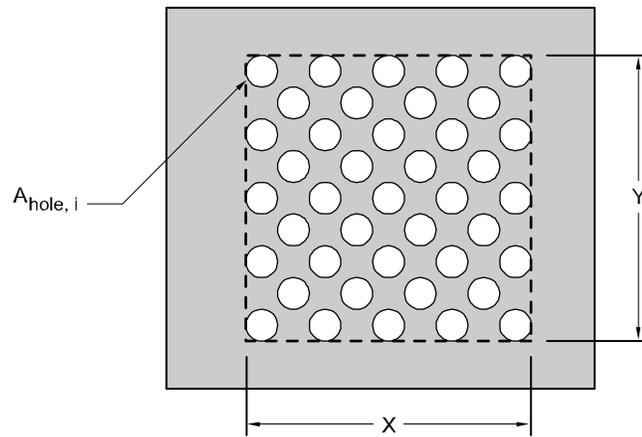
c. Calculate the Reynolds Number:

Equation 17: Reynolds Number

$Re = (W \cdot d_H) / (FAR \cdot A_{grid} \cdot \nu)$

> where W is the volume flow rate approaching the vent
> ν is the kinematic viscosity of air (0.000016 m^2/s)

d. Calculate the ratio of the channel length (e.g., the sheet metal thickness of the chassis panel) to the hydraulic diameter:

Equation 18: Channel Length to Hydraulic Diameter Ratio

$f = \text{thickness} / d_H$

e. For $f \geq 0.015$, use Equation 19 to calculate the vent loss coefficient (ξc); otherwise, use Equation 20:

Equation 19: Vent Loss Coefficient (f ≥ 0.015)

$$\zeta c := \frac{\left[\dfrac{(1 - FAR)^{0.375}}{\sqrt{2}} + 1 - FAR\right]^2}{FAR^2}$$

Equation 20: Vent Loss Coefficient

$$\zeta c := \frac{\left[\left(0.5 + \tau \cdot \sqrt{1 - FAR}\right) \cdot (1 - FAR) + (1 - FAR)^2 + 64 \cdot \dfrac{f}{Re}\right]}{FAR^2}$$

Note: The loss coefficient is inversely proportional to the square of the FAR.

f. Calculate the vent impedance:

Equation 21: Vent Impedance

$dP = 0.5 \cdot \rho \cdot \xi_c \cdot (W / A_{grid})^2$

where ρ is the density of air (1.177 kg/m^3)

Notice that the vent impedance is proportional to the loss coefficient, which is proportional to the vent FAR. Increasing the open area at any vent location will reduce its impedance. Figure 37 illustrates two vent pattern options with equivalent A_{grid}. Since the packing density (number of holes per unit area) is higher for hexagonal holes than circular holes, the FAR is higher and the impedance is lower for the hexagonal pattern shown on the left.

Figure 37: Vent Patterns Illustrations
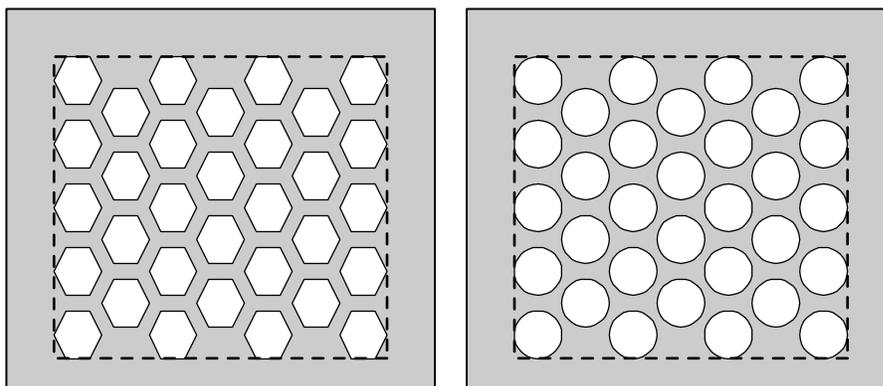
OM16741

Notice, also, that the impedance is proportional to the square of the volumetric airflow and inversely proportional to the square of the vent area. Increasing the allocation of a chassis panel to vent area (increasing A_{grid}) reduces the impedance contribution of the vent. A reduction in system impedance through a reduction in vent impedance typically increases the volumetric airflow that a Thermal Module fan can force through the system. This increase in system airflow will result in an increase in the volumetric flow approaching the vent (increasing W), which increases the vent impedance. This complex relationship is simplified, however, when the chassis designer or system integrator understands that the increase in volumetric airflow approaching the vent is typically a small portion of the total system airflow; therefore, vent impedance will decrease as the vent area is increased.

Finally, the information and equations used in the impedance calculations above were based on a review of a reference fluid dynamic engineering text (Chapter 8 of *Handbook of Hydraulic Resistances*, 3rd edition, I.E. Edelchik). A chassis design or system integrator thermal engineer will find this reference text quite useful, since the equations used above have a limited domain of sheet metal thickness in which they are valid.

2.6 Front Bezel Design

The open area and shape of the front bezel inlet airflow opening will have a significant influence on the system impedance. A low impedance bezel design assures that the Thermal Module can deliver the required system airflow at low fan operating speed. As the bezel's impedance increases, the system airflow will reduce or the fan's operating speed will need to increase to achieve the required airflow. Higher fan operating speeds increase system noise.

Of course, the front bezel is often the most important industrial design element of the chassis. Bezel industrial designers are often reluctant to concede that the front bezel requires airflow inlet openings because they perceive that the opening erodes the bezel's cosmetic value. Also, bezel access for drive bays and front panel I/O often reduces the remaining bezel area allocated for airflow openings.

The bezel industrial designer must work with the system thermal engineer to design a bezel that does not substantially increase the bezel impedance. Failure to do so will erode the system's acoustic performance and may impact its ability to maintain compliant operating temperatures within the system.

Numerical CFD analysis studies conducted by Intel have demonstrated that a reasonably low impedance bezel can be designed with side or bottom inlet ventilation patterns, leaving the cosmetic front surface of the bezel unencumbered with vent openings. This is discussed further in Section 2.7.3.1.2 and Table 10.

2.6.1 Bezel Impedance

Bezel impedance is a function of the total open area, the shape of the inlet airflow passages, and the number of turns from the bezel to the chassis inlet. As with any airflow channel, decreasing the channel cross-section area, increasing the channel length, and forcing the airflow to turn will increase the impedance.

Larger systems have larger front bezels that allow the bezel industrial designer and system thermal engineer more options. In larger systems, the airflow channel may be longer but the bezel design can compensate for this impedance increase by minimizing the number of turns or by increasing the channel cross-sectional area.

In small system designs, there are fewer design options for the location and density of inlet airflow openings. Small system designs are also more likely to use a vented portion of the front panel instead of a duct to provide the required Thermal Module Interface. With the smaller area offered by a small front bezel, the bezel inlet opening must often be located immediately in front of the Thermal Module Interface. The industrial designer will likely want design the bezel inlet opening in a way that eliminates the "line of sight" view of the sheet metal chassis front panel that the opening otherwise might provide. Figure 38 illustrates inlet vent channel design options that eliminate "line-of-sight."

Figure 38: Bezel Airflow Channel Design Options That Eliminate "Line of Sight"
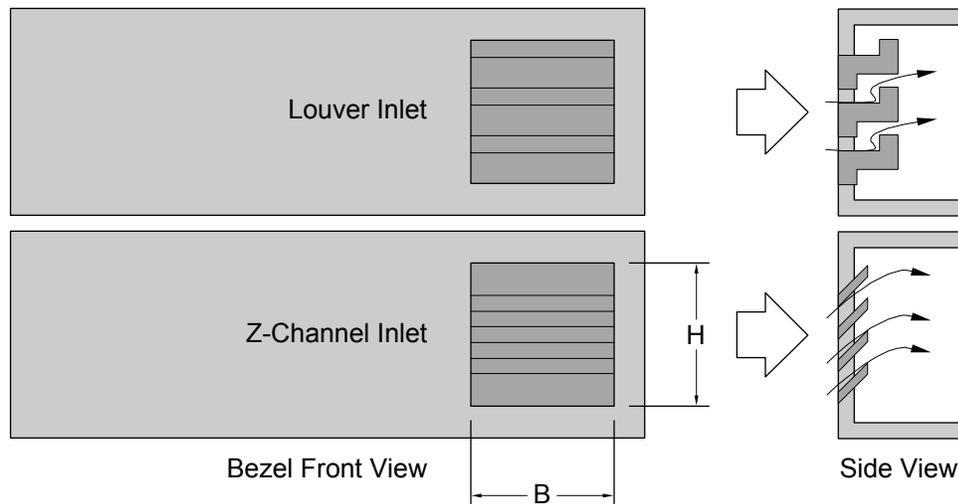
A system integrator will need to balance the tooling cost, industrial design, and impedance of the front bezel. The impedance of the Louver and Z-Channel inlet geometries can be calculated using the methods outlined below. The information and equations used in the bezel impedance calculations were based on a reference fluid dynamic engineering text (Chapters 3 and 6 of *Handbook of Hydraulic Resistances*, 3rd edition, I.E. Edelchik). A system integrator thermal engineer (and industrial designer) will find this reference text quite useful, since the equations and loss coefficient figures used below have a domain limited to these two inlet geometries.

The inlet geometry dimensions used in the impedance calculations are illustrated in Figure 39 (Z-Channel) and Figure 40 (Louver).

Figure 39: Z-Channel Front Bezel Inlet Geometry
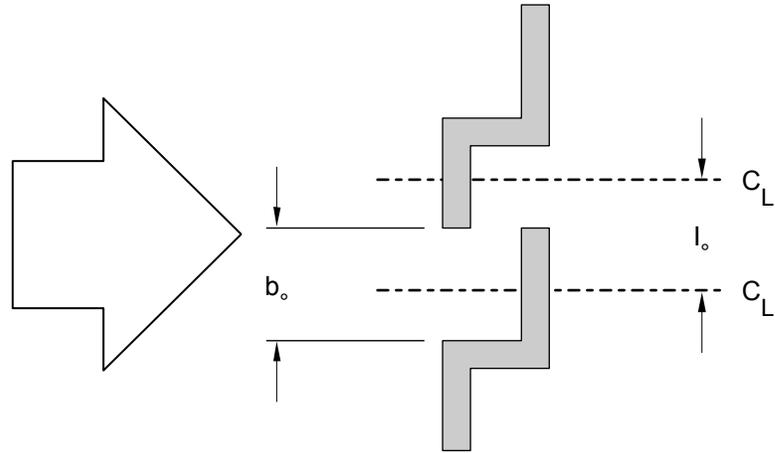
OM16743

Figure 40: Louver Front Bezel Inlet Geometry
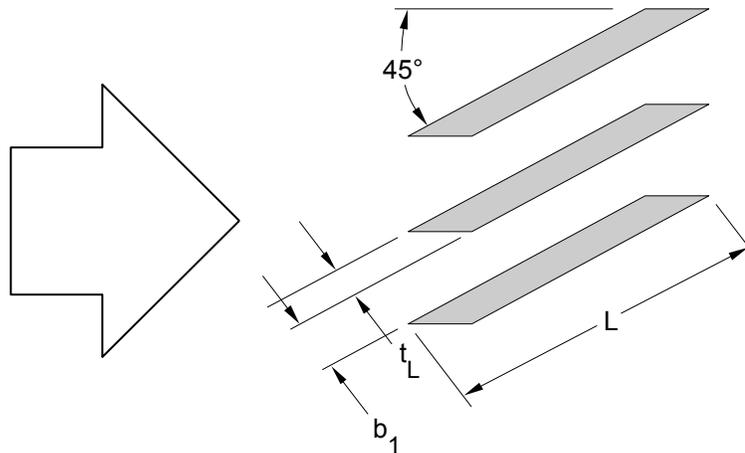
OM16744

Z-Channel vent channel impedance calculation method:

a. Calculate the open area of one slot:

Equation 22: Slot Open Area

$A_{slot} = b_o \cdot B$

where B is the width of the vent opening (see Figure 38)

b. Calculate the FAR:

Equation 23: Slot Free Aspect Ratio

FAR = $n \cdot A_{slot} / (B \cdot H)$

where B and H are the width and height of the bezel vent opening (see Figure 38)

c.   Calculate the hydraulic diameter:

Equation 24: Slot Hydraulic Diameter

$d_H = 2 \cdot A_{slot} /(B + b_o)$

d.   Calculate the Reynolds number:

Equation 25: Reynolds Number

Re = $(W \cdot d_H)/ (n \cdot A_{slot} \cdot v)$

where W is the volume flow rate approaching the bezel vent
n is the number of slot openings in the bezel vent
v is the Kinematic viscosity of air (0.000016 m^2/s)

e.   Calculate the friction loss coefficient:

Equation 26: Friction Loss Coefficient

$\xi_{friction}$ = $64 \cdot l_o /(Re \cdot b_o)$

f.   Select a turning loss coefficient, $\xi_{turning}$, from Figure 41 based on the ratio of l_o / b_o.

Figure 41: Turning Loss Coefficient as a Function of l_o / b_o
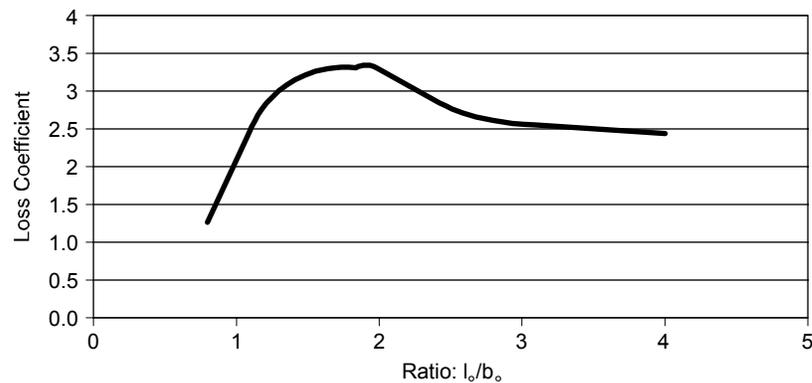
OM16745

g.   Select the contraction and expansion loss coefficient from Figure 42 based on the FAR.

Figure 42: Expansion and Contraction Loss Coefficients as a Function of FAR
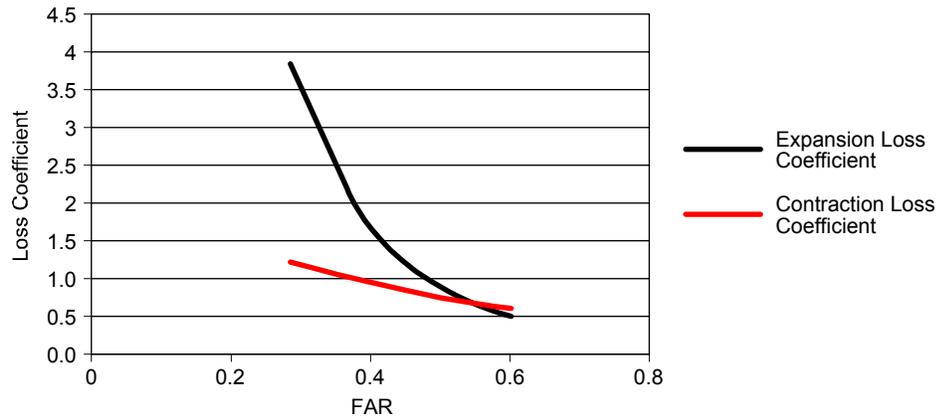
OM16746

h. Calculate the total loss coefficient:

Equation 27: Total Loss Coeffiecient

$$\xi_{total} = \xi_{contraction} + \xi_{expansion} + (\xi_{friction} + \xi_{turning}) / FAR^2$$

i. Calculate the front bezel vent impedance loss:

Equation 28: Front Bezel Vent Impedance Loss

$$dP = 0.5 \cdot \rho \cdot \xi_{total} \cdot FAR \cdot W / (n \cdot A_{slot})$$

Louver vent channel impedance calculation method:

a. Calculate the open area of one slot:

Equation 29: Slot Open Area

$$A_{slot} = b_1 \cdot B$$
where B is the width of the vent opening (see Figure 38)

b. Calculate the FAR

Equation 30: Slot Free Aspect Ratio

$$FAR = b_1 / (b_1 \cdot t_L)$$

c. Calculate the hydraulic diameter:

Equation 31: Slot Hydraulic Diameter

$$d_H = 2 \cdot A_{slot} / (B + b_1)$$

d. Calculate the Reynolds number:

Equation 32: Reynolds Number

$Re = (W \cdot d_H)/ (n \cdot A_{slot} \cdot v)$

where W is the volume flow rate approaching the bezel vent

n is the number of slot opening in the bezel vent

v is the kinematic viscosity of air (0.000016 m^2/s)

e. Calculate the friction loss coefficient:

Equation 33: Friction Loss Coefficient

$\xi_{friction} = 64 \cdot L /(Re \cdot b_1)$

f. Calculate the Louver vent loss coefficient:

If $L/ b_1 < 11$ (1 – FAR), use Equation 27, otherwise use Equation 28.

Equation 34: Louver Vent Loss Coefficient (L/ b1 < 11 (1 – FAR))

$\xi_{louver} = k (0.85+(1-FAR)^2+\xi_{friction})/FAR^2$

Equation 35: Louver Vent Loss Coefficient

$\xi_{louver} = k (0.85+(1-FAR)^2+\xi_{friction})/FAR^2 + 0.5(11 (1-FAR) – L/b_1)$

where k is determined from Figure 43

Figure 43: Louver Inlet Design Impact on Variable *k*

$k = 1.0$ $k = 0.6$

OM16747
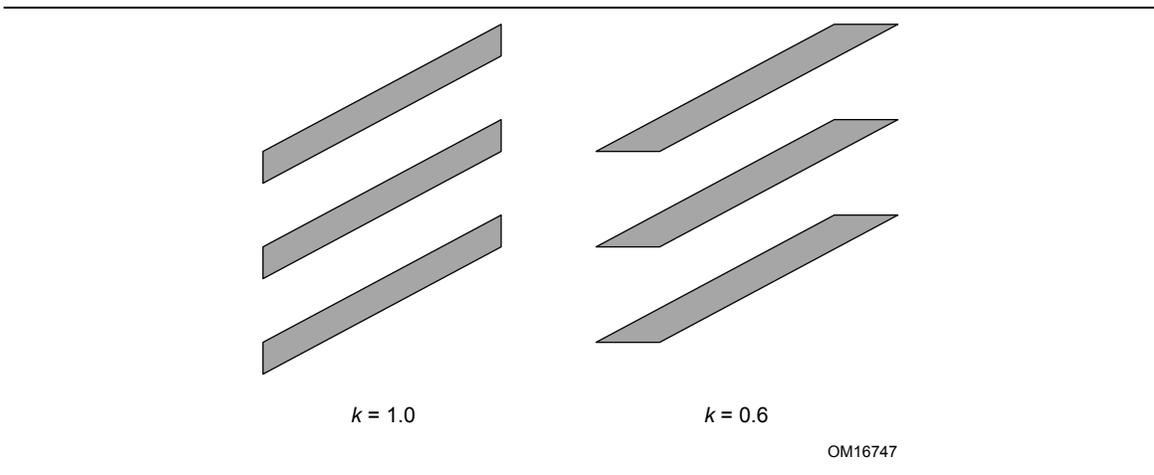
Notice that the Louver design on the right in Figure 43 will have a lower impedance than the one on the left because the value of *k* is smaller. The impedance and loss coefficients are directly proportional to *k*.

g. Calculate the total vent loss coefficient by adding the contraction and expansion loss coefficient estimates:

Assume contraction loss coefficient, $\xi_{contraction} = 0.5$.

Assume expansion loss coefficient, assume $\xi_{expansion} = 0.2$.

Equation 36: Total Loss Coefficient

$$\xi_{total} = \xi_{friction} + \xi_{contraction} + \xi_{louver} + \xi_{expansion}$$

h. Calculate the front bezel vent impedance loss:

Equation 37: Front Bezel Impedance Loss

$$dP = 0.5 \cdot \rho \cdot \xi_{total} \cdot FAR \cdot W / (n \cdot A_{slot})$$

2.6.2 The Bezel's Role in Managing Recirculation

As discussed briefly in Section 2.5.2, the bezel design may need to include features to manage the airflow exhausted by the front panel vents, since there may be fan speed settings that create exhaust of higher temperature air from some front panel vent locations (e.g., drive bay vents). If this exhaust airflow is allowed to immediately recirculate and join the Thermal Module inlet airflow, it will increase the system inlet ambient temperature. This will have an immediate and detrimental effect on all system temperatures. As described Chapter 1, the system may respond to an increase in the temperature at its monitor points by increasing fan speeds, which will increase system noise.

The bezel design can include simple features that segregate the airflow within the bezel, such that the potential exhaust airflow from other front panel vent locations does not immediately join the Thermal Module inlet airflow (see Figure 44). Forcing the vent airflow to completely exhaust from the bezel before it rejoins the Thermal Module inlet airflow will reduce the impact of front panel vent recirculation.

Figure 44: Illustration of Bezel Features That Prevent Within-Bezel Recirculation of Front Panel Vent Exhaust Airflow
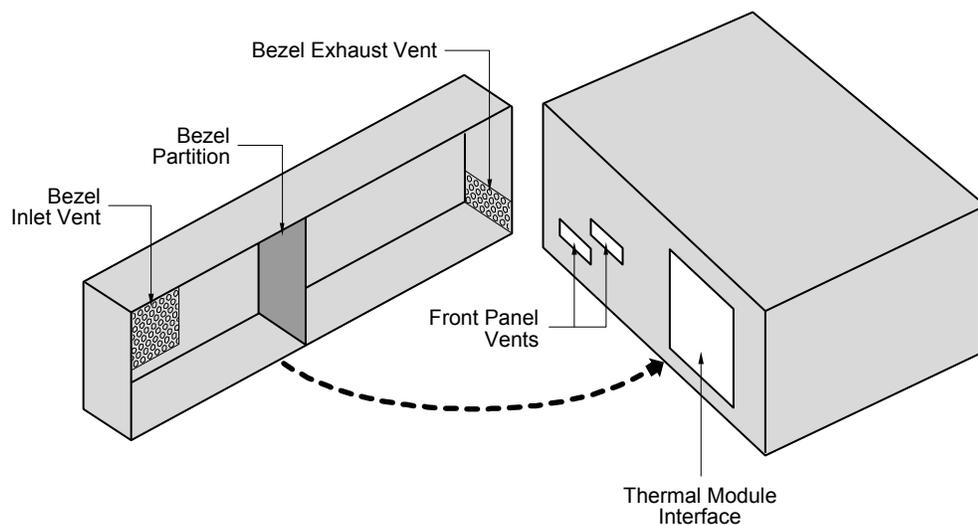
2.7 Case Study – Intel® Reference Design System Airflow Engineering

2.7.1 Thermal Module Effective Fan Curve

Intel fluid dynamic engineering resources have developed an efficient impeller and stator design that substantially improves the axial flow and pressure capability of the Thermal Module fan. The exceptional performance allows the Thermal Module fan to provide sufficient airflow to the CPU heatsink and system at a lower fan speed. Use of less efficient fan and /or stator designs will require higher operating speeds.

2.7.1.1 Type I Reference Design Thermal Module

The 90-mm fan engineered by Intel has the maximum operating speed stand-alone fan curve shown in Figure 45.

Figure 45: Type I Reference Design Thermal Module Fan - Stand-Alone Fan Curve
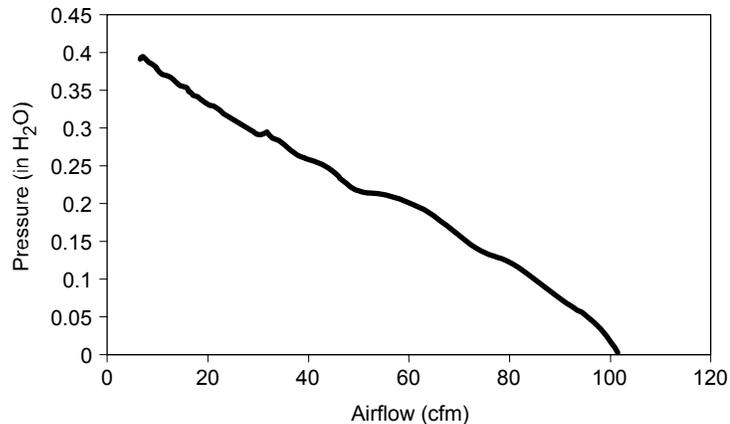
OM16792

The impedance curve for the LGA775, Pentium® 4 Processor Platform Compatibility Guide '04 B compliant Voltage Regulation, and Intel engineered post-and-plate heatsink reference design have the impedance curves shown in Figure 46.

Figure 46: Type I Reference Design Thermal Module and RVP VR Impedance Curve
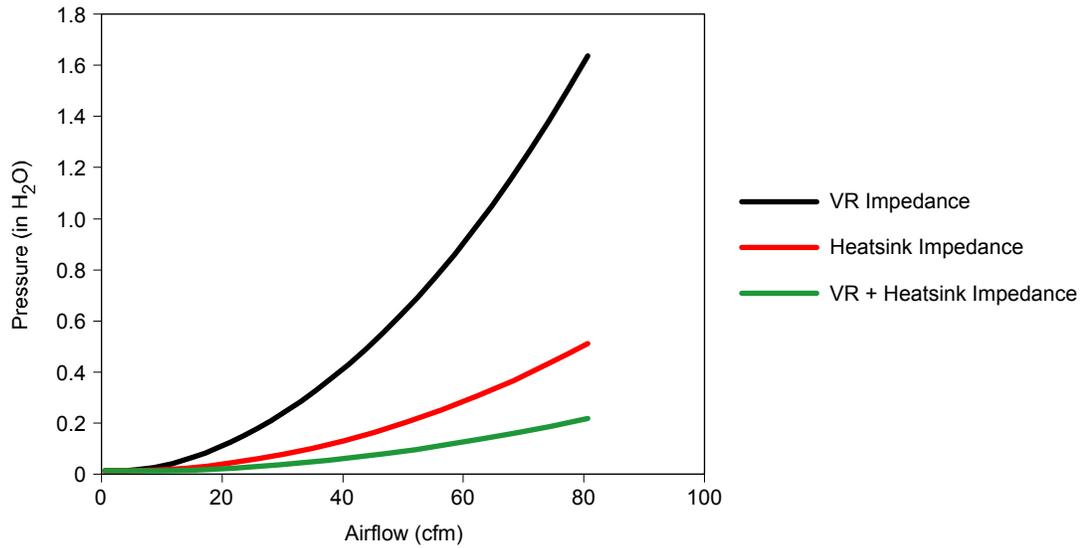
OM16793

The resulting reference design Type I Thermal Module Effective fan curve is created by the subtraction of the combined impedance curve from the stand-alone fan curve, as shown in Figure 47.

Figure 47: Type I Reference Design Thermal Module Effective Fan Curve
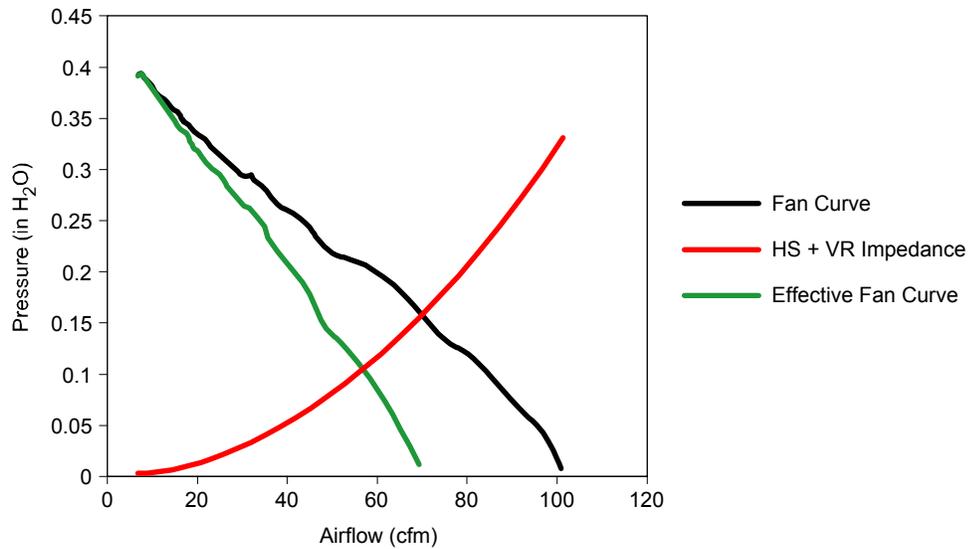
OM16794

The remaining elements in the 12.9 liter reference design system, engineered by Intel, has the impedance curve shown in Figure 48.

Figure 48: 12.9 Liter Reference Design System Impedance (Less Thermal Module)
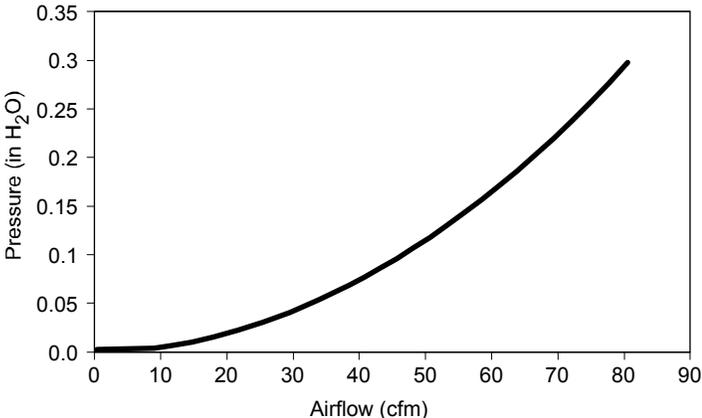
OM16795

The operating point for the 12.9 liter reference design system can be determined by the intersection of the reference design Thermal Module Effective fan curve and the system impedance curve, as shown in Figure 49.

Figure 49: 12.9 Liter Reference Design System Operating Point
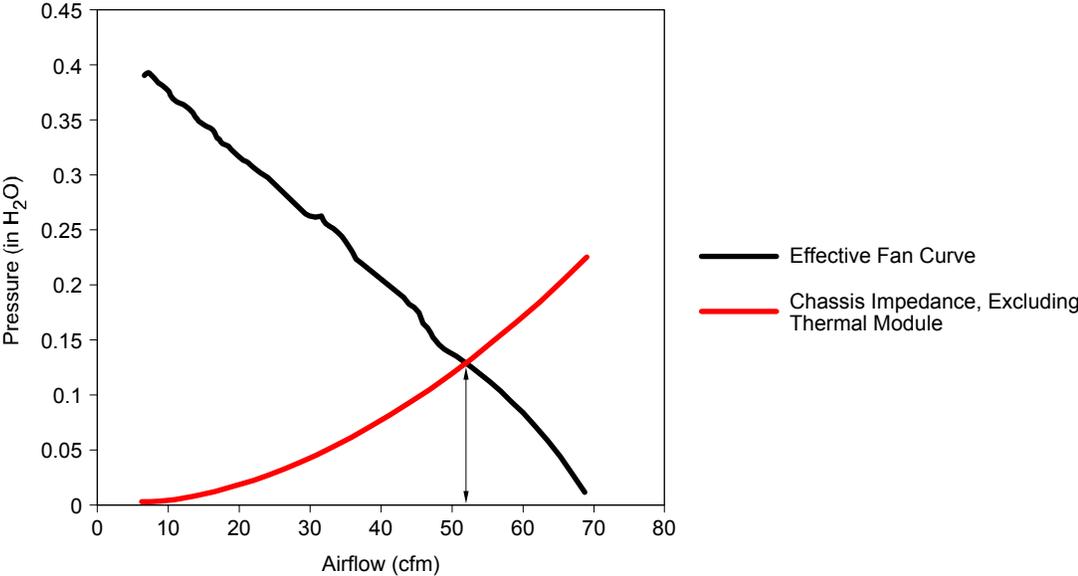
OM16796

2.7.1.2 Type II Reference Design Thermal Module

The 70-mm fan engineered by Intel has the maximum operating speed stand-alone fan curve shown in Figure 50.

Figure 50: Type II Reference Design Thermal Module Stand-Alone Fan Curve
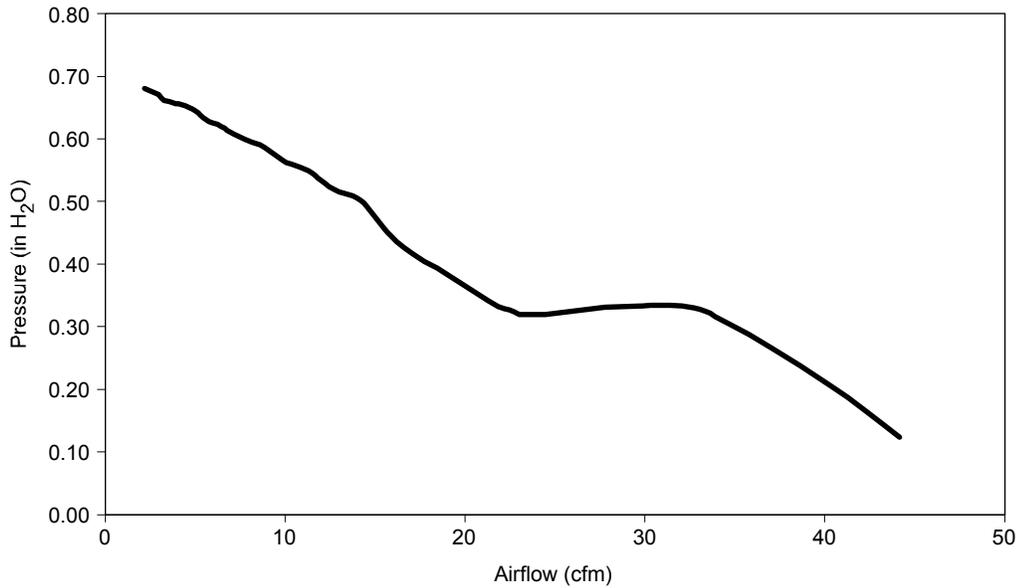
The impedance curve for the LGA775, Intel® Pentium® 4 Processor Platform Compatibility Guide '04 B compliant Voltage Regulation, and Intel engineered flat base vertical fin heatsink reference design have the following impedance curves shown in Figure 51.

Figure 51: Type II Reference Design Thermal Module and RVP VR Impedance Curves
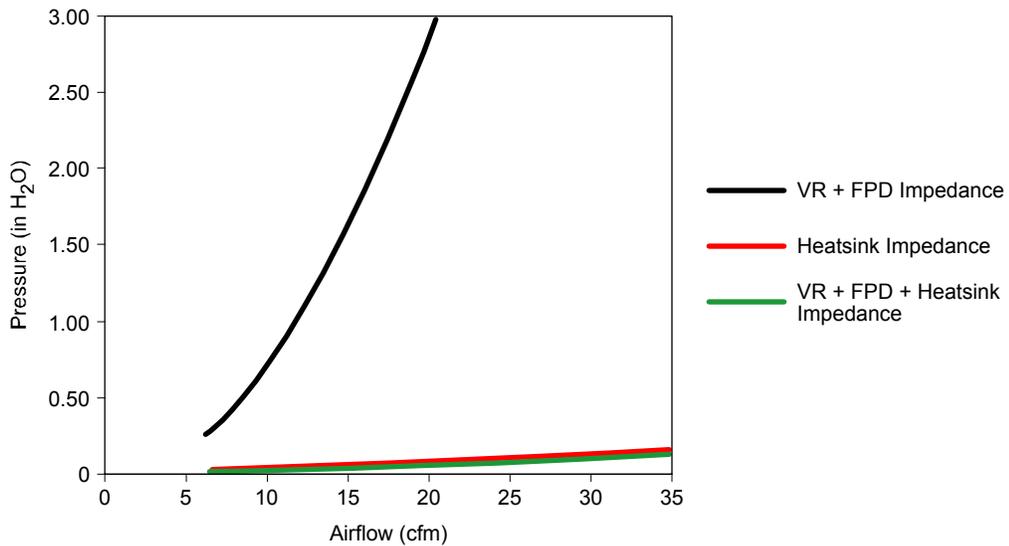
The resulting reference design Type II Thermal Module Effective fan curve is created by the subtraction of the combined impedance curve from the stand-alone fan curve, as shown in Figure 52.

Figure 52: Type II Reference Design Thermal Module Effective Fan Curve
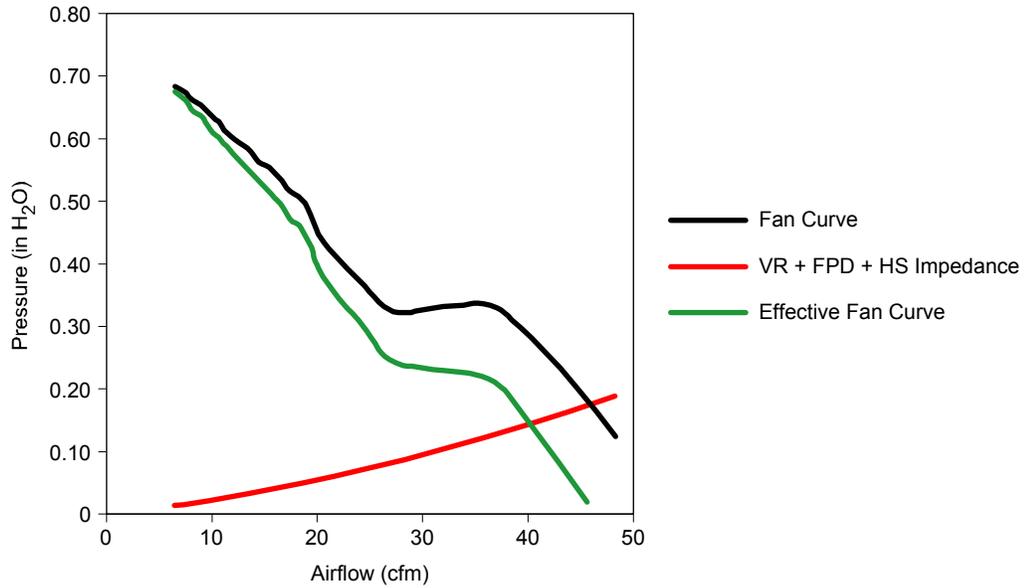
OM16799

The remaining elements in the 6.9 liter reference design system engineered by Intel have the impedance curve shown in Figure 53.

Figure 53: 6.9 Liter Reference Design System Impedance (Less Thermal Module)
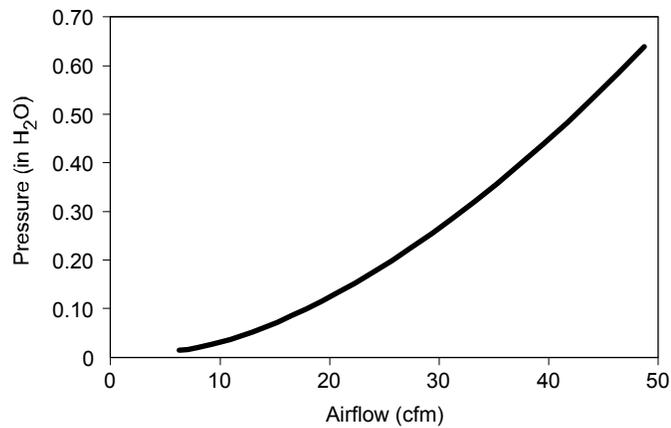
OM16800

The operating point for the 6.9 liter reference design system can be determined by the intersection of the reference design Thermal Module Effective fan curve and the system impedance curve, as shown in Figure 54.

Figure 54: 6.9 Liter Reference Design System Operating Point
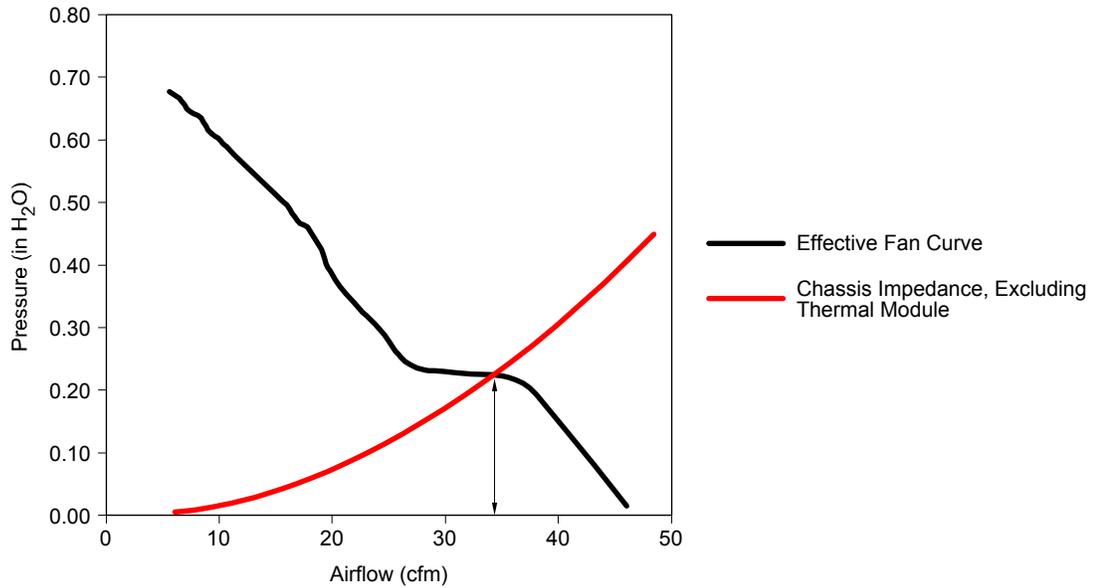
OM16801

2.7.2 Thermal Module Airflow Partitioning

The Intel reference design Type I Thermal Module uses a post-and-plate (vertical copper core and horizontal aluminum fin) configuration. When placed on a typical LGA775 BTX motherboard, the Thermal Module creates airflow that results in Table 4 natural airflow partitioning through the heatsink, above the motherboard, and below the motherboard.

Table 4: Type I Thermal Module Flow Partitioning

Operating Condition	Heatsink Airflow (CFM)	Above Motherboard Airflow (CFM)	Below Motherboard Airflow (CFM)
Maximum fan speed	39.0	7.1	6.2
Minimum fan speed	9.1	2.1	2.1

The Intel reference design Type II Thermal Module uses a more standard heatsink (horizontal copper base and vertical copper stacked fin) configuration. The resulting natural airflow partitioning does not provide sufficient airflow through the CPU heatsink without the use of a Flow Partition Device, as shown in Table 5.

Table 5: Type II Thermal Module Airflow Partitioning (With and Without Flow Partition Device)

Operating Condition	Heatsink Airflow (CFM)	Above Motherboard Airflow (CFM)	Below Motherboard Airflow (CFM)
Without Flow Partition Device			
Maximum fan speed	19.1	9.3	5.7
Minimum fan speed	4.3	2.1	1.3
With Flow Partition Device			
Maximum fan speed	28.8	3.4	2.1
Minimum fan speed	6.3	0.8	0.5

2.7.3 Subsystem Airflow and Temperature

The two Intel reference design systems use an integrated Fan Speed Control (FSC) circuit design that monitors the CPU diode temperature and a System Monitor Point thermistor. The temperatures reported to the FSC analog circuit are compared to their respective limit temperatures to determine the Pulse Width Modulation (PWM) signal that is sent to the Thermal Module fan. The PWM and voltage signal provided the fan motor will determine the Thermal Module fan speed. Obviously, the total system airflow and subsystem airflow will depend on the Thermal Module fan speed.

The three operating conditions in which the reference design performance has been predicted are outlined in Table 6.

Table 6: Standard Load and Ambient Temperature Descriptions Used for System Evaluation

	Representing	External Ambient	Power Load	Anticipated Fan Speed
Worst Case	Maximum required capability Extreme ambient and power loading	35 °C	All components operating at maximum power	PSU: maximum Thermal Module: maximum
Idle	Acoustic test condition Sustained typical environment and typical use	23 °C	All components operating at idle power	PSU: minimum Thermal Module: minimum
Typical Intermittent	Intermittent high power application in typical use environment	23 °C	All components operating at maximum power	PSU: maximum Thermal Module: intermediate

Both Intel reference design systems were engineered to meet Intel Pentium® 4 Processor Platform Compliance Guide '04 B guidelines in a performance system configuration. Feature descriptions, power loads, and assumed performance requirements are described in Table 7.

Table 7: Reference Design System Loads and Boundary Condition Assumptions

Component	Description	Qty.	TDP (watts)			Requirement Boundary Condition
			Worst	Idle	Inter	
Ambient	External		35 °C	23 °C	23 °C	
CPU	Platform Compatibility Guide '04 B	1	115	115	57.5	72 °C T_{case}
Voltage Regulation	Platform Compatibility Guide '04 B	1	38.3	38.3	19.2	91 °C motherboard under socket Component specification
MCH	Next Generation	1	17.7	17.7	6.9	T_j = 110 °C
ICH	Next Generation	1	4	4	1.2	T_j = 120 °C
Memory	DDR400		13.9	13.9	6.4	T_j = 93 °C
HDD	Standard	1	10	10	8	55 °C Natural Convection (NC)
ODD	12.9 L Standard	1	10	10	5	45 °C NC
	6.9 L Mobile	1	3	3	1	50 °C NC
FDD	12.9 L Mobile		0.3	0.3	0.3	55 °C NC
	6.9 L None					
x16 PCI Express	12.9 L Full Ht, 1/2 Length	1	75	75	37.5	104 °C T_{case}
	6.9 L None					
PCI Cards	12.9L Low Profile	2	5	5	3	55 °C NC
	6.9L None					
ExpressCard	12.9L None	1	2.1	2.1	1	65 °C NC
	6.9L Wide					
Motherboard		1	10	10	5	55 °C NC
PSU	CFX12V	1	280.9	280.9	142.5	50 °C Inlet
	LFX12V	1	226	226	112.2	

Table 8 illustrates the anticipated fan operating speeds and resulting airflow 9and temperature boundary conditions for the 12.9 liter reference design system.

Table 8: 12.9 Liter Reference Design System Airflow and Temperature Performance

		Worst	Idle	Intermittent
	Thermal Module Fan RPM	4500	900	2600
	PSU Fan RPM	3600	1550	3600
Anticipated Operating Point				
System Airflow	CFM	52.3	13.3	33
Operating point	Pressure (in. H_2O)	0.22	0.017	0.082
PSU Fan	CFM	17.7	6.4	15.8
CPU Heatsink Airflow	CFM	39	9.1	23.8
Above Motherboard Airflow	CFM	7.12	2.1	4.7
Below Motherboard Airflow	CFM	6.2	2.1	4.6
Subsystem Airflow and Temperature Boundary Conditions				
MCH	(LFM)/Temp (°C)	280/44	53/39	163/36
ICH	(LFM)/Temp (°C)	160/45	50/40	120/38
Graphics Add-In	(CFM)/Temp (°C)	3.5/43	1.2/34	3.2/32
ODD	Inlet Temp (°C)	43	39	36
HDD	Inlet Temp (°C)	46	41	39
FDD	Inlet Temp (°C)	48	44	41
PCI cards - LP	Inlet Temp (°C)	43	38	37
PSU	Inlet Temp (°C)	45	42	39
Component Temperature Performance				
CPU	Case Temp (°C)	71	45	53
MCH	Junction Temp (°C)	88	86	90
ICH	Junction Temp (°C)	112	60	105
PCI Express Graphics	Case Temp (°C)	103	87	95
DDR400 Memory	Case Temp (°C)	84	61	77

Table 9 illustrates the anticipated fan operating speeds and resulting airflow and temperature boundary conditions for the 6.9 liter reference design system.

Table 9: 6.9 Liter Reference Design System Airflow and Temperature Performance

		Worst	Idle	Intermittent
	Thermal Module Fan RPM	6800	1500	3400
	PSU Fan RPM	3000	1400	3000
Anticipated Operating Point				
System Airflow	CFM	34.4	7.6	16.9
Operating point	Pressure (in. H_2O)	0.274	0.017	0.071
PSU Fan	CFM	9.6	3.3	9
CPU Heatsink Airflow	CFM	28.8	6.3	14
Above Motherboard Airflow	CFM	3.4	0.8	1.8
Below Motherboard Airflow	CFM	2.1	0.5	1.1
Subsystem Airflow and Temperature Boundary Conditions				
MCH	(LFM)/Temp (°C)	352/45	62/38	164/43
ICH	(LFM)/Temp (°C)	483/50	97/47	235/45
Temperature				
ODD	Inlet Temp (°C)	43	45	43
HDD	Inlet Temp (°C)	43	50	43
PSU	Inlet Temp (°C)	48	47	53
Component Temperature Performance				
CPU	Case Temp (°C)	71	47	55
MCH	Junction Temp (°C)	86	81	97
ICH	Junction Temp (°C)	117	67	112
ExpressCard	Inlet Temp (°C)	48	53	64
DDR400 Memory	Case Temp (°C)	85	67	84

2.7.3.1 System Airflow Performance Sensitivities

2.7.3.1.1   Power Supply Operating Point

System airflow in the Intel reference designs is not strongly dependent on the PSU fan speed. The PSU fan is primarily responsible for providing sufficient airflow to the internal PSU components to maintain their operating temperature within specification.

Generally, at their maximum operating speeds, the PSU fan will evacuate approximately 30% of the airflow generated by the Thermal Module fan. Because the range of operating speeds for the Intel reference design Thermal Module designs are larger than that of the PSU fan, at their minimum operating speeds, the PSU fan will evacuate approximately 50% of the airflow generated by the Thermal Module fan. If the PSU FSC circuit were designed for the same proportional reduction in fan speed (from maximum to

minimum fan speed), the proportion of airflow that the PSU evacuates from the system would be similar at the maximum and minimum operating speed.

2.7.3.1.2   Thermal Module Interface Inlet Loss Coefficient

For the both the 6.9 liter and 12.9 liter reference design configurations, the Thermal Module Interface is provided by the front panel sheet metal; therefore, no internal system ducting is required. In addition to the impedance introduced by the front panel vent design, the airflow into the Thermal Module is impeded by the front bezel. Both reference design bezels were engineered to demonstrate the BTX acoustic value proposition. Specifically, the bezel open area for Thermal Module inlet airflow was located immediately in front of the Thermal Module Interface to eliminate turning losses and a louver inlet channel geometry was selected to minimize contraction and expansion losses.

As the Thermal Module interface inlet impedance is increased, the operating point airflow will decrease at any fixed fan speed. In order to meet the system temperature requirements, it will likely be necessary to compensate for increasing inlet impedance by increasing the Thermal Module fan speed. This, of course, increases the system's acoustic noise.

For the 12.9 liter reference design, the following table illustrates how an increase in bezel loss coefficient impacts the system operating point. Also included in the table are estimates of the Thermal Module fan speed increase required at the Worst and Idle operating conditions, and the impact on system acoustic Sound Power. For reference, the 12.9 liter reference design bezel has a bezel loss coefficient equal to 1.

Table 10: Impact of Bezel Impedance on System Operating Point, Thermal Module Fan Speed, and System Sound Power

Bezel Loss Coefficient	Operating Point (CFM)	Thermal Module RPM Required to Meet 52.3 CFM (Worst Case)	Thermal Module RPM Required to Meet 13.3 CFM (Idle Case)	Idle Case Sound Power (BA)
1	52.3	4500	980	3.85
2	51	4615	995	3.86
4	50	4707	1040	3.91
6	48	4903	1060	3.93
8	46	5116	1100	3.98
10	44	5349	1130	4.02
12	42.5	5538	1167	4.06
16	40.5	5811	1215	4.12
20	39	6035	1250	4.16

2.7.3.1.3   Ventilation

Increasing the ventilation open area in the rear panel will increase the operating point airflow and the proportion of airflow that exits the rear panel vents. As an illustration, in the 12.9 liter reference design, the PCI-Express x16 and x1 full height, half-length cards on a riser were replaced by two low profile cards perpendicular to the motherboard. With this card configuration change, the rear panel open area that is in line with the Thermal Module airflow is increased because the riser card rear panel slots are replaced with ventilation (see Figure 55 and Figure 56 for an illustration of vent designs for these two add-in card configurations).

Figure 55: 12.9 Liter Reference Design System with Graphics Add-in Card on a Riser

Power Supply Unit

Rear Vent 4

Rear Vent 1

Rear Vent 3

Rear Vent 2

Front Vent 1

Front Vent 2

Thermal Module
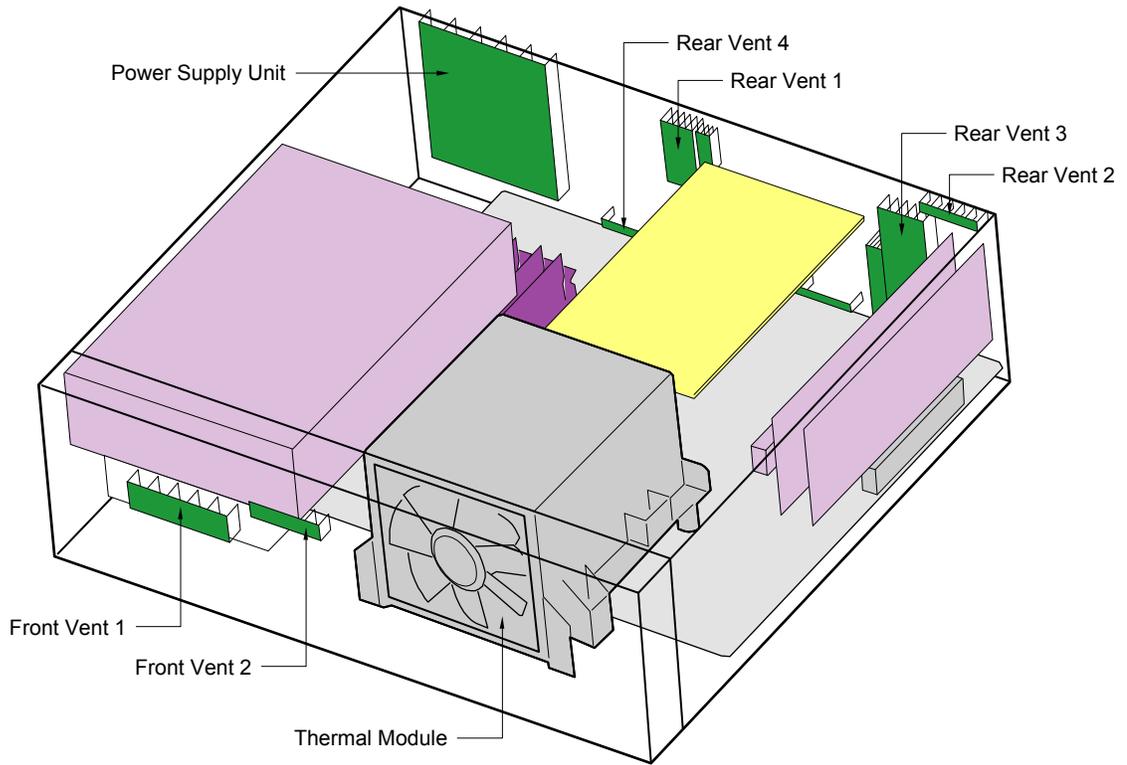
OM16802

Figure 56: 12.9 Liter System with Four Low Profile Cards
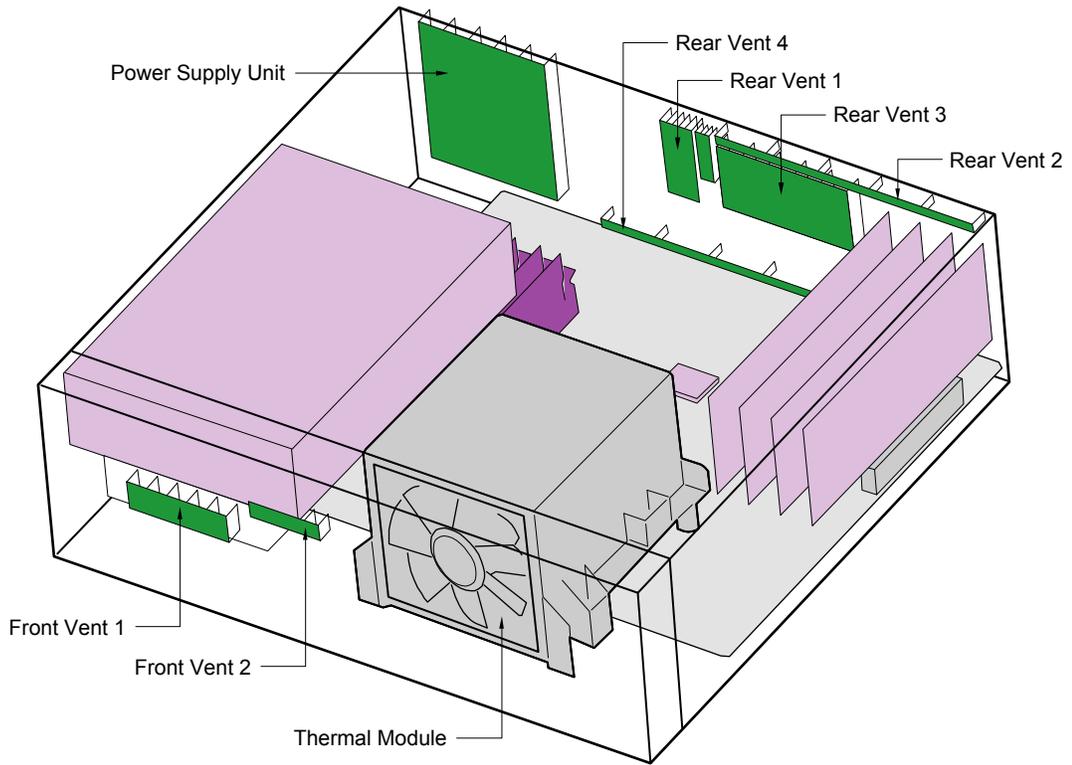
OM16803

The resulting 12.9 liter system operating point airflow and mass flow balance changes are reflected in Table 11.

Table 11: Mass Flow Balance for 12.9 Liter Reference Design System with Four Low Profile Cards

Description	Units	Graphics Card Position		
		Riser	Low Profile	% Change
System Airflow	CFM	52.3	54.5	4%
Operating Point	Pressure (in. H_2O)	0.22	0.21	-5%
PSU Fan	CFM	17.7	16.8	-5%
Rear Vent 1	CFM	-6.9	-6.4	-7%
Rear Vent 2	CFM	-1.1	-4.9	345%
Rear Vent 3	CFM	-17.5	-19.3	10%
Rear Vent 4	CFM	-3.3	-2.8	-15%
Front Vent 1	CFM	-3.8	-2.8	-26%
Front Vent 2	CFM	-1.9	-1.4	-26%

Front panel ventilation plays an important role in managing drive bay temperatures. For example, in the CFD analysis of a tower system drive bay spacing, the drive bay temperatures increased by 7-10 °C if it used the same front panel ventilation as the 12.9 liter reference design. Increasing the ventilation open area will reduce the drive bay temperatures in the tower configuration. Front panel ventilation reduces airflow stagnation in the drive bay area, especially in configurations where the drive bays are farther away from the airflow created by the Thermal Module.

2.7.3.2 MCH Performance Sensitivities

MCH heatsink performance is a strong function of inlet airflow speed and heatsink bypass. The Intel reference design MCH heatsink performance as a function of airflow is shown in Figure 57.

Figure 57: MCH Heatsink Performance versus Inlet Airflow
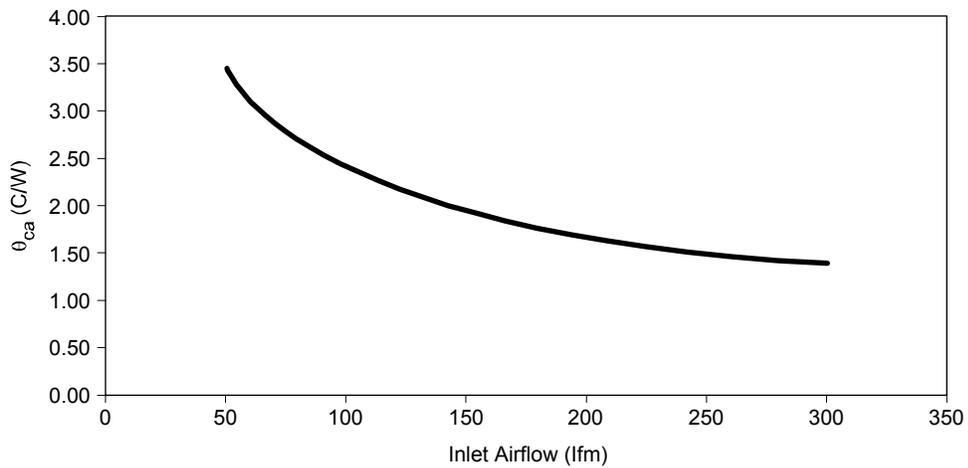
OM16804

MCH heatsink inlet airflow is a function of Thermal Module operating point. Table 12 illustrates the system operating point airflow and MCH heatsink inlet airflow in the 6.9 liter and 12.9 liter Intel reference design systems as a function of Thermal Module fan speed. Note that the 12.9 liter system uses a Type I Thermal Module and the 6.9 liter system uses a Type II Thermal Module.

Table 12: MCH Heatsink Inlet Airflow versus System Operating Point

	Thermal Module Fan Speed (RPM)	Operating Point (CFM)	MCH Heatsink Inlet Airflow (LFM)
12.9 Liter System			
	4500	52.3	280
	2600	33.0	163
	1200	13.3	53
6.9 Liter System			
	6800	34.4	352
	3400	16.9	164
	1500	7.6	62

Note, in Figure 58, that despite the lower operating point in the 6.9 liter system (relative to the 12.9 liter system), the MCH heatsink inlet airflow is higher. This is due to the difference in MCH heatsink bypass in the 6.9 liter reference design, which has the HDD above the MCH heatsink, versus the 12.9 liter reference design, which has the x16 and x1 PCI Express graphics add-in cards above the MCH heatsink.

Figure 58: MCH Heatsink Inlet Airflow versus System Operating Point
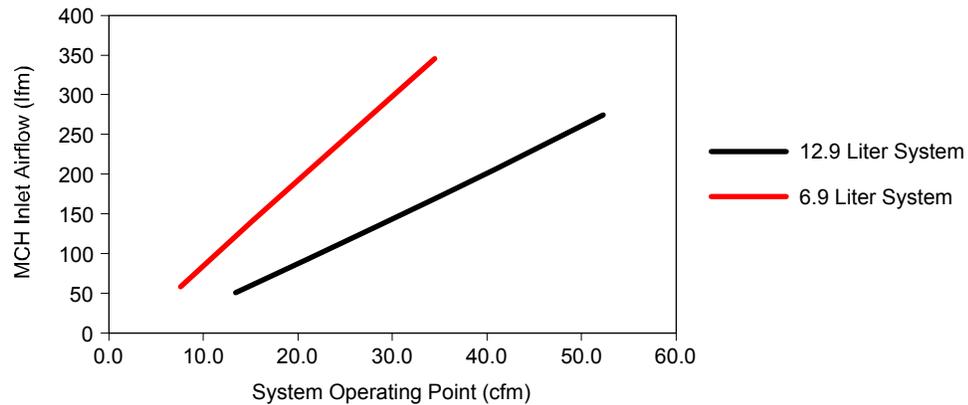
OM16805

2.7.3.3 Graphics Add-in Card Performance Sensitivities

The 12.9 liter reference design illustrates the use of a high performance graphics add-in card in a low profile system. Since it is unlikely that low profile high performance graphics will become available, the full height half length PCI-Express card was put on a riser so that it is parallel to the motherboard. In this position, the powered components on the card are facing into the Thermal Module exhaust airflow path. In this orientation, the cards block potential rear panel ventilation area above the Rear Panel I/O; the turning impedance introduced reduces the system operating point airflow.

In tower system configurations, the system thickness will allow the full height PCI-Express card to be placed directly in slot 1 so that it is perpendicular to the motherboard. In this position, the powered components on the card are also facing into the Thermal Module exhaust airflow path. This perpendicular card orientation allows more ventilation above the Rear Panel I/O because the rear panel area blocked by the riser card openings can be replaced with vent area. This decrease in rear panel impedance increases the system operating point airflow.

When the graphics add-in card is moved to the perpendicular position in slot 1, the Thermal Module exhaust airflow experiences more extensive expansion due to the increase in expansion open area. The turning associated with that expansion creates some amount of impinging airflow on the graphics card heatsink, in addition to the increased axial flow that results from lower rear panel impedance. The increase in airflow in the parallel and impinging direction improve graphics card heatsink performance. The graphics add-in card heatsink illustration in Figure 59 establishes the heatsink surfaces through which airflow can pass. Table 13 shows the expected graphics heatsink airflow at the maximum operating speed of the Thermal Module fan (note that positive values indicate airflow into the surface and negative values indicate airflow exhausting from the surface).

The important conclusion from this analysis is that rear panel ventilation near the graphics add-in card heatsink exhaust will substantially improve graphics card heatsink performance. In the 12.9 liter system, the proximity of the graphics card heatsink to the Rear Panel I/O leaves no room for rear panel ventilation,

and the airflow that enters the graphics card heatsink exhausts out of the top instead of the rear of the heatsink, reducing the heatsink convective performance.

Figure 59: Graphics Add-in Card Illustration of Airflow Surfaces
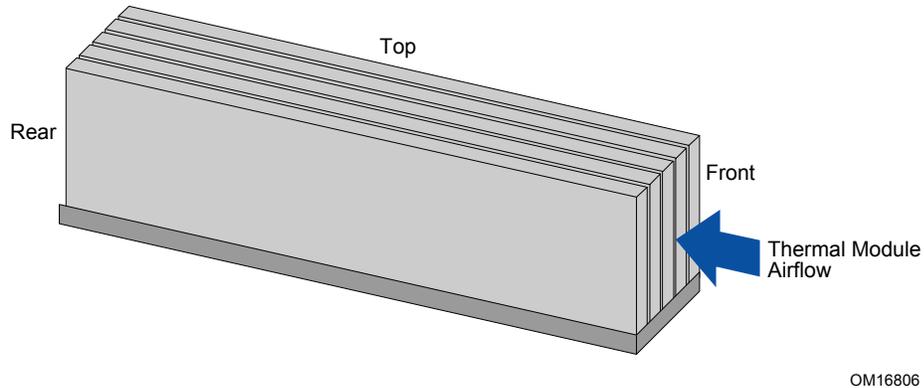
OM16806

Table 13: Graphics Add-in Card Heatsink Airflow and Performance versus System Configuration

Configuration	System Operating Point Airflow (CFM)	Front Airflow (CFM)	Top Airflow (CFM)	Rear Airflow (CFM)	Heatsink Performance ($^{\circ}$C/W)
12.9 liter Full-height, half-length card on riser	52.3	3.5	-3.1	-0.4	1.1
Tower Full height card in slot 1	56.6	0.4	0.8	-1.2	0.9

2.7.3.4 Voltage Regulation Performance Sensitivities

Intel engineered the CPU Voltage Regulation on the BTX Reference Validation Platform (RVP) to be compliant with the Pentium 4 Processor Platform Compatibility Guide '04 B requirements (refer to the processor datasheet for more information). A complete description of the initial three phase RVP CPU Voltage Regulation design is available in the RVP Drop-In Core Layout and RVP Schematic.

Numerical thermal analysis of Voltage Regulation component, motherboard, and socket was conducted and the predicted temperatures as a function of airflow are shown in Table 14 and Figure 60. The temperature predictions were made from a detailed numerical model that included:

- Motherboard, socket, voltage regulation components material conduction properties

- Motherboard, socket, voltage regulation component, and SRM physical size and location

- Type I reference design Thermal Module fan curve and velocity profile

- Input power for voltage regulation components, motherboard power planes and traces, and the socket based on characterized efficiencies (or impedances) at Platform Compatibility Guide '04 B maximum current.

The motherboard, socket, and voltage regulation component temperatures are the maximum predicted temperature for the worst-case location.

Table 14: Voltage Regulation Component Temperature versus Voltage Regulation Airflow

	Voltage Regulation - Component Descriptions					
	Drivers	**High Side FETs**	**Low Side FETs**	**Motherboard**	**Output Inductor**	**Socket**
Temperature Specification (C)	125	150	150	120	85	91
Inlet Airflow (CFM)	Voltage Regulation Component Temperatures					
2.8	108	134	130	117	100	75
6.0	96	120	117	103	86	65
9.4	88	110	106	94	78	61

Figure 60: Voltage Regulation Component Temperature versus Voltage Regulation Airflow
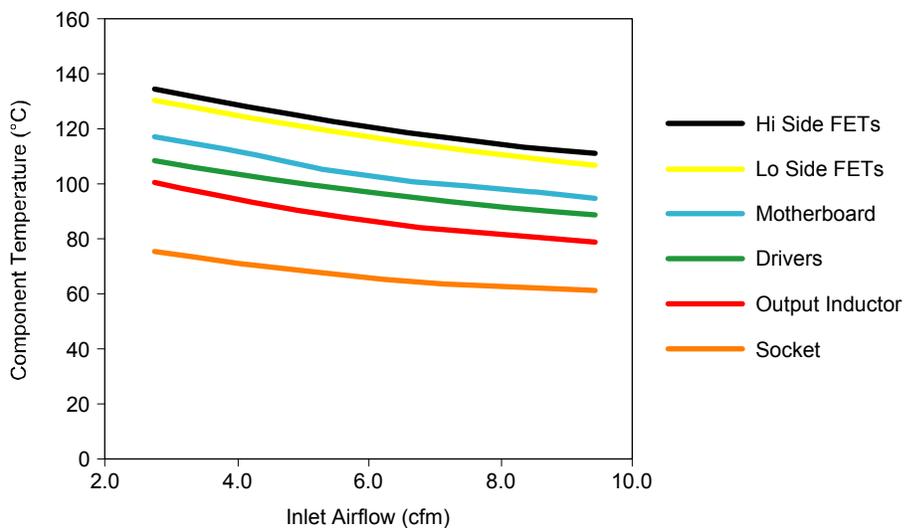
OM17071

Note that Voltage Regulation temperatures increase as inlet flow to the area is reduced. Voltage Regulation inlet flow is a function of the Thermal Module operating point airflow. Therefore, it is important to understand the Fan Speed Control circuit's role in managing Voltage Regulation temperature.

In a system that does not regulate the Thermal Module fan speed, the fan speed and operating point airflow will always be at their maximum. In systems where the Thermal Module fan speed is regulated, one of the temperature inputs to the Fan Speed Control circuit will be the CPU diode temperature. Since the CPU diode temperature will increase as CPU power increases, the Thermal Module fan speed and operating point airflow will be proportional CPU input power. CPU input power is provided by CPU Voltage Regulation, meaning that the temperature of Voltage Regulation will be strongly correlated to the temperature of the CPU. Therefore, there will not be conditions in which the Thermal Module fan is not providing sufficient airflow to manage Voltage Regulation temperature because if the Thermal Module fan

is providing sufficient airflow to meet the CPU temperature requirement, a properly design Thermal Module will simultaneously provide the appropriate airflow to Voltage Regulation.

The Intel® 2004 Performance BTX RVP Voltage Regulation supports processors that require VRD10.1 (Reference Yellow Cover Document #14294). Refer to the Voltage Regulator-Down (VRD) 10.1 Design Guidelines for Socket 775 Based Desktop and Transportable Motherboards for more information. Table 15 below is an excerpt from VRD10.1.

Table 15: VRD10.1 Requirements

VR Configuration	Iccmax	VR TDC	Dynamic Icc	R_{LL}	TOB	Maximum VID
775_VR_CONFIG_04A	78 A	68 A	55 A	1.40 m	+/-25 mV	1.4 V
775_VR_CONFIG_04B	119 A	101 A	95 A	1.00 m	+/-19 mV	1.4 V

To meet the VRD 10.1 requirements, the 2004 Performance BTX RVP implements a 3 phase Voltage Regulation that is similar to those implemented on the Intel 2004 Performance ATX RVP. BTX allows Voltage Regulation cost optimization relative to ATX because of the improved thermal environment. Cost reduction is achieved by using less efficient components. Since more Voltage Regulation power can be dissipated in a BTX system, less efficient FET components can be used to support a given Total Design Current (TDC). A 3 phase BTX Voltage Regulation is approximately 15% cheaper than an equivalent solution in ATX.

Additional Voltage Regulation cost reduction opportunities depend on several key factors, including system fan speed setpoints and AC-to-DC power supply limitations. First, at lower fan speeds the airflow to through BTX Voltage Regulation limits the power that can be dissipated. Second, as Voltage Regulation efficiency decreases, the amount of power required from the PSU 12V rail increases. The board and system designer must, therefore, work together to determine the appropriate component efficiencies. Figure 61 illustrates input current versus the power delivered to the processor across a range of component efficiencies.

Figure 61: Voltage Regulation Power Vs Current
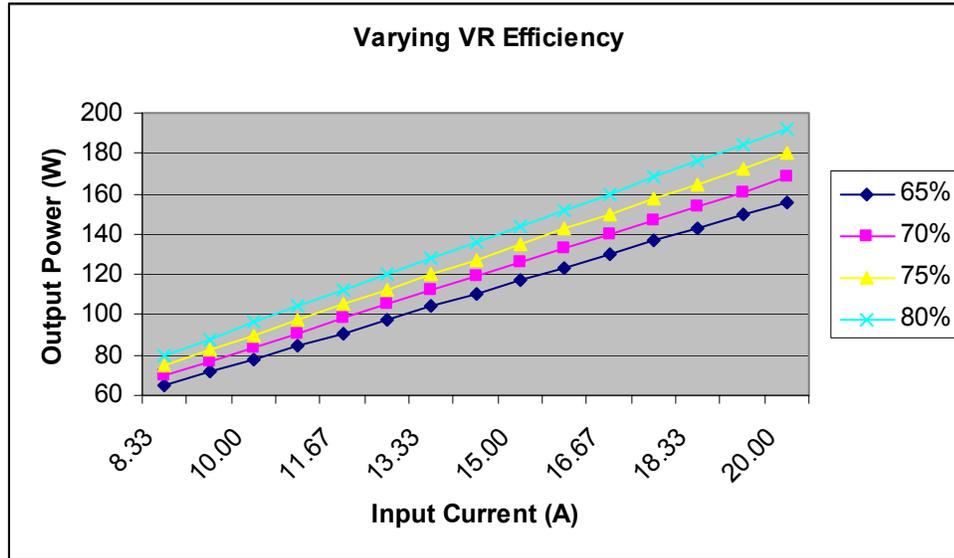
The 2004 Performance BTX RVP has been validated to meet thermal and electrical performance requirements with 2 phase Voltage Regulation in the 12.9 liter reference design system. A typical Voltage Regulation controller requires that phases be shut down in a certain order. The initial 3 phase design used on the BTX RVP, however, did not allow operation as a 2 phase VR with component changes only. Figure 62 shows the phase order on the 2004 Performance BTX RVP.

Figure 62: BTX RVP Voltage Regulation Phase Diagram
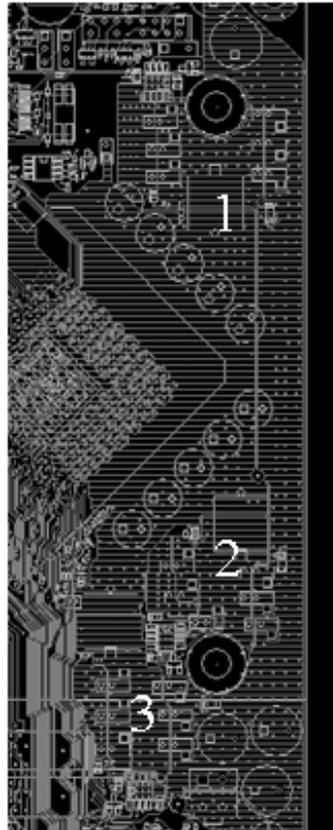
A design that would allow 2 or 3 phase operation would require the following:

1. Set phase order per Figure 62.
2. Dual site the output inductors for phases 1 and 2 to allow for higher current inductors.
3. Move thermal compensation thermistor between phases 2 and 3 per Figure 63.
4. Place feedback components around the controller based on the number of phases implemented.

Figure 63: BTX RVP 2 and 3 Phase Operation
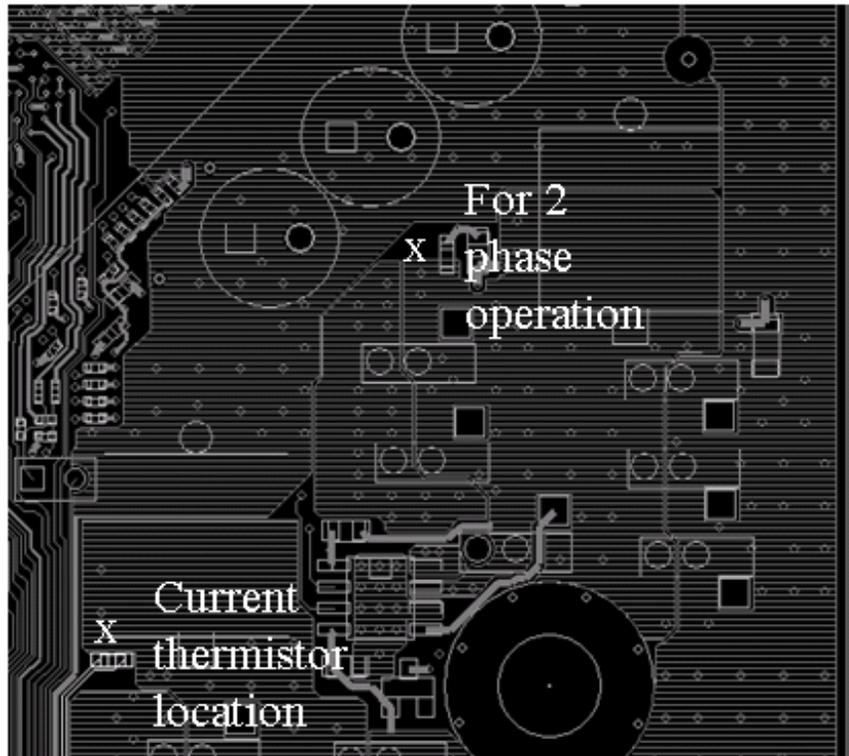
The potential cost advantage of a 2 phase solution relative to 3 phase will depend on component pricing. Intel recommends that board layouts allocate sufficient space and pursue design practices that would allow 3 phase implementations, taking into account the guidance offered above that may allow an easy migration to a 2 phase implementation.

2.7.3.5 Memory Performance Sensitivities

Memory is the most sensitive subsystem in a BTX system design from a thermal performance perspective. Airflow to the memory channels is influenced by not only the Thermal Module operating point airflow, but also by:

- PSU location
- CPU Voltage Regulation open area (see Section 2.4.5.2)
- Voltage Regulation exhaust airflow management
- Bypass around the memory channel
- Rear panel ventilation
- Memory position

2.7.3.5.1 Operating Point

The approach speed of airflow to the memory subsystem is a function of Thermal Module operating point, as illustrated in Table 16 and Figure 64.

Table 16: Memory Airflow versus System Operating Point

	Thermal Module Fan Speed (RPM)	Operating Point (CFM)	Memory Inlet Airflow (LFM)
6.9 Liter System			
	6800	34.4	223
	3400	16.9	153
	1500	7.6	55

Figure 64: Memory Airflow versus System Operating Point
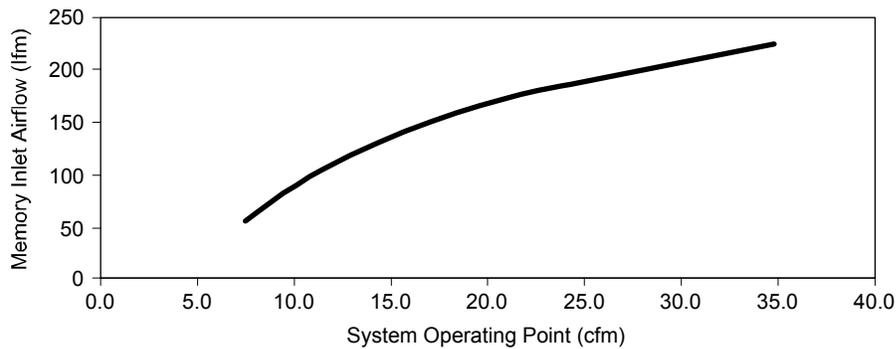
OM17072

2.7.3.5.2 PSU Location

In both the 6.9 and 12.9 liter reference designs, the PSU is located near the rear panel and immediately next to the Zone C fixed motherboard edge – the board edge closest to the memory components. Additionally, both reference design Power Suppliers – the Desoto CFX12V PSU and the Stockbridge LFX12V PSU – use a pressurizing fan. The proximity of the PSU to the memory location in both reference systems improves the airflow in the memory channels.

In CFD analysis of a BTX tower configuration, Intel chose to place the PSU with a pressurizing fan at the rear of the system and at the Zone C board edge. As indicated in the following table, the channel airflow did not change substantially. As noted in Section 2.7.3.5.6, rear panel ventilation can influence memory channel airflow. The fact that the tower CFD analysis indicates equivalent memory channel airflow to the 12.9 liter reference design is due to the dominant effect of the PSU.

2.7.3.5.3 Voltage Regulation Exhaust Airflow Management

The airflow that goes through CPU Voltage Regulation and toward memory will approach the memory channels at an oblique angle. The amount of airflow that enters the memory channel from Voltage Regulation will increase if there are any components that force it to turn toward the rear of the system as it nears memory. For instance, a BTX tower configuration may place drive bays perpendicular to the motherboard near the front of the system and along the Zone C motherboard edge (see Figure 65).

Figure 65: BTX Tower
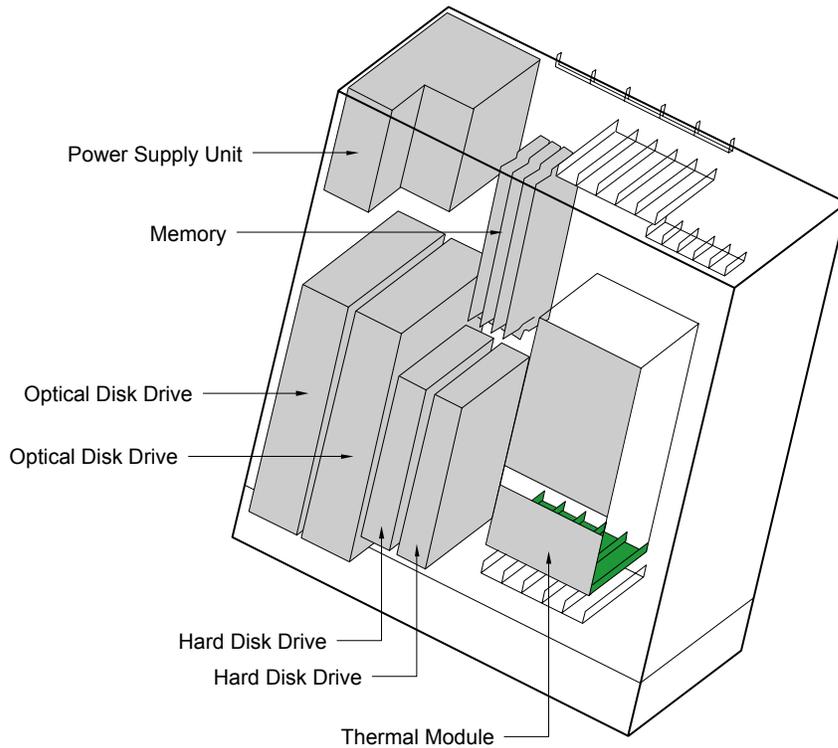
The HDD in Figure 65 nearest the Thermal Module will turn the airflow that exhausts from the Voltage Regulation toward the rear of the system and increase the airflow that goes into the memory channel. Recall from Section 2.7.3.5.2 that the a PSU near the Zone C board edge will create airflow that enters the memory channel from the rear. Therefore, the airflow that enters the front of the memory channel from Voltage Regulation and the airflow that enters the rear of the memory channel from the rear panel will meet and become stangant at some position in the memory channel. This stagnation point will be the location where the memory component temperatures will be the highest. Figure 66 illustrates how the stagnation point changes based on the strength of airflow that enters the front of the memory channel.

Figure 66: Memory Channel Airflow Pattern – Dependence on Voltage Regulation Exhaust Airflow Strength
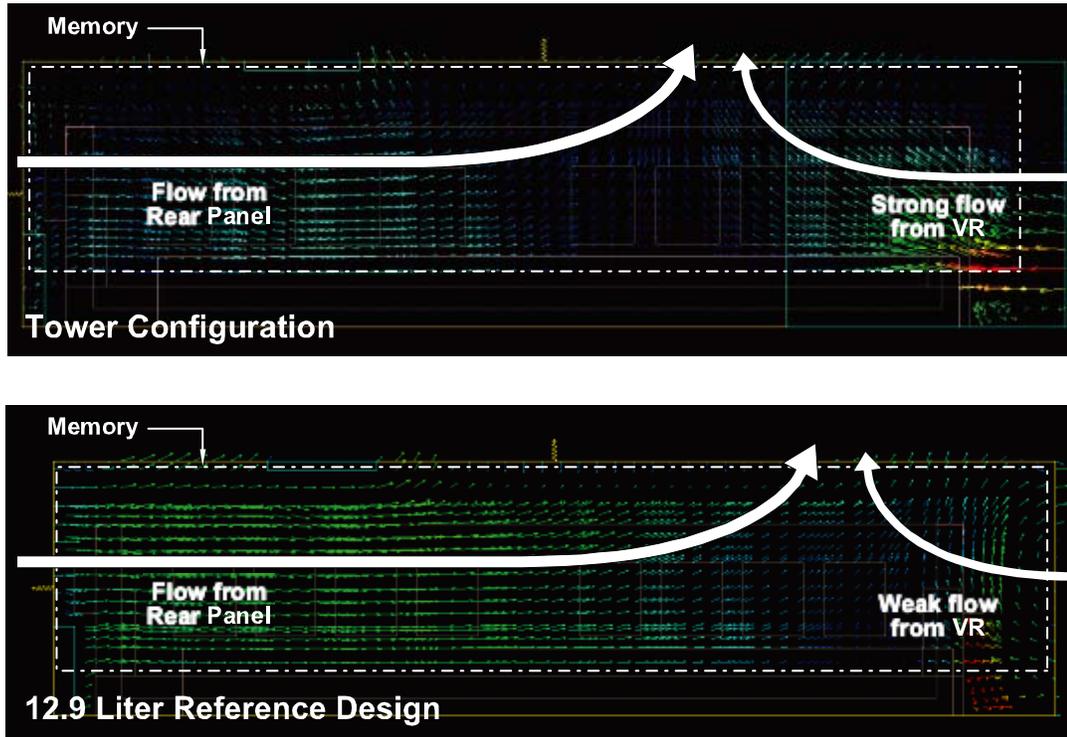
OM17074

2.7.3.5.4   Voltage Regulation Open Area

Voltage Regulation component frontal area and position can effect the memory inlet airflow speed. Taller, wider components in the airflow path between the Thermal Module fan and memory inlet will reduce inlet airflow speed. Intel will include more information on the impact based on two-phase versus three-phase Voltage Regulation thermal and fluid dynamic analysis and testing in future revisions of this guide.

The approach speed of airflow to the memory subsystem is a function of Thermal Module operating point, as illustrated in the following table and graph from the 6.9 liter reference design system CFD analysis:

2.7.3.5.5   Memory Bypass

Cables, drive bays, and brackets are examples of components that can create bypass over the memory components. The proximity of these components to memory will influence the airflow behavior in the memory channel. For example, in the 6.9 liter reference design, the Hard Disk Drive (HDD) bracket spans from the left side panel to the right side panel, and a portion of the bracket extends over memory. This bracket effectively forces that Voltage Regulation exhaust airflow and Thermal Module exhaust expansion airflow to stay in the memory channels. Figure 67 illustrates how memory channel airflow changes in the 6.9 liter reference design as the HDD bracket position above memory is moved toward the top panel.

Figure 67: Memory Channel Airflow versus Drive Bracket Bypass
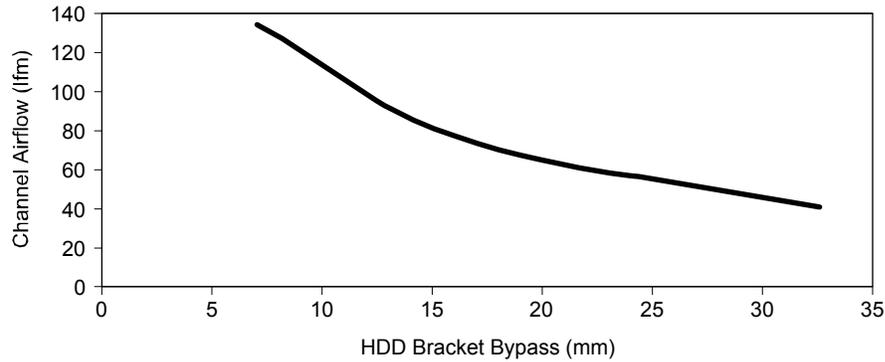
OM17075

Cables that extend from the PSU or drives are often routed over memory as they transition from the component to their motherboard connectors. These cables act as bypass for airflow that enters the memory channel from the Voltage Regulation or Thermal Module exhaust expansion.

2.7.3.5.6   Rear Panel Ventilation

When the rear panel is well ventilated, the Thermal Module exhaust airflow has a low impedance path that allows it to cleanly exit the system. As discussed previously, this can improve the system operating point airflow. However, this may also change the memory channel airflow. If the rear panel is not well ventilated (as is the case in the 12.9 liter reference design system where the PCI-Express cards on the riser block potential vent area), a larger portion the Thermal Module exhaust airflow is forced to turn at the rear panel. In the 6.9 liter and 12.9 liter reference designs, this airflow is then pulled through the memory channel by the PSU fan. This airflow pattern – from the rear of the system into the rear of the memory channel – is strongly influenced by the location of the PSU and the PSU fan curve.

The following table illustrates the impact on memory channel airflow and worst case DDR-400 memory component temperature in the 12.9 liter reference design system when the two riser PCI-Express cards are replaced by two low profile cards, and the rear panel in that area is ventilated.

Also included in the table is the memory channel airflow and temperature behavior for a full tower configuration; one in which the full height PCI Express card is placed in slot 1 perpendicular to the motherboard. In this configuration, the rear panel is ventilated above the Rear Panel I/O.

Table 17: Memory Airflow and Maximum Component Temperature versus System Configuration

Configuration	Memory Inlet Airflow (LFM)	Memory Channel Airflow (LFM)	Memory Component Temperature ($^\circ$C)
12.9 liter Two full height, half length cards on riser	40-80	50	84
12.9 liter Two low profile cards	50-80	15	88
Tower Full height card in slot 1	50-80	50	84

2.7.3.5.7   Memory Position

As memory moves closer to the front of the motherboard, less airflow will be able to enter the memory channels. The below heatsink /above motherboard airflow is diverted toward memory by the CPU socket. If the socket is oriented perpendicular to the fan axial flow direction, the flow speed diverted toward memory will be lower. If the socket is angled, the flow speed diverted toward memory will be higher. Each memory channel will receive more of this airflow if the memory channel opening is positioned in this flow path. See Figure 24, Figure 25, and Figure 34 for illustrations of airflow patterns that enter memory and memory thermal analysis boundary condition descriptions, for reference.

2.8 Case Study – BTX Tower PC Configuration

A BTX tower system was designed and numerically evaluated to determine if a large volume, high performance tower system could be adequately cooled without a rear panel fan and using the Intel® Type I Reference Design Thermal Module Assembly with a 90mm fan.

The system mechanical design was based on conventional placement of subsystem components. A thermal numerical model was constructed to analyze the airflow and subsystem temperatures.

2.8.1   BTX Tower Design and Thermal Model Construction

Figure 68 illustrates the component placement and vent scheme for the 36 liter system. Directly below the ODD drive bay are 3.5" drive bays (including one that sits over the Thermal Module) for multiple FDD and HDD placements. An ATX12V PSU is located immediately behind the drive bays, near the intersection of the side and rear panel.

The thermal numerical model also included a 7-slot BTX board with the component power descriptions from Table 7. The Thermal Module Interface complies with the BTX Interface Specification requirements at the appropriate location, but widens to 120mm square at the front panel.

Figure 68: BTX Tower Component Placement and Vent Scheme
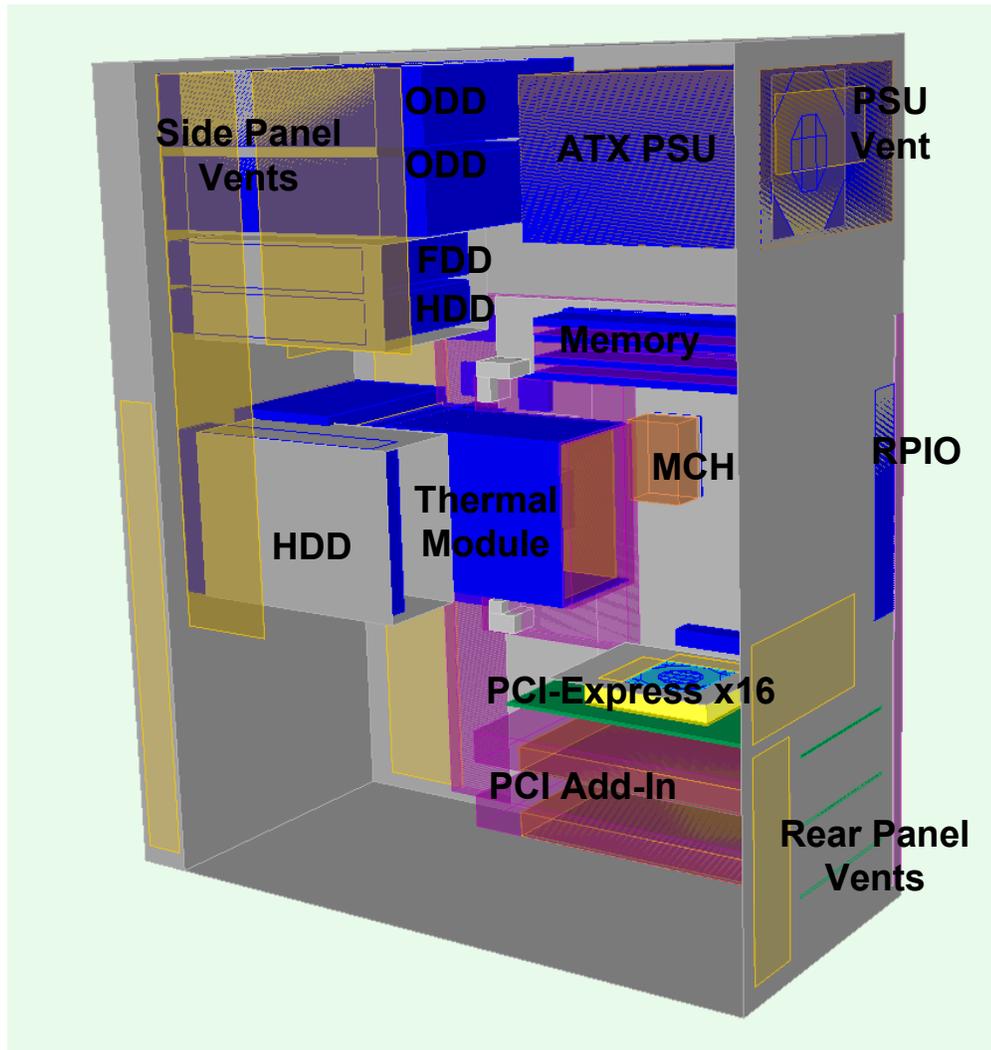
2.8.2 BTX Tower Thermal Simulation – Vent Design Investigation

Figure 69 illustrates the airflow temperature parallel to the board through a cut plane at the midpoint of the TMA and PSU fans.

Figure 69: BTX Tower Thermal Simulation Airflow and Temperature Plot (Top View)
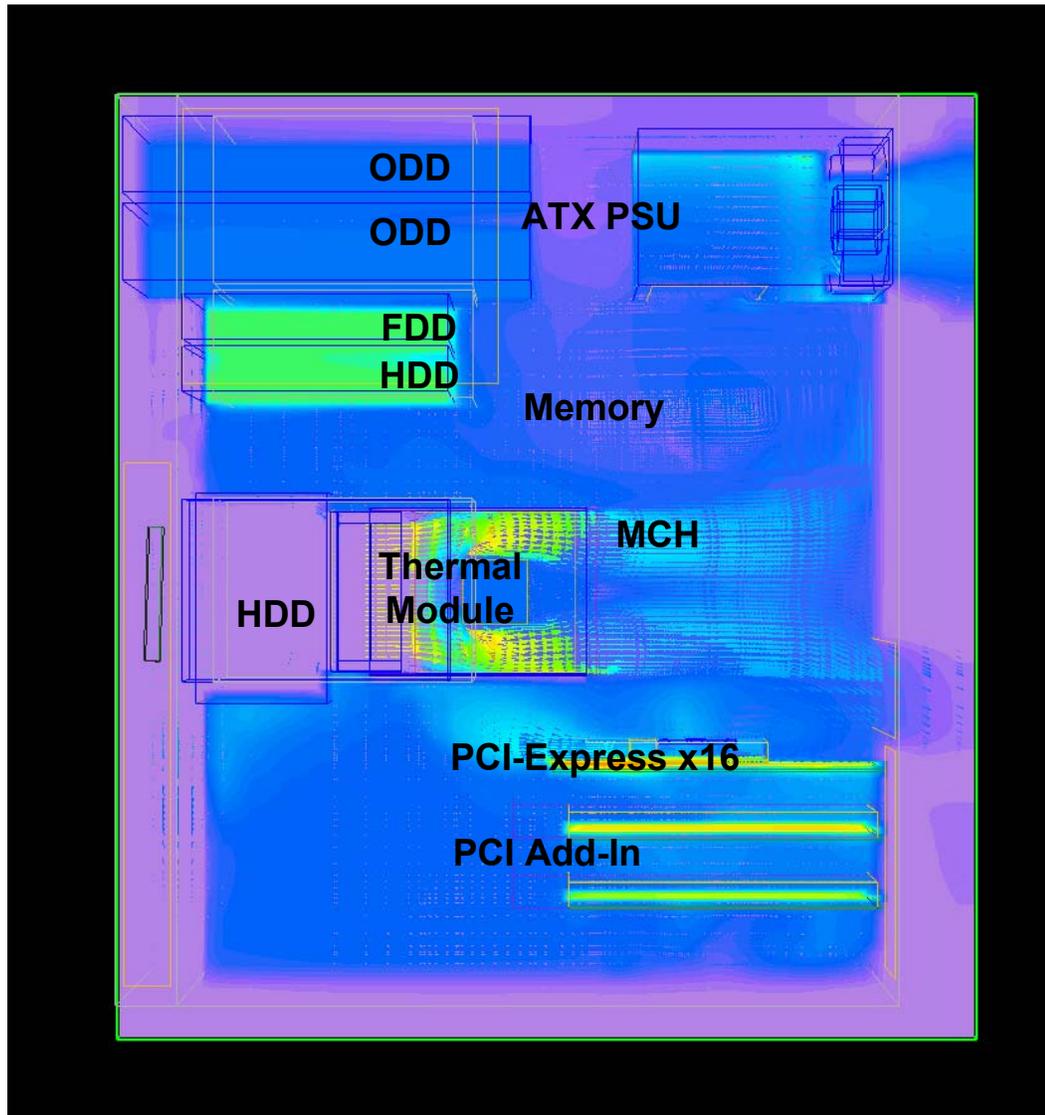
Figure 70 illustrates the airflow temperature through a cut plane at the midpoint of the TMA and PSU fans but perpendicular to the board. It also includes a particle source trace that illustrates the flow pattern entering the bezel, going through the Thermal Module to the rear panel, and exhausting from the front and side panel after moving through the drive bay area. Of particular note is that the front bezel uses a side inlet instead of a front inlet position, which illustrates that the system performance targets can be met with a bezel whose front surface is incompromised by inlet venting.

Figure 70: BTX Tower Thermal Simulation Airflow and Temperature Plot (Side View)
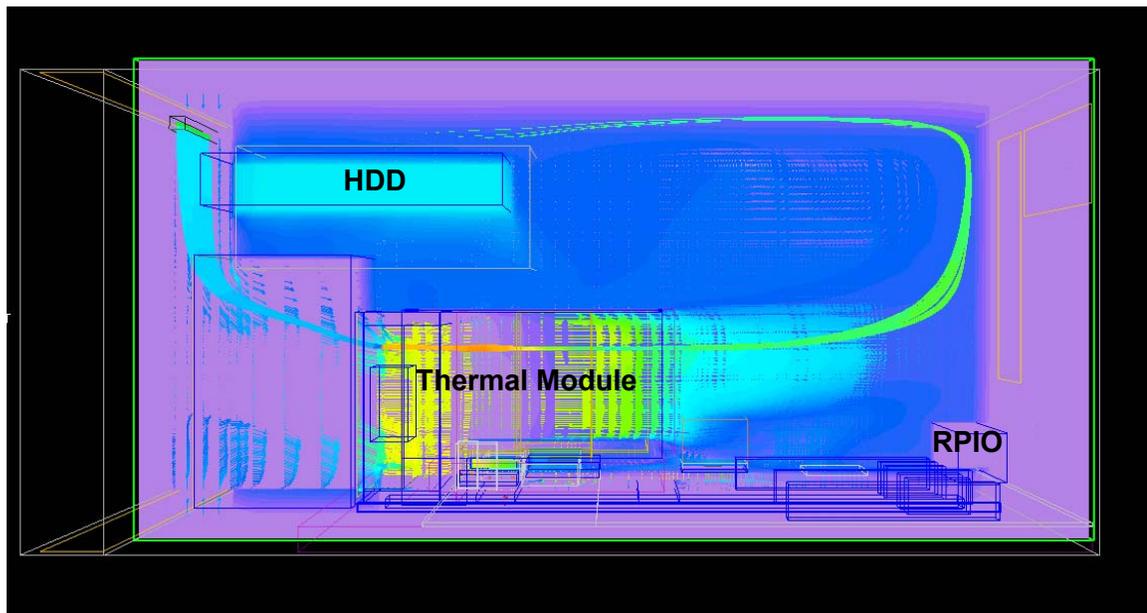
Several side panel vent options were evaluated and the results indicate that drive bay component performance can improve with side panel ventilation. Figure 71 shows the airflow pattern with a vent near the drive bays that extends the entire length of the side panel. Since the exhaust area extends along the entire length of the drive bay, the airflow exhausts before it effectively lowers the drive bay temperature.

Figure 72 shows the airflow pattern with a side panel vent limited to an area near the ODD. This exhaust location improves airflow through the drive bays and reduces drive bay ambient temperature.

Figure 71: BTX Tower Extensive Side Panel Vent Scheme Thermal Simulation Airflow Plot (Iso-View)
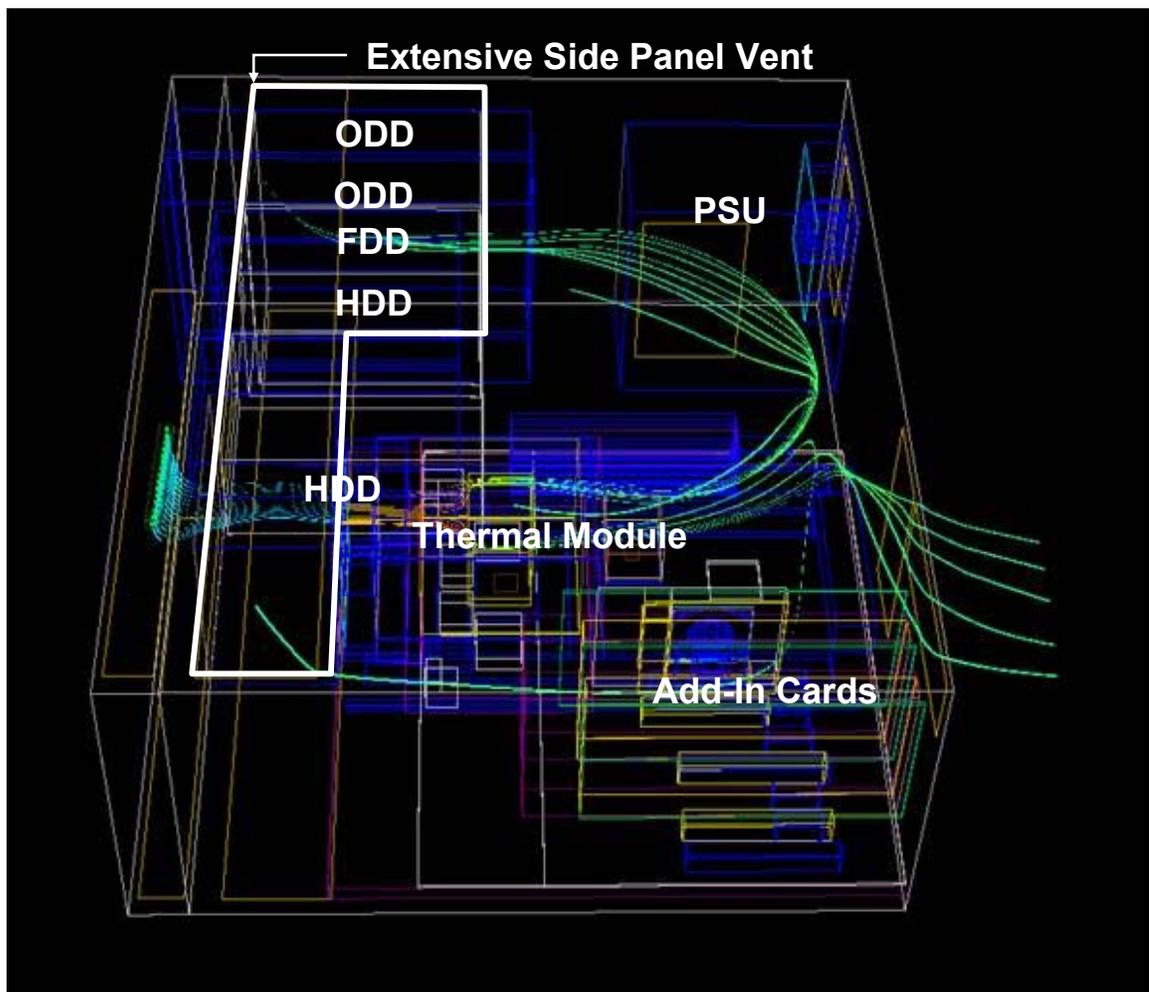
Figure 72: BTX Tower ODD-Only Side Panel Vent Scheme Thermal Simulation Airflow Plot (Iso-View)
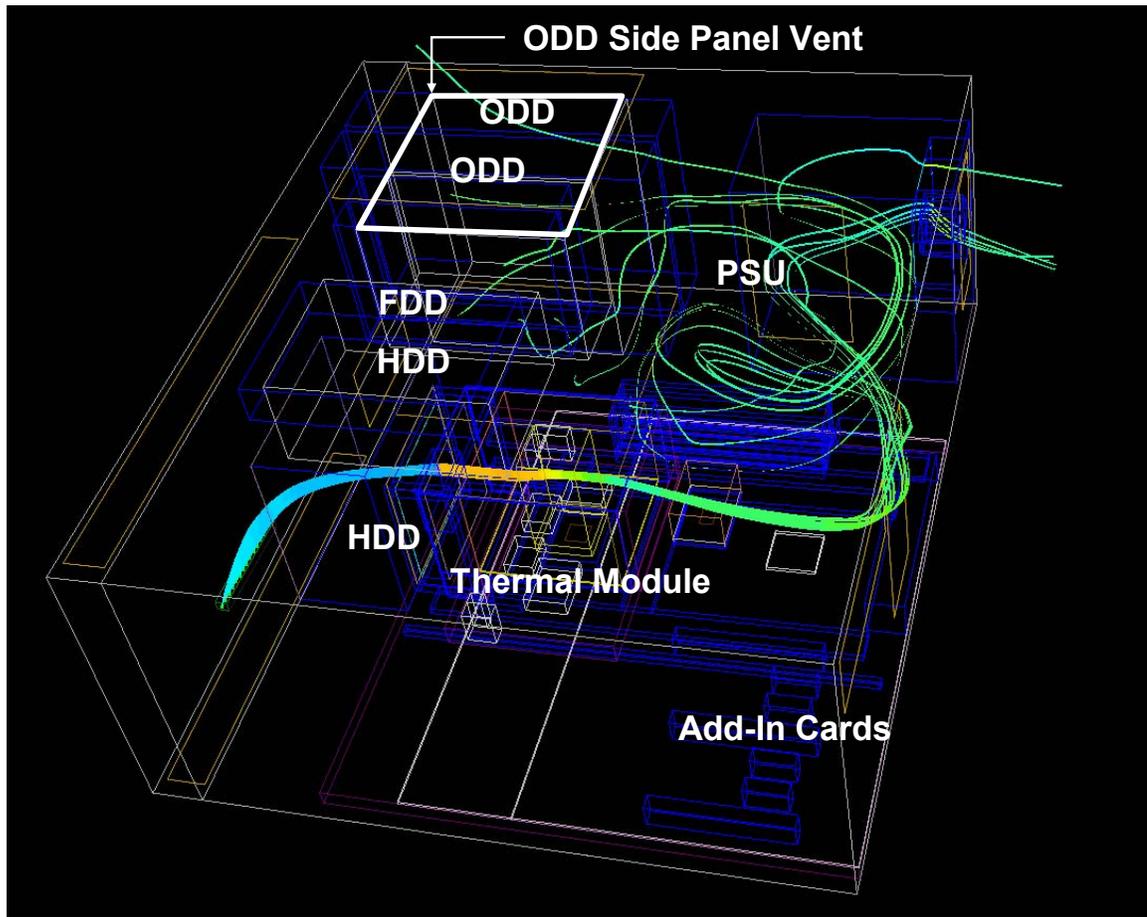
2.8.3 BTX Tower Thermal Simulation – Rear Panel Fan Investigation

A 90mm rear panel fan was included in the baseline BTX Tower thermal numerical model to determine its impact to system airflow and temperature. The rear panel fan was located near the PSU and RPIO.

In order to determine the relative impact, the system performance was evaluated for three rear panel fan fixed volumetric flow conditions: equivalent to the Intel Type I Reference Design Thermal Module operating point in the BTX Tower, then above and below that operating point. The thermal simulation operating point and drive bay temperature results for these three different rear panel fan flow conditions are outlined in

Table 18. It can be concluded that a rear panel fan is unnecessary in a properly ventilated BTX Tower when using the Intel Type I Reference Design Thermal Module. In fact, the addition of a rear panel fan will not only increase the acoustic sound power but also it may increase drive bay temperatures in some operating conditions. This, of course, would also require more complex fan regulation schemes.

Table 18: BTX Tower Rear Panel Fan Investigation Results

Rear Panel Fan Flow Rate	N/A	20 CFM	50 CFM	80 CFM
Thermal Module Fan Operating Point	59 CFM	61.5 CFM	61.7 CFM	61.9 CFM
ODD Ambient (°C)	44	49	41	37
HDD Ambient (°C)	47	51	43	45

2.9 Case Study – BTX Cube PC Configuration

A BTX cube system was designed and numerically evaluated to determine if a small form factor, high performance cube-shaped system could be adequately cooled without a rear panel fan and using the Intel Type I Reference Design Thermal Module Assembly with a 90mm fan.

The system mechanical design was based on tear-down analysis of existing cube-shape systems. A thermal numerical model was constructed to analyze the airflow and subsystem temperatures.

2.9.1 BTX Cube Component Placement and Ventilation Scheme

An analysis of existing cube-shaped system designs lead to the use of a two-slot BTX motherboard, one slot greater than the standard definition of a picoBTX board. This board size was selected based on the anticipated prevalence of both a PCI and a PCI-Express slot. The system included full height add-in cards, a CFX12V PSU, standard drive bay components, and a Type I Thermal Module.

The resultant cube system volume is approximately 14 liters (to the extended face of the ODD), or 15 liters measured to the front bezel. Using a picoBTX board and Type II Thermal Module, the system volume could be reduced to approximately 12 liters; however, this system option was not analyzed.

Figure 73 illustrates the component placement and ventilation scheme for the 15 liter cube system. The standard ODD is placed below the two side-by-side standard HDD placements, and above the Type I Thermal Module. A CFX12V PSU is located above the Rear Panel I/O and the board Zone B volumetric keep-in identified in the BTX Interface Specification.

The thermal numerical model included a 2-slot BTX board with the component power descriptions from Table 7. The Thermal Module Interface at the front sheet metal panel complies with the BTX Interface Specification requirements at the appropriate location.

The 35mm deep front bezel is ventilated on the bottom and each side to allow sufficient airflow through Thermal Module Interface and into Thermal Module. Again, this illustrates that an bezel without front ventilation can be used in BTX system designs and still meet the thermal and acoustic performance targets. In this particular design, airflow enters the vents on the side panel near the front of the system. That airflow is then channeled by a sheet metal wall into the front bezel cavity where it is then drawn into ventilated Thermal Module Interface on the front sheet metal panel.

Figure 73: BTX Cube Component Placement and Ventilation Scheme
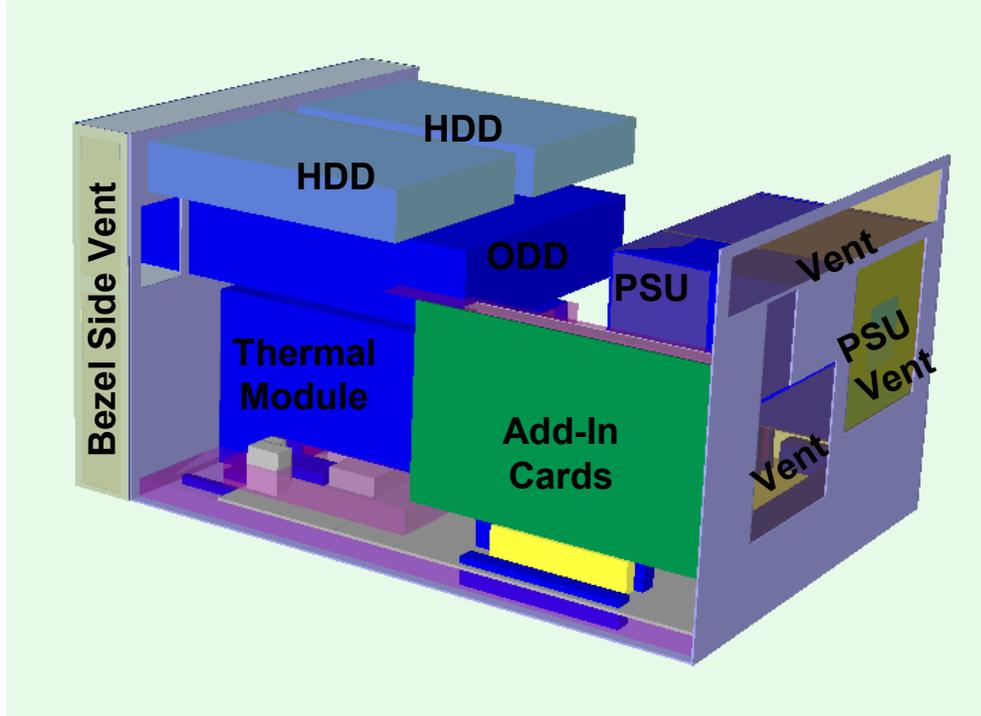
Figure 74 shows the airflow through the system on a plane 25mm above the motherboard. The airflow can be seen entering the bezel intake vents and being drawn into the TMA. The airflow through the TMA is high velocity and exits across the motherboard behind the TMA. A portion of the TMA exhaust airflow is deflected off of the rear chassis wall and forced to recirculate through the chassis to cool the memory components. However, the majority of the TMA exhaust airflow exits the chassis directly out the rear chassis wall after passing through the graphics card heatsink.

Figure 74: Cube Case Study – Airflow and Thermal Contour Plot
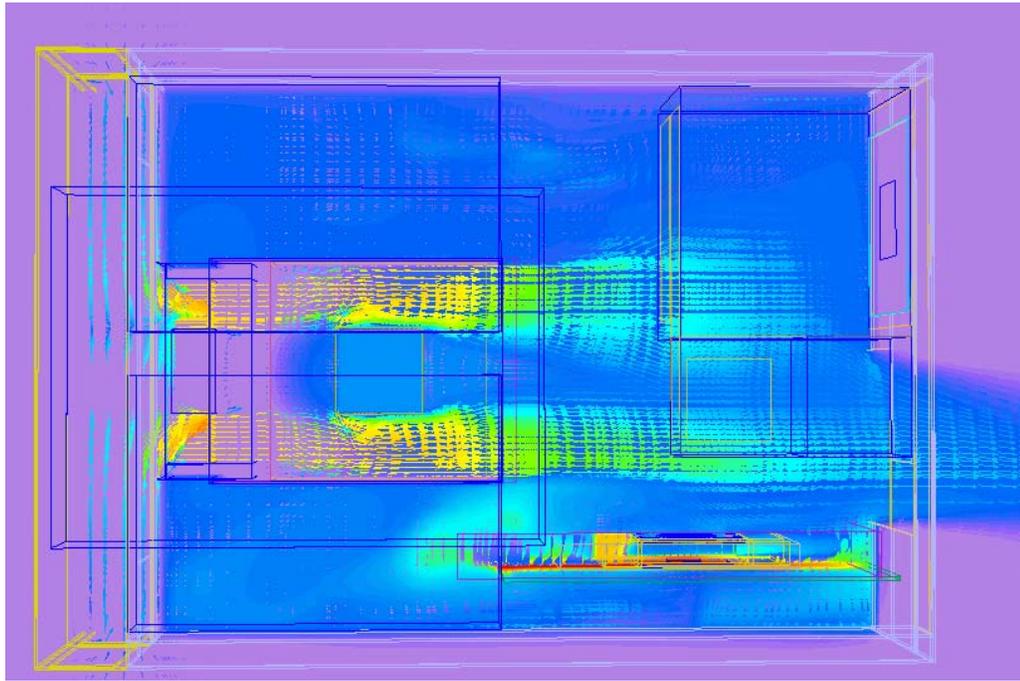
Care must be taken when selecting the graphics card. There are cards with active heatsinks that exhaust the air in all directions . In this Cube Case Study illustration, this omni-directional exhaust airflow would likely increase the ambient temperature near the ODD.

Figure 75 (side view of the system) and Figure 76 (top view of the system) illustrate the airflow behavior through the front bezel, TMA, and graphics card heatsink. The unique bezel and chassis side wall inlet to the Thermal Module Inlet are clearly seen with the particle trace coming in from the left side. The TMA exhaust airflow through the graphics card and out the rear panel is visible with the particle trace existing slightly toward the   right.

Figure 75: Cube Case Study - System Front View with Air Particle Trace On (Inlet, TMA, Graphics)
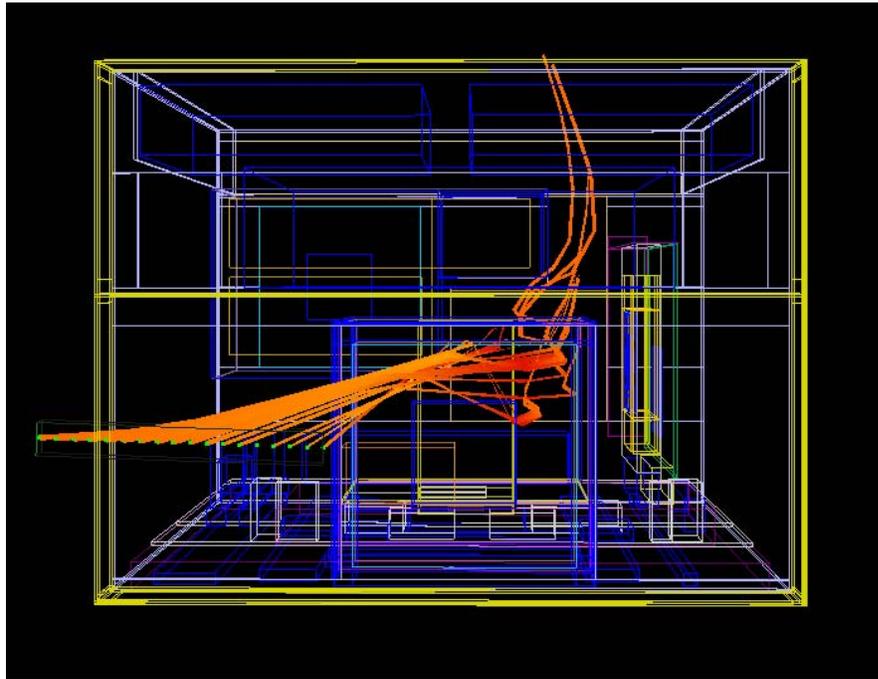
Figure 76: Cube Case Study - System Top View with Air Particle Trace On (Inlet, TMA, Graphics)
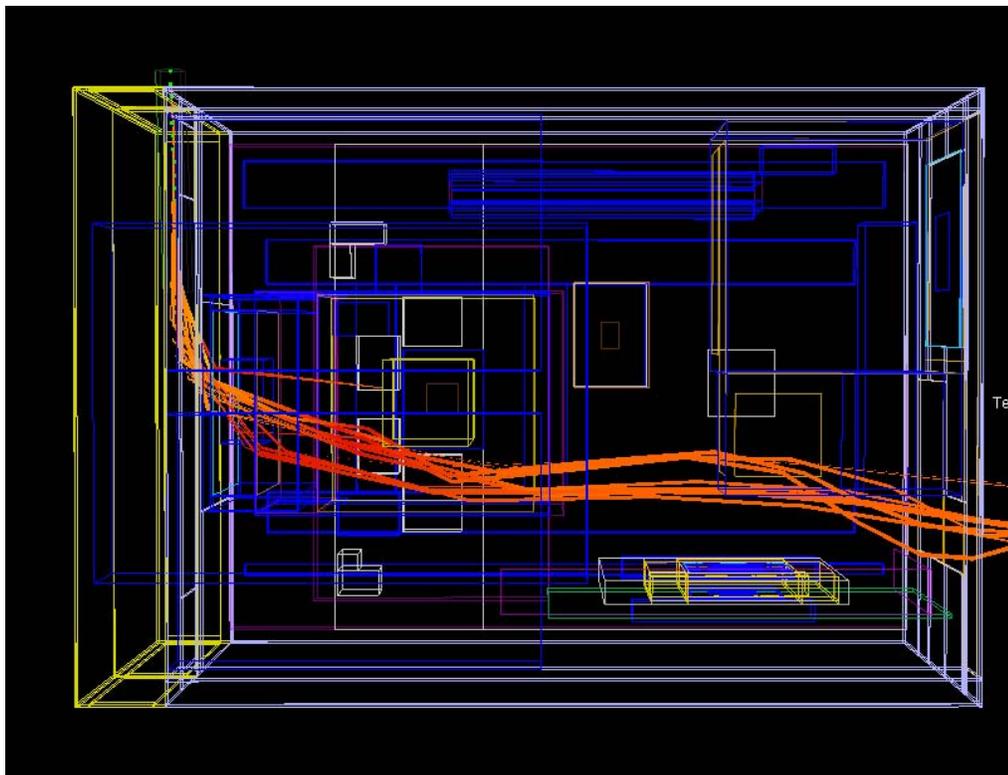
In Figure 77 (front view of the system) and Figure 78 (side view of the system), the particle trace illustrates airflow near the memory modules, which are to the left of the TMA. Memory airflow comes from the front of the system by way of processor voltage regulation exhaust, from the rear of the system as the TMA exhaust turns at the rear panel, and from the top of the system after it recirulates from the right side of the TMA. This last airflow pattern also offers cooling to the drive bays, which are located above the TMA.

Figure 77: Cube Case Study - System Front View with Air Particle Trace On (Memory, Drive Bays)
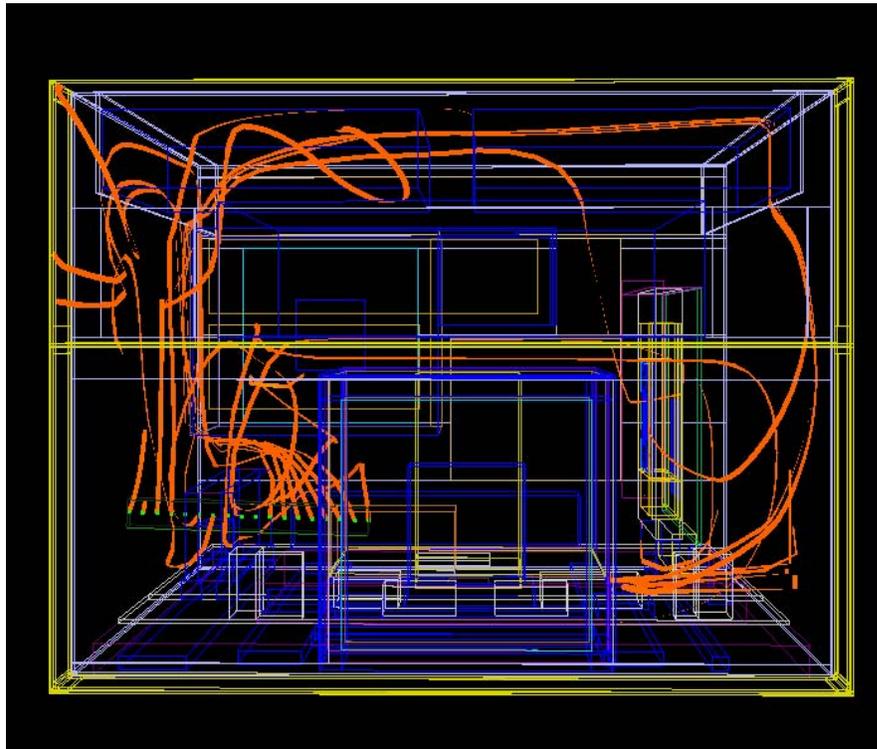
Figure 78: Cube Case Study - System Top View with Air Particle Trace On (Memory, Drive Bays)
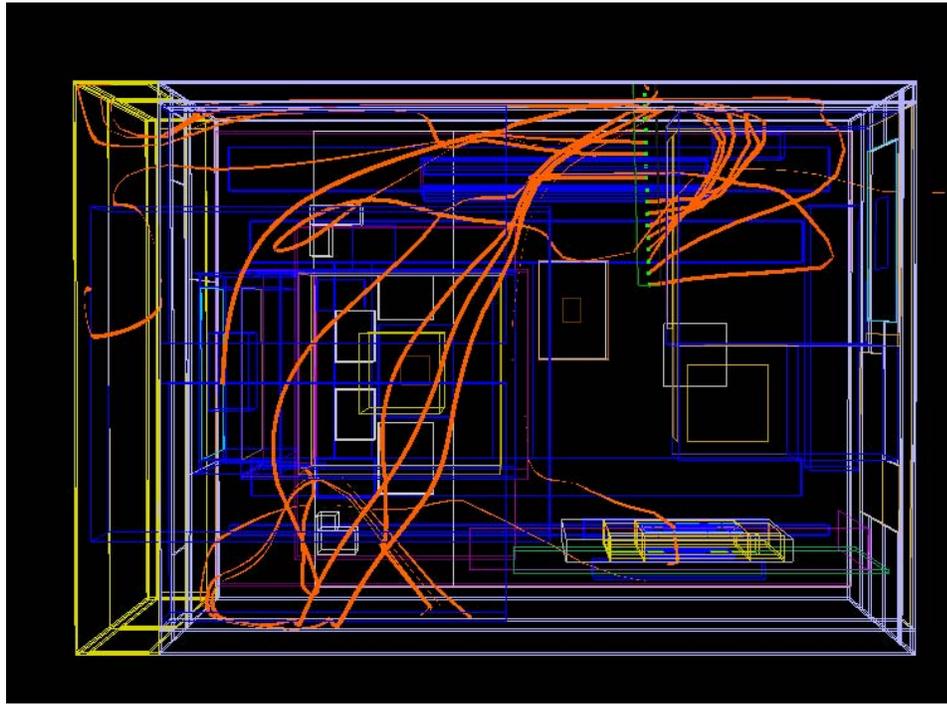
2.9.1.1 Summary

This cube case study demonstrates the scalability of the BTX into an system profile that is currently an expensive, custom system domain. It clearly illustrates that unique system profiles, ones other than "standard" desktop and tower systems, can be designed and built that in a way that allows the use of standard BTX components and complies with the *BTX Interface Specification*. This case study also illustrates that inlet airflow to the Thermal Module Interface can be provided in a way that does not compromise the appearance of a unique front bezel, since this case study used only chassis side panel and front bezel side and bottom inlet ventilation. The thermal numerical models used in this case study predicted that all component temperatures were compliant with their specifications at fan speeds that were equivalent to those used in the 12.9L Reference Design System Case Study.

2.10 Thermal Module Engineering

2.10.1 Definition of Airflow and Venting Terms

Table 19: System Airflow and Venting Terms

Term	Symbol	Definition
Fan		A motor driven impeller that generates airflow
CFM	CFM	Measure of volumetric airflow, cubic feet per minute
Flow Partitioning		Feature/geometry that introduces impedance in one of two parallel flow paths to control the distribution of flow between

Term	Symbol	Definition
		these two paths
Free Area Ratio	FAR	Ratio of open/total area within the perimeter of a vent
Heatsink		A component designed to provide efficient conductive heat transfer from a heat source
Airflow Impedance		Resistance to airflow, in terms of static pressure as a function of air speed
Stator		Stationary airfoil used to improve axial flow and pressure capability performance of a fan

2.10.2 Principles of the Thermal Module

The Thermal Module is a key element of the BTX form factor in that it directly affects the system and its component thermal solution performance. The Thermal Module is a unitized subassembly that acts as a system fan and flow distributor while providing immediate and direct cooling to the processor heatsink and processor voltage regulation. In typical implementations, air is drawn in through the front of the system to minimize the temperature rise into the Thermal Module and is partitioned into the heatsink and to the motherboard (above and below board) flow paths. The airflow is then distributed to the rest of the system components providing a relatively low temperature, high-velocity airflow. From the system perspective, the Thermal Module is effectively a system fan providing airflow at two fundamental locations, at the processor heatsink exhaust and above and below the motherboard at the front of the system.

The fundamental strategy for the Thermal Module is to maximize and effectively use the airflow from the Thermal Module fan for the CPU heatsink, VR, and system. This is accomplished by minimizing preheated air into the Thermal Module, maximizing fan performance that creates the low temperature, high-velocity airflow. This direct airflow path that is parallel to the board creates less turning , reduces impedance, and provides advantageous thermal boundary conditions for the system components..

The Thermal Module is intended to be a subassembly that consists of the Thermal Module fan, the CPU heatsink, and the airflow management components. The individual Thermal Module components will be designed and integrated by the thermal solution supplier for a simple integration of that assembly into the chassis.

2.10.3 Thermal Module Design

2.10.3.1 Airflow Split Above and Below Board

Within the Thermal Module the airflow is split into three paths. It is important to understand how the flow will be distributed down the paths so that proper heatsink and system performance can be achieved. The flow split is determined by the relative impedance between the voltage regulation flow path, both above and below the board, and the heatsink flow path. The impedance of the voltage regulation above the motherboard is a function of the voltage regulation component type and layout. The voltage regulation impedance below the motherboard is primarily a function of the gap to the chassis wall and the Support and Retention Module (SRM) geometry. Refer to Sections 2.4.3and 2.4.4 for more details.

The impedance of the CPU heatsink has a strong influence over the natural flow partitioning that result. Since the CPU heatsink impedance is a strong function of the airflow cross-sectional area available, the CPU heatsink impedance is quite different for Type I versus a Type II Thermal Module. This difference affects how airflow is distributed in systems that use either a Type I or a Type II thermal module. For example; the Type I Thermal Module in the Intel Reference Design system, the natural flow split is

approximately 70% through the heatsink and 30% through the voltage regulation above and below the board.

For the Type II Thermal Module, the area ratio between the heatsink and voltage regulation cross sectional areas is reduced because the Type II available heatsink height is 25.4 mm shorter than the Type I height. This reduced ratio will change the natural flow split to approximately 55% through the heatsink and 45% through the voltage regulation above and below the board, as evidenced in the Intel Reference Design system. It may be necessary to add an impedance element in the Thermal Module to increase the above and below motherboard impedance that also results in more air forced through the CPU heatsink.

For proper heatsink and system cooling, it is necessary to understand the flow split for the Thermal Module design. The Thermal Module flow split will determine what type of fan is necessary to meet the system and heatsink volumetric airflow requirements. The flow split will determine what steps, if any, need to be taken to augment airflow to the heatsink path or to the voltage regulation airflow path.

2.10.3.2 Temperature Rise for Heatsink and Voltage Regulation

The temperature rise into the inlet of the fan for the Thermal Module is expected to be approximately 0.5 °C relative to the external ambient temperature. The temperature rise is due to reintroduction of heated air from the following three sources:

- Recirculation of exhaust airflow out of vents in the front of the system
- Slight leaks between the Thermal Module seal and the chassis
- Leakage between the unsealed bottom of the Thermal Module fan and the chassis

The Thermal Module should be designed to minimize reintroduction of heated airflow within the system into the thermal module. The reference designs utilize rubber gaskets to accomplish this though other design approaches may also be suitable. The minimal temperature rise provides a good environment for the CPU heatsink and voltage regulation as well as the remainder of the system.

The temperature rise exiting the Thermal Module heatsink is important for system cooling, in particular the cooling of the MCH, ICH, and graphics components. When the Thermal Module and PSU fan speeds are maximized and approximately 115 W power is applied to the processor, the mean temperature rise through the CPU heatsink and voltage regulation will be on the order of 5 °C to 7 °C. This will result in an average inlet temperature to system components such as graphics, memory, and MCH of approximately 10 °C lower than a typical ATX chassis.

The CPU Voltage Regulation directly benefits from the low ambient air temperature provided by the Thermal Module fan. The Voltage Regulation air temperature is approximately 15 °C to 20 °C lower than typical ATX chassis. This allows the number of voltage regulation components and/or phases to be reduced, lowering motherboard cost. The temperature rise for a BTX flow path such as the voltage regulation or heatsink can be approximated with the following formula.

Equation 38: Temperature Rise Approximation

$$\Delta T(^oC) = \frac{Q(Watts)}{\rho(kg/m^3) \cdot \forall(m^3/s) \cdot Cp(J/kg^oC)}$$

Where: Q is the input power

ρ is the density of the fluid (air)

∀ is the volumetric airflow speed

Cp is the specific heat of the fluid (air)

Which can be approximated by the following equation when considering the properties of air at 35 °C.

Equation 39: Temperature Rise Approximation (Simplification)

$$\Delta T(^oC) \approx 0.263 \cdot \frac{Q(Watts)}{\forall (m^3/s)}$$

2.10.4 Type I Thermal Module Size Relations to Airflow Design

The Type I Thermal Module is designed to accept a standard 92 mm box fan. The depth of the fan can be up to 40 mm. The distance between the front of the motherboard and the Thermal Module Interface is defined in the *BTX Interface Specification* and this distance constrains the fan depth. High performance fan solutions that optimize flow through the heatsink and voltage regulation will be beneficial for thermal and acoustic performance – see Chapter 1 for more details.

2.10.5 Thermal Module Element Details

2.10.5.1 Fan Design Strategy

The system fan used in the Thermal Module should be a high performance fan so that this single fan can provide the majority of the airflow needed for the system while overcoming the impedance of the heatsink and system. This high performance fan will usually need a fan speed control that uses pulse width modulation so that a large range from high to low RPM can be achieved. The large RPM range will allow the fan to be acoustically acceptable at typical use conditions while meeting the thermal requirements in the extreme power and ambient temperature conditions. For more information on fan speed control and acoustics refer to Chapter 1.

The fan is composed of the housing, impeller and motor which all reside within the Thermal Module duct. The design strategies for these components specifically for the Thermal Module are discussed in the following sections.

2.10.5.2 Housing Design Strategy

A key consideration of the fan design for the Thermal Module is to fully utilize all velocity components exiting the fan. Specifically, it is desirable to efficiently turn the swirling flow from the impeller and straighten it allowing it to flow through the vertically or horizontally parallel fins of the CPU heatsink. This can be accomplished using a fixed blade stator that is attached to or integral to the fan housing downstream of the impeller. The stators are designed so that the flow vector coming off of the impeller blade is turned into axial flow to better align with the CPU heatsink fins. By straightening the flow at the stator exit, the turning losses at the leading edge of the heatsink fins are greatly reduced resulting in a reduction of overall pressure loss in the Thermal Module. Detailed stator design is dependent on the impeller design characteristics such as number of blades, blade angle and airfoil type. Detailed stator design, if employed, will typically be done by the Thermal Module supplier.

Figure 79: Cross-Sectional View of an Impeller and Stator Blade

Impeller Blade
(moving)

motion

Stator
(stationary)

OM16749
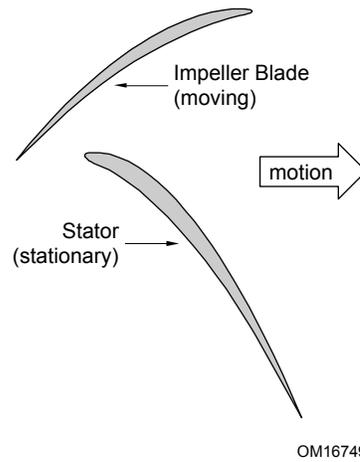
The velocity vector entering the stator is related to the axial direction by angle ϕ_2 and exits the stator at angle ϕ_2 as illustrated in Figure 80.

Figure 80: Description of Airflow Vectors Entering and Exhausting From the Stator Blade
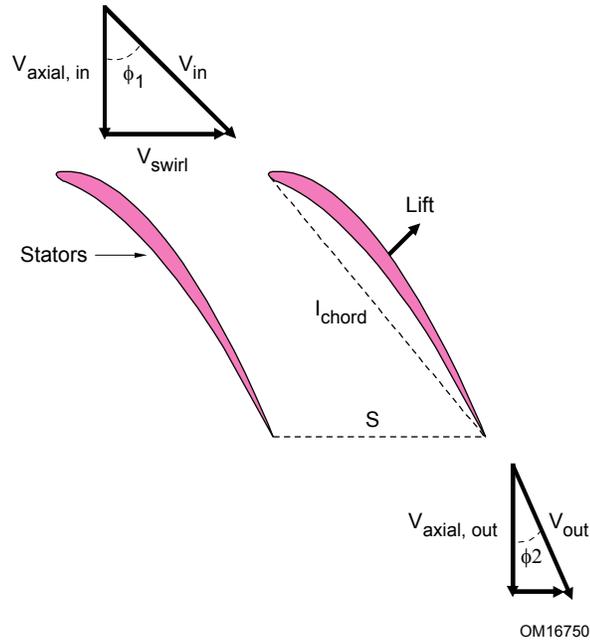
The design intent of the stator is to turn the airflow such that angle ϕ2 is minimized. The mechanism by which this is accomplished can be simply described by looking at the momentum balance on the stator blade.

Equation 40: Stator Momentum Balance

$$\rho \cdot s \cdot v_{axial}^2 \cdot \left[\tan(\phi_1) - \tan(\phi_2)\right] = Lift \cdot \cos(\omega)$$

Whereas ρ is the density of air, s is the pitch between stators, ω is the angle between the chord of the stator blade and the axial flow direction, and the Lift force is determined as:

Equation 41: Lift Force

$$Lift = C_L \cdot \tfrac{\rho}{2} \cdot v_{in}^2 \cdot Area_{chord}$$

Whereas C_L is the lift coefficient of the stator blade and Area$_{chord}$ is the area described by the chord length and blade depth (assumes into page depth is 1m). Substituting Equation 41 into Equation 40 provides the following expression.

Equation 42: Tangential Expression

$$\tan(\phi_1) - \tan(\phi_2) = \left[\tfrac{1}{2} \cdot C_L \cdot \cos(\omega) \cdot \frac{l_{chord}}{s}\right] / \cos(\phi_1)^2$$

Since the design intent is to minimize ϕ2, it is desirable to maximize the difference between the tangents of angles ϕ1 and ϕ2. There are two strategies to maximize the right side of Equation 42. The maximum lift coefficient for airfoil geometries is known to occur for angles of attack in the range of 10°-20°, whereas a

stall condition typically occurs beyond an angle of attack of 20°. Therefore, if the inlet angle, $\phi1$, is known, the first approach would be to bound the stator entrance angle, $\phi2$.

Equation 43: Angle of Attack

$$\omega = \phi_1 - (15° \; to \; 20°)$$

The second approach maximizes the chord to pitch ratio by increasing the number of stator blades or increasing the stator depth. However, it is necessary to balance this approach against impedance and acoustic detriments. See Chapter 1 for more details.

Another advantage of a good stator design is the improvement in pressure capability of the fan. In process of turning the flow, the stator converts swirl energy into pressure head capability, which can be illustrated by Bernoulli's equation.

Equation 44: Bernoulli's Equation

$$p_{in} + \frac{\rho}{2} \cdot v_{in}^2 = p_{out} + \frac{\rho}{2} \cdot v_{out}^2$$

Whereas Pin and Pout are the static pressures at the inlet and outlet of the stator blade, respectively. For an ideal stator, the flow is turned completely and V_{out} is equal to V_{axial}. Substituting back into Equation 44 gives:

Equation 45: Inlet Velocity

$$v_{in}^2 = v_{axial}^2 + v_{swirl}^2$$

and

Equation 46: Inlet Pressure

$$p_{in} - p_{out} = \frac{\rho}{2} \cdot v_{swirl}^2$$

Improvements in pressure capability will vary with fan speed and the impedance in which it operates, so it is important that the stator is designed for the operating condition in which the improvements are most needed. This is most often at the fan's maximum operating speed and at the impedance where the fan's performance curve and the system impedance curve intersect.

For the Type I Thermal Module, the inlet is of similar size to the chassis inlet vent and contraction or expansion losses can generally be ignored. For the exhaust side of the fan housing near the entrance to the ducted CPU heatsink, the upper housing corners will create expansion that will cause some turning recirculation and pressure loss. This loss can be reduced by providing a smooth transition at the fan-to-heatsink duct housing upper corners - between the fan impeller exit and entrance of the CPU heatsink.

For the Type II Thermal Module, the inlet and exhaust side of the fan is 70 mm by 90 mm. The management of airflow into and exhausting from a rectangular to circular fan area and back to a rectangular area needs to be carefully considered in the design of the Thermal Module. The fan housing should be design to draw air from the full Thermal Module Interface area available to minimize pressure losses. Typically the Thermal Module duct will have to be designed to manage the airflow into and out of a standard 70 mm box fan. Custom fan housings may be a better solution as they can provide a smooth

transition to the fan impeller on both the inlet and exhaust side of the fan and, where applicable, stator blades.

2.10.5.3 Impeller

The impeller for the Thermal Module should be designed to produce as much axial flow and as little swirl flow as possible, since the air will be impinging directly on the heatsink fins which are directly down stream of the fan exhaust (or near-field.) If the velocity of the air from the impeller has a large entrance angle relative to the heatsink fins, the airflow will be reduced due to the pressure loss associated with the impinging and turning of airflow required for the airflow to travel through the heatsink fins.

The blade design should consider the effects of the near-field inlet vent and heatsink fins, as well as the impedance of the heatsink and typical system in which the Thermal Module will be installed.

2.10.5.4 Motor

High performance fans require motors that are capable of delivering relatively high torque, which efficiently turn the DC power into fluid kinetic energy. It is also acoustically beneficial to design or select a motor with a wide range of available speeds for the maximum flow is needed under different circumstances than acoustic performance (low fan speed) is desired.

High torque motors typically require a large hub diameter, as the torque of a motor is proportional to the square of the diameter. Torque is linearly proportional to the length along the axis of the motor. Magnet type and wire gap also have a significant affect on torque capability of motors. Also, if the efficiency of the motor is too low, the motherboard may not be able to supply sufficient current, causing the motor to overheat. This will also increase the inlet air temperature to the CPU heatsink due to the air heating from the extra motor power. Motherboard fan headers (at 12V) are typically limited to 1 A current draw in steady state operation.

To maximize the performance capability of a fan and minimize its Sound Power contribution in important operating conditions, a wide operating speed range is needed. Motors that vary speed by varying supply voltage are typically limited to a minimum of 5 V (12 V. maximum). The voltage range limits the minimum operating speed of the fan to 42% of its maximum operating speed. For instance, for a fan that has a 4500 RPM maximum operating speed, the minimum RPM would be 1900 RPM. Alternatively, digital speed control such as pulse width modulation (PWM) can be integrated into the motor controller. PWM can allow a minimum operating speed as low as 18% of maximum operating speed. For instance, for the maximum operating speed of 4500 RPM, the minimum RPM would be 800 RPM. It is recommended that PWM implementations use a four-wire motherboard fan header to provide an isolated PWM control signal.

2.10.6 Airflow Management Design Strategy

2.10.6.1 Duct

The duct is a contributing component to the heatsink thermal performance. Airflow bypass around a heatsink occurs when the air generated by the Thermal Module chooses the easier path around the dense fins of the heatsink. This reduces airflow across the fins and lowers the heat transfer coefficient. In order to minimize the bypass around the processor heatsink and force the air through the fins, the duct is used to direct the air through the heatsink. The duct also provides the structural load path from the heatsink to the chassis structure to maintain the pressure for thermal interface as well as carry the initial preload and heatsink inertial loads resulting from system shock loading. Information on the structural role played by the duct is available in Chapter 4.

In the absence of bypass reduction, the high-density fin configuration on high performance heatsinks create significant impedance that will cause the airflow to move around the outside of the fins; thereby reducing the performance of the heatsink.

2.10.6.2 Thermal Module Interface Seal

The seal is a compressible gasket that seals the inlet of the Thermal Module and the chassis-provided Thermal Module Interface described in the *BTX Interface Specification*. The seal is a design detail implemented to minimize the amount of air that would otherwise recirculate from the chassis interior back into the Thermal Module fan. The seal also minimizes noise that could be generated when air passes through this small potential airflow path. The seal design needs to be compliant enough to take up the expected dimensional tolerance between the Thermal Module and Thermal Module Interface. It is recommended that the seal compliance or compression range be 2 mm to 4 mm.

2.10.6.3 Flow Partitioning Device

As described previously, partitioning of the airflow may need to be designed into the Thermal Module to ensure there is sufficient airflow through the CPU heatsink to meet the thermal performance requirements. The flow partitioning needs to consider the effects on total system flow since adding impedance to the above and below motherboard airflow path will increase the Thermal Module impedance and reduce the total airflow out of all the exhaust paths of the Thermal Module.

A Flow Partitioning Device (FPD) can be incorporated into the fan housing or the duct. Whether a separate piece or an integral part of the fan housing, the FPD needs to be integrated into the Thermal Module assembly since it needs be tuned to meet the CPU heatsink airflow.

2.11 Heatsink Design Strategy

2.11.1 Constraints

BTX Type I and Type II Thermal Modules have the same motherboard mechanical interface and motherboard keep-out footprint. The motherboard keep-in volume described in the Socket and Platform Design Guides describe the constraints for the design of the CPU heatsink base. The current LGA775 socket keep-in allows the CPU Voltage Regulation to be installed under the CPU heatsink and that volume requirement is the same whether a Type I or Type II Thermal Module is used.

The *BTX Interface Specification* provides the same width and depth constraint for both Type I and Type II Thermal Modules. The height of the Type II Thermal Module heatsink is 25.4 mm shorter than the Type I Thermal Module, as defined in the *BTX Interface Specification*.

As part of the heatsink design, the thickness of the duct and heatsink to duct assembly clearances need to be considered. Also, the relative deflections between the motherboard and the heatsink during preload application needs to be considered.

The mass limit for BTX heatsinks that use Intel reference design structural ingredients is 900 grams. The BTX structural reference component strategy and design is reviewed in depth in Chapter 4.

2.11.2 Thermal Performance

The overall thermal performance of the heatsink can be broken down into the following lower level performance characteristics:

- Case to sink resistance: Represents the Thermal Interface Material (TIM) conduction resistance. The term, when used in estimates of integrated performance of a Thermal Module or heatsink, typically includes the 3σ performance variation due to manufacturing, assembly, and material heat transfer behavior tolerances.

- Conduction resistance: Represents the conduction resistance between the heatsink to TIM interface and the cumulative resistances to the base of the fins (spreading and conduction). This term can also include a fin to base contact resistance.

- Convection resistance: Represents the fin efficiency and the convective resistance associated with transfer of heat from the fin to the air.

- Air heating resistance: Represents the reduced sensible heat capability of the air due to the increase in air temperature as heat is added to the air from the fan and fins.

2.11.3 Thermal Interface Material

The thermal interface material selection in an important consideration to the overall thermal performance as is can represent over 25% of the total resistance. Higher performance materials, such as thermal grease, provide a performance improvement over a typical phase change material. Also, thermal interface material performance improves with increased mechanical pressure between the heatsink and IHS as this reduces the bond line thickness and shortens the TIM conduction path.

The application of preload on the CPU through the attachment of the CPU heatsink is described in detail in Chapter 4. The pressure on the TIM that this preload generates typically improves the TIM thermal performance.

2.11.4 Conduction

Reducing the effective length that heat must travel from its source to the base of the fins and then to the air can minimize conduction resistance. The physical length divided by the conductivity of the material can characterize the effective length. The conduction value is difficult to approximate with first order equations, as it is an iterative problem that is also dependent on the convective heat transfer coefficient distribution.

For the Type I Thermal Module, the available CPU heatsink volume is relatively high which increases the importance of creating an efficient conduction path to the top of the heatsink. A vertical core (spreader) can be used to improve conduction for taller heatsink designs – they can be more effective than a flat base (spreader) conducting laterally, even though a vertical core heatsink may require shorter fins than a flat base heatsink.

For both Thermal Module types, all of the many heatsink manufacturing methods available for core, fin, and their assembly can be used to design heatsink solutions the meet the overall performance requirements.

2.11.5 Convection

Convection resistance is primarily a function of mass flow of air through the heatsink and the total fin area. Other factors such as, fan airflow velocity distribution, boundary layer development, separation, and fin surface characteristics may also influence the convective resistance in a second order effect.

Engineering low convective resistance typically involves an optimization of the fin pitch, fin thickness, and fan performance to achieve the highest performance. Increasing the fin pitch (number of fins) increases the fin area but increases airflow impedance, which reduces the mass flow rate. Increasing fin thickness improves the fin efficiency but also increases heatsink mass and reduces the mass flow rate. These design optimizations are typically done using computation fluid dynamic (CFD) numerical tools. The convection and conduction optimization should be done simultaneously or iteratively to determine the best split of the available heatsink mass between the heatsink base (conduction) and the fins (convection).

For BTX Thermal Modules, the airflow velocity is non-uniform since the fan blades span from the top of the heatsink to below the motherboard and the fan hub does not generate airflow. The highest airflow velocity from the fan is typically in the last 70 to 90% of the blade span. This results in a mass flow that is distributed in an annular shape around the hub of the fan. The velocity directly downstream of the fan hub is typically relatively low. This airflow distribution can be considered in the CFD analysis and heatsink optimization for convection as well as conduction performance.

Figure 81: Simplified Fan Blade Velocity Distribution at Which the Maximum Velocity Occurs at 85% of the Blade Span
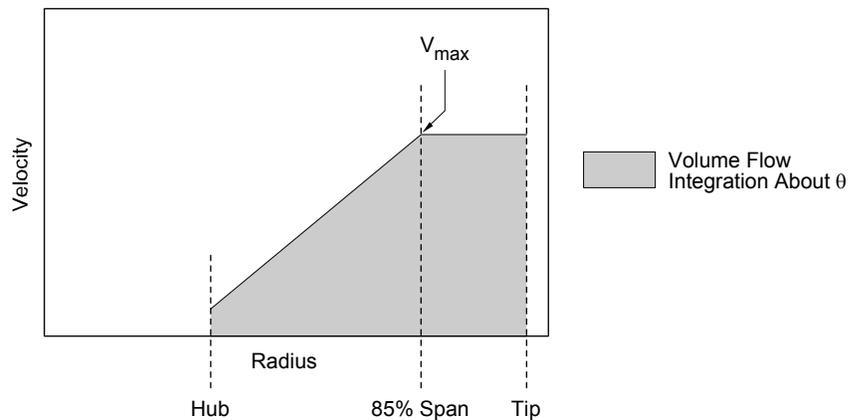
2.12 Characterizing the Thermal Module

2.12.1 Impact of System Impedance

In addition to the many design factors within the Thermal Module, the thermal performance of the heatsink is also dependent on the impedance of the system into which it is installed. A system that has high airflow impedance will reduce the airflow through the Thermal Module and that will reduce the CPU heatsink thermal performance.

It is suggested that the Thermal Module performance be met when installed into a system that has known impedance, excluding Thermal Module impedance, described below. The impedance capability should be specified as part of the Thermal Module performance estimate.

Figure 82: System Impedance Curve Excluding the Contribution from the Thermal Module for Which the Heatsink Performance Targets Are Achieved
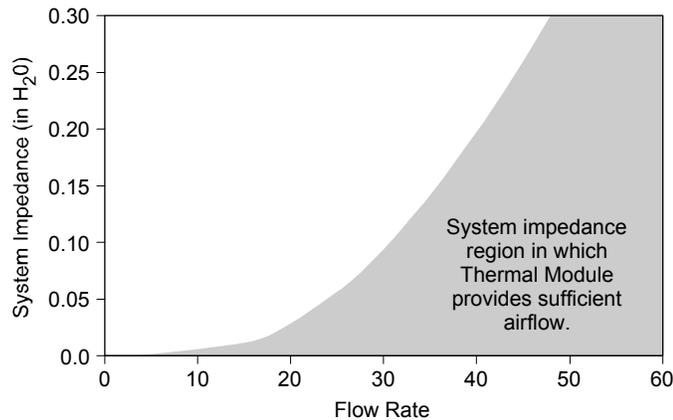
2.13 Thermal Module Case Studies

2.13.1 Type I Thermal Module Design (for 2004 Performance Targets)

A Type I Thermal Module reference design was developed utilizing these guidelines and is provided as an example demonstrating the design considerations and engineering approach. The Thermal Module fan is comprised of a 90-mm impeller, high performance motor, and an integrated stator. The heatsink is a "post and plate" design utilizing a central copper core with embedded, helically-wound, aluminum fins. The duct interfaces to the chassis with a rubber seal and provides the structural retention to the heatsink. Details of each component are described in the following sections.

The Thermal Module case study that follows, showcases the design of the Intel Reference Design Type I Thermal Module to meet the Platform Compatibility Guide '04 B targets. Section 2.13.1 describes the design of the Intel Reference Design Type II Thermal Module for 2004 . After reading these sections, refer to Section 2.13.3 for engineering suggestions for Platform Compatibility Guide '04 A Type I and Type II Thermal Module development.

Figure 83: Cross-Sectional View of a Type I Thermal Module Reference Design Assembled to a MicroBTX Motherboard
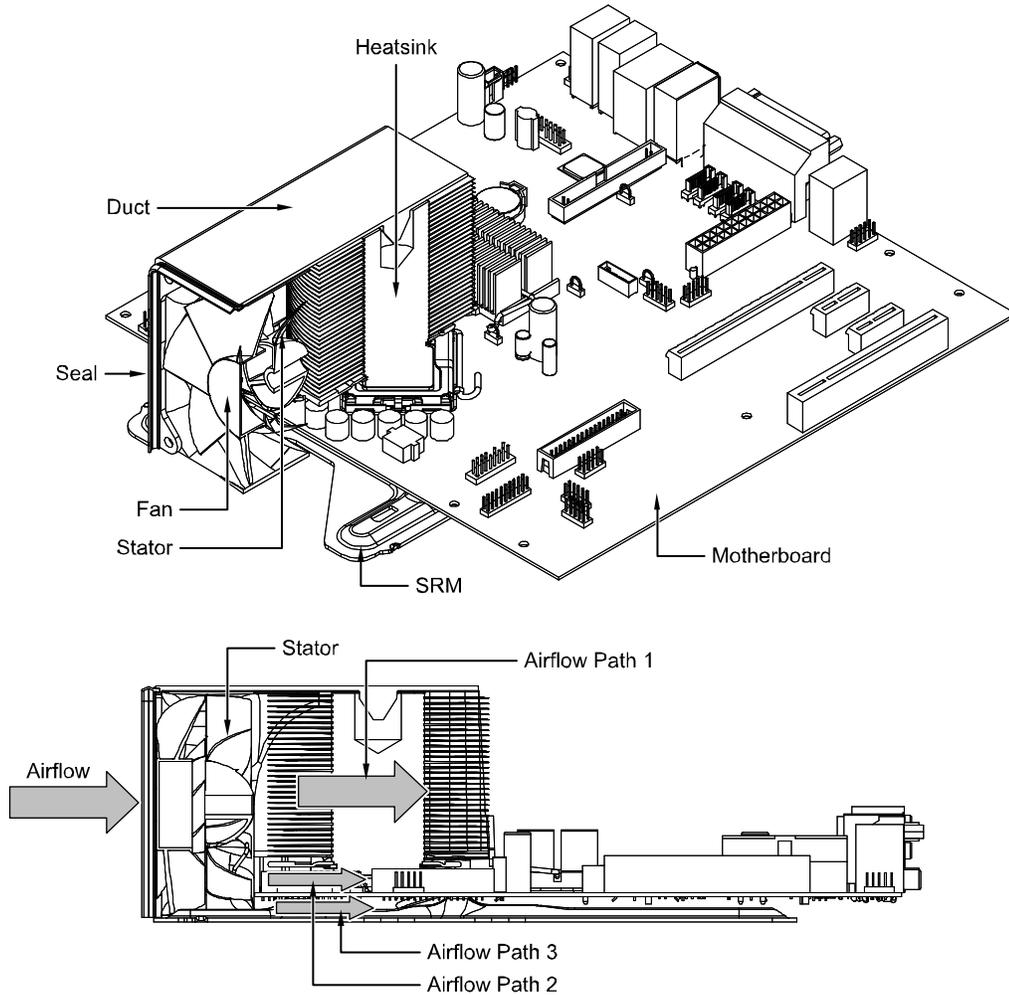
OM17076

2.13.1.1 Fan Assembly

The design intent of the reference design fan assembly was to provide sufficient volume of airflow to the processor as well as the system while maximizing the axial component of the airflow from the fan. In order to provide sufficient airflow, considerable attention was given to the shape of the impeller blades as well as the motor design. In addition, the fan design for the Type I Thermal Module was intended to interface to the duct utilizing the industry standard mounting geometry for a 92-mm fan.

An airfoil shape was selected for its flow performance and evaluated at various blade angles in terms of maximizing the blade lift coefficient at minimal sound power levels. Also, the blades are attached to the hub at a sweep angle of 20°. The sweep angle is formed by the radius of the fan and the leading edge of the impeller blade. While it would be desirable to create the fan with no sweep angle to provide best performance, the manufacturing process required an angle in order to remove the impeller from the mold. Therefore, the impeller design was based on a trade-off assessment of performance and manufacturability.

Figure 84: Impeller Design

Blade Angle
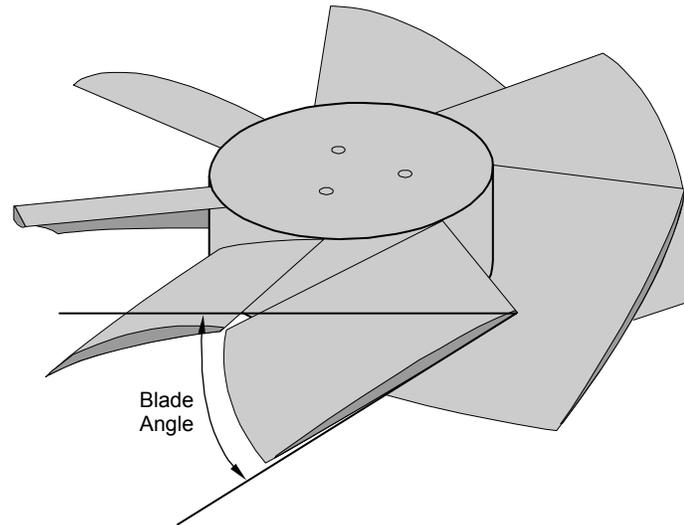
OM17077

The impeller diameter was selected to be 90 mm and three impeller designs were fabricated and tested in a wind tunnel and acoustic chamber in order to assess the flow performance versus acoustics. Table 20 presents flow rate sensitivity results for three impeller designs having the same airfoil shape. In this example, the blade angle was adjusted and the flow rate results are presented at the same sound power level.

Table 20: ISO-Acoustic Flow Comparison of an Impeller with Three Blade Angles

Blade Angle	Airflow
30^o	36.6 CFM
35^o	35.2 CFM
40^o	32.8 CFM

Requirements for the fan motor were based on estimates of the total impedance of the system, including the Thermal Module. The fan requirements were then translated into a torque requirement for a given RPM (see Table 21).

Table 21: Torque Settings for the Fan Motor

RPM	Pressure	Torque
4500	Free stream	11.2 Nmm
4150	0.15 in H_2O	11.8 Nmm
3800	0.25 in H_2O	12.7 Nmm

In addition, it was desirable that the fan operate at speeds less than 1000 RPM to take full advantage of any acoustic opportunity at low system power conditions. In order to meet the high torque requirements, the motor required a hub of 20 mm thickness and 38 mm diameter in order to house the windings. Instead of controlling the fan speed by voltage, which is typically limited to incremental steps between 5 and 12 volts, the fan is controlled by a Pulse Width Modulation signal operating at a fixed frequency typically between 21 kHz and 28 kHz. This control approach provides added benefit over sending a pulsed 12V signal to the fan in that it eliminates fan chatter. The PWM signal is transmitted to the fan by a fourth wire attached to

the new 4-pin fan header. In the case of the reference design fan, the PWM signal allows the fan to operate at speeds as low as 900 RPM. By allowing a wide range of RPM, 900 to 4500, the reference design Thermal Module is able to operate at low acoustic levels when the thermal demands of the processor and system are low. For more information regarding the reference design implementation of the PWM signal and fan settings, refer to Section 3.4.4. Additional details about the 4-pin fan header can also be found in the *Fan Specification for 4-wire PWM Controlled Fans.*

The final component of the fan design is the housing. For the reference design fan, the housing was intended to provide further improvements to the airflow by integrating a stator at the exiting side of the impeller. The principles of the stator operation are described in the previous sections.

The first step in establishing the stator design involved determination of the airflow exit angle off the impeller blades. This step was accomplished through CFD analysis.

Figure 85: CFD Analysis of Impeller Airflow
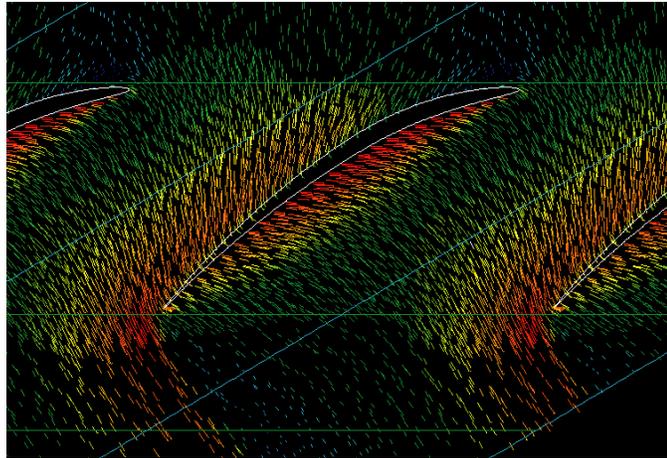
OM17110

It was found that the exit angle varies with radial distance from the fan hub and the downstream pressure. Analysis results provided the relationships for exit angles at the anticipated operating points.

Figure 86: Exit Angle as a Function of Radius and Pressure

OM17078
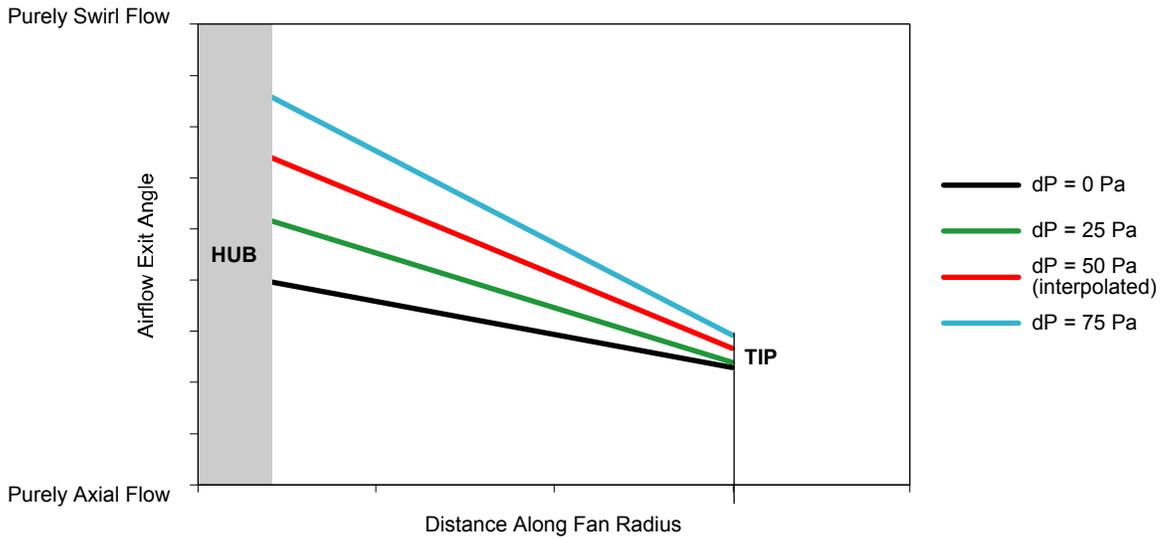
Utilizing these results, the stator blade angle and twist from hub to tip were determined in order to provide the most effective straightening of the flow. The stator blades were then integrated into the housing as shown in Figure 87.

Figure 87: Stator Blades Integrated into the Fan Housing
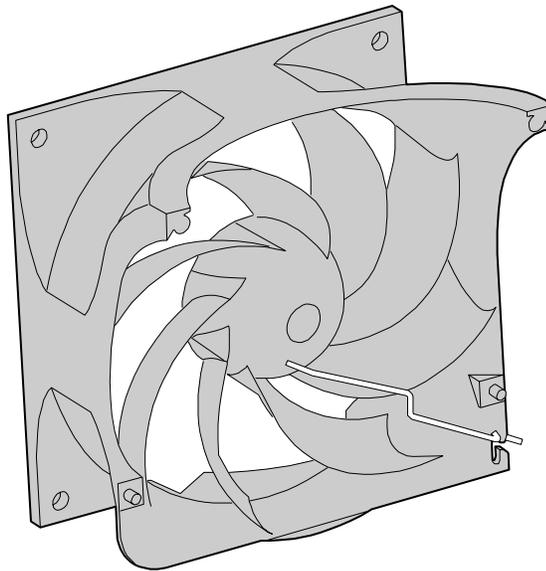
OM17079

Several stator-integrated housings were fabricated for testing in order to realize the benefit of the stator. Results of the testing indicated approximately 19% improvement in air delivery from the fan assembly (see Figure 88).

Figure 88: Fan Curves With and Without Stators
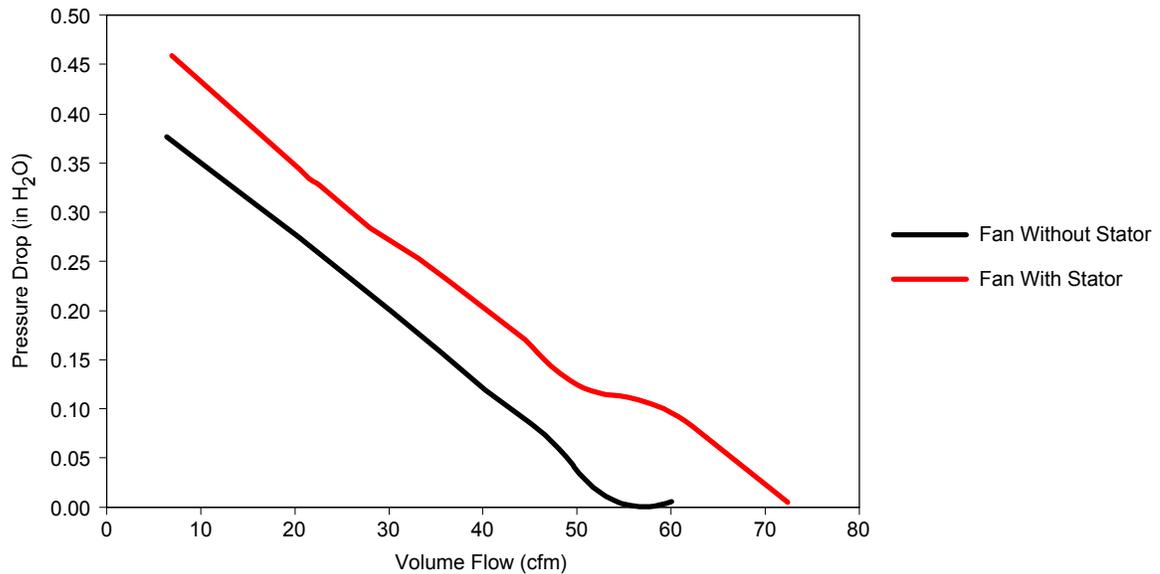
OM17080

Overall, the fan assembly was developed to provide large amounts of axial airflow to the Thermal Module. The impeller design was selected based on analysis of various airfoil types, a high torque motor was utilized with PWM control in order to provide a wide range of operating speeds (900 to 4500 RPM), and the housing design incorporated a stator and was designed for integration with the TM duct and interface to the chassis.

2.13.1.2 Heatsink Design

In order to demonstrate the cost-savings opportunity and structural capability of a BTX-compliant system reference design, the heatsink was designed to meet the thermal performance requirements utilizing relatively inexpensive materials and manufacturing technology, which in turn, resulted in a mass nearing the maximum mass requirement of 900 g.

Since the heatsink is directly downstream of the fan, it was necessary to consider the non-uniform airflow profile entering the heatsink. In particular, the 'no-flow' wake region directly downstream of the fan hub required consideration by the heatsink design. For this reason, it was determined to utilize a "post and plate" heatsink design that incorporated a post in the center of the heatsink that takes advantage of airflow at the sides of the heatsink as depicted in the CFD image in Figure 89.

Figure 89: Top View of Streamlines through a Post and Plate Heatsink
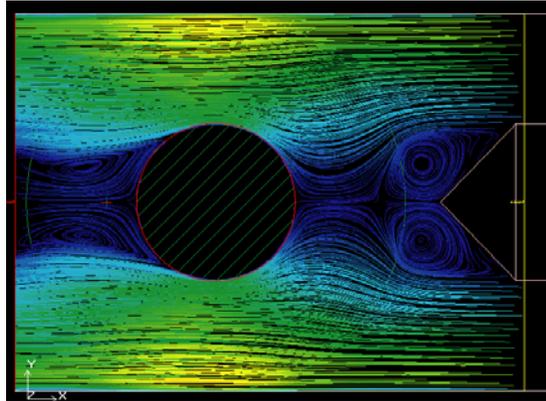
OM17111

The technology utilized to create the Type I processor heatsink design was determined by considering techniques utilized in heat exchanger construction. In particular, the heatsink utilizes a fin-embedding technique such that a single, helically-wound fin is embedded into the solid core to create a good structural and thermal interface (see Figure 90).

Figure 90: Embedded Fin Manufacturing Technology
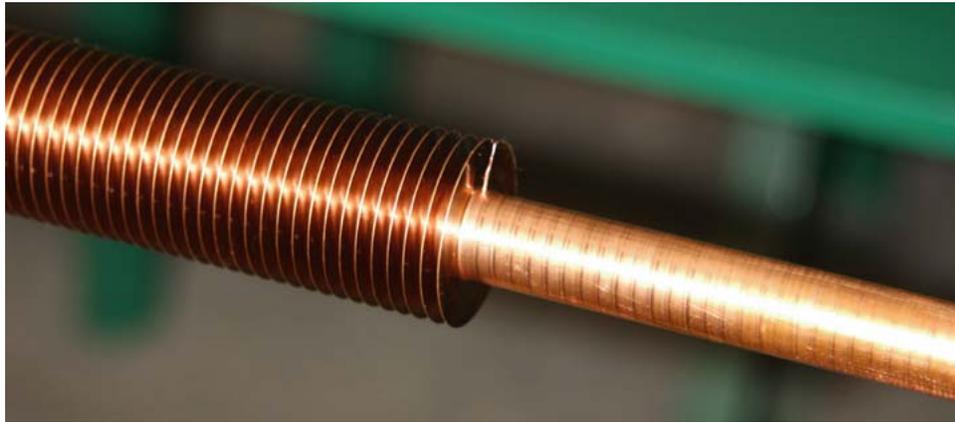
OM17112

Unfortunately, the heatsink dimensions were constrained by current embedded fin manufacturing capabilities. Therefore, in order to maximize the volume within the Type I Thermal Module keep-in, the core diameter was chosen to be 38 mm with 26.4 mm long fins. The heatsink design was also constructed to integrate a structural retention mechanism. Figure 91 shows a top view of the heatsink and Figure 91 shows a bottom view.

Figure 91: Processor Heatsink Design with Cu Core and Al Fins (Top View)

Core cutout

Fins   Core
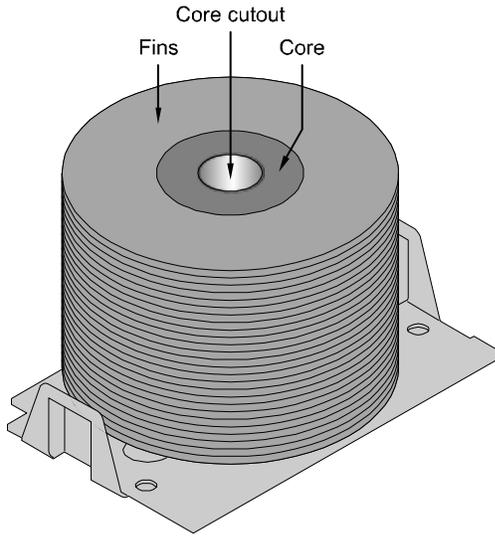
OM17081

Figure 92: Processor Heatsink Design with Cu Core and Al Fins (Bottom View)

Retention Clip

IHS contact area
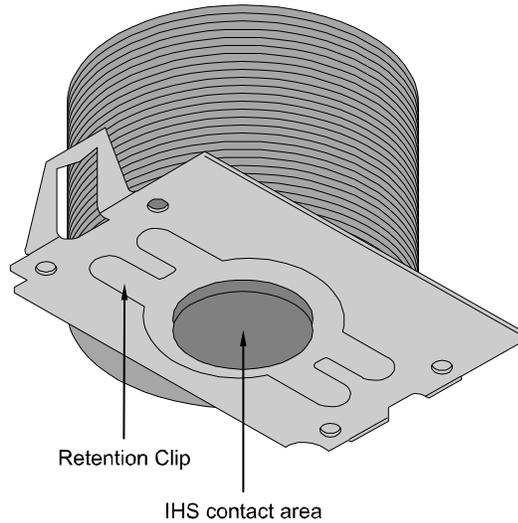
OM17082

The core was chosen to be copper in order to minimize spreading and conduction resistance from the IHS of the processor to the fins. Due to the high density of copper, a small portion of the core was cutout to reduce the heatsink mass below the limit of 900 g. The cutout volume was determined through numerical analysis of the conduction resistance within that region, to minimize impact to performance.

The fins were constructed of 0.4-mm thick aluminum with a pitch selected to be 2.21 mm/fin in order to optimize the convection resistance versus heatsink impedance.

Expressed in terms of percent contributions, the resistance stack of the heatsink when operating at a high power, high flow rate condition is illustrated in the bar chart in Figure 93.

Figure 93: Resistance Contribution within Embedded Fin Heatsink
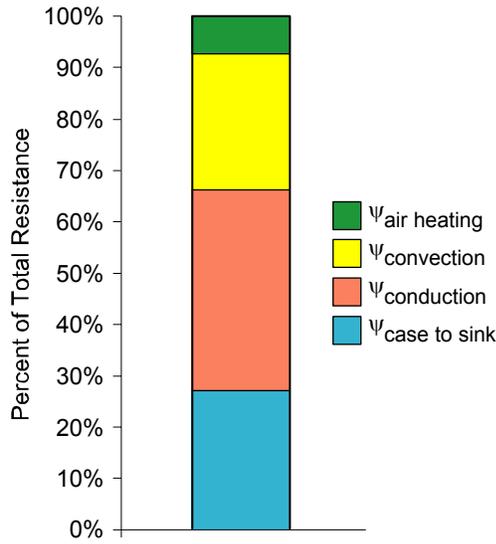
OM17083

2.13.1.3 Duct Design

As mentioned in preceding sections, it is necessary to minimize the airflow bypass to the heatsink in order to meet the flow requirements of the heatsink and the system. For the Type I reference design, the duct was intended to provide an interface to the fan and heatsink while providing the structural retention support for the TM to the system chassis. In the case of the embedded fin heatsink, it was especially important to provide a fully-enclosed duct solution to prevent the escape of air from the sides of the heatsink. It was also determined, that the duct could provide additional performance improvements for the heatsink by diverting air at the downstream portion of the fins to minimize the wake region behind the post. The flow benefit can be easily understood by inspection of the CFD image in Figure 94.

Figure 94: Top View of Flow Vectors through Embedded Fin Heatsink with Downstream Flow Diverters
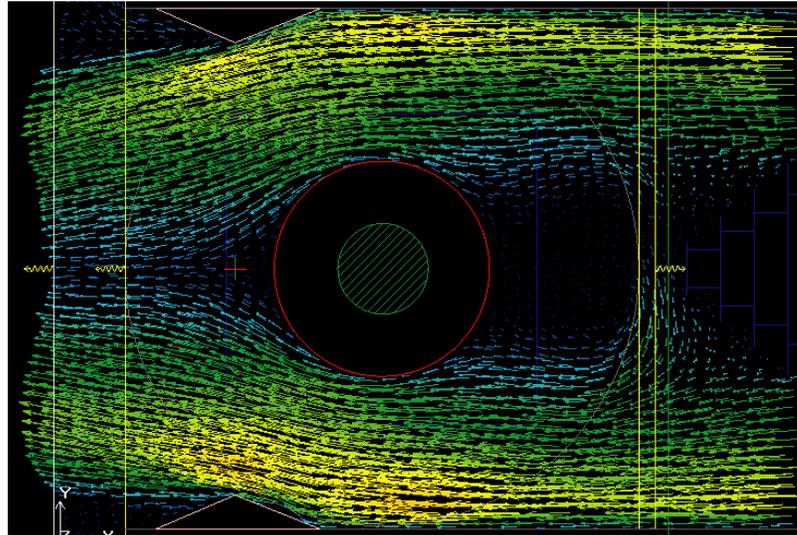
OM17113

The flow diverters within the Type I reference design were incorporated into the design of the duct and are illustrated in Figure 95.

Figure 95: Flow Diverters at Downstream Side of Heatsink
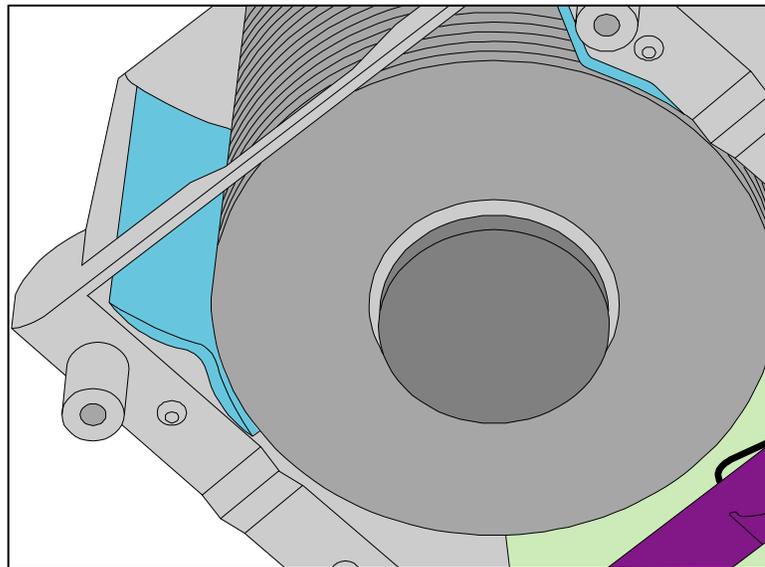
OM17084

The duct also provides a feature by which to retain a seal that interfaces the duct to the interior chassis wall. The seal is integral in preventing recirculation of heated air from the system back into the intake of the Thermal Module. The rubber seal also provides structural compliance at the interface of the TM and

chassis. It should also be noted that since the seal attaches to the duct, and not the fan housing, it provides the opportunity to use industry-standard 92-mm fans within the Thermal Module.

2.13.1.4 Performance Validation

The performance of the Type I Thermal Module was measured experimentally within a 12.9 liter reference design system. Performance was based on the ψ_{ca} performance as measured by a thermocouple attached to the IHS, measuring Tcase, and four ambient thermocouples at inlet to the Thermal Module Interface.

At the Thermal Module maximum operating point and a Design Ambient condition of 35C, the reference design meets the Platform Compatibility Guide '04 B requirement.

2.13.2 Type II Thermal Module Design (for 2004 Performance Targets)

A Type II Thermal Module reference design was also developed utilizing these guidelines and is provided as an example demonstrating the design considerations and engineering approach. The Thermal Module fan is comprised of a 67-mm impeller, high performance motor, an integrated stator, and an attached flow partitioning device. The heatsink is a stacked fin, parallel plate fin design utilizing a forged copper base and copper fins. The duct interfaces to the Thermal Module Interface with a rubber seal and also provides the structural retention to the heatsink. Figure 96 shows a Type II Thermal Module attached to a picoBTX motherboard. Details of each component are described in the following sections.

Section 2.13.1 showcases the design of the Intel Reference Design Type I Thermal Module for 2004 Performance targets. This section describes the design of the Intel Reference Design Type II Thermal Module for 2004 performance targets. After reading these sections, refer to Section 2.13.3 for engineering suggestions for 2004 Mainstream Type I and Type II Thermal Module development.

Figure 96: Type II Thermal Module Attached to a PicoBTX Motherboard
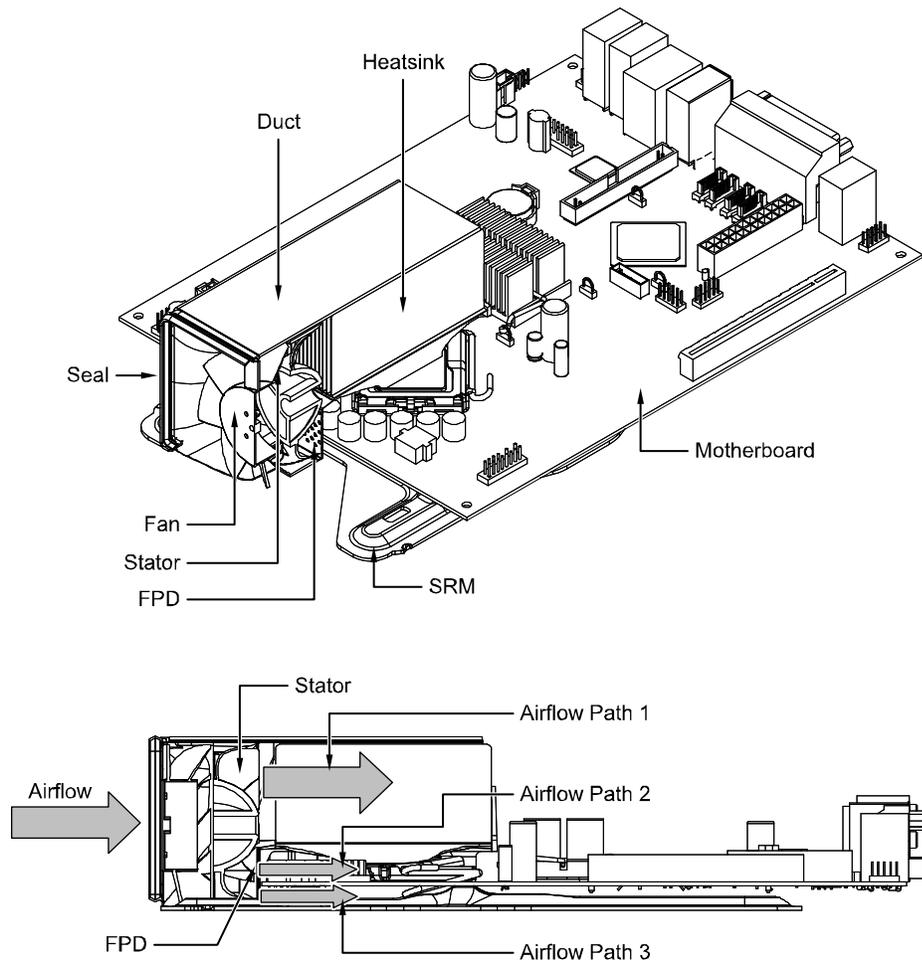
OM17085

2.13.2.1 Fan Assembly

The design objective of the reference design fan assembly for the Type II Thermal Module was identical to that of the Type I Thermal Module; that is, to provide sufficient volume of airflow to the processor as well as the system while maximizing the axial component of the airflow from the fan. The Type II Thermal Module utilized the same impeller blade design as the Type I reference design, at a blade diameter of 67 mm, and utilized the same motor with a slight winding change to obtain the correct RPM values at the high and low values of the operating speed range.

The Type II reference design impeller also utilizes the same airfoil shape and 30° blade angle as the Type I reference design (see Figure 97).

Figure 97: Type II Reference Design Impeller
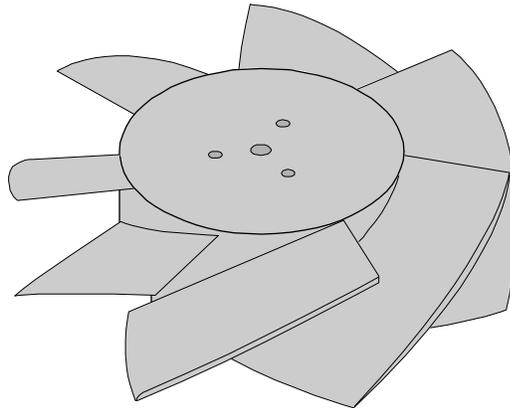
OM17086

The engineering approach utilized to assess the stator geometry was also similar to Type I. However, the resulting stator geometry is different than that used in the Type I reference design stator because the speed, pressure, and airflow vector from the Type II impeller is different than those for the Type I impeller. Specifically, the stator blade angle and tip-to-hub twist are different. The Type II reference design stator is shown in Figure 98.

Figure 98: Stator Design for Type II Reference Design Thermal Module
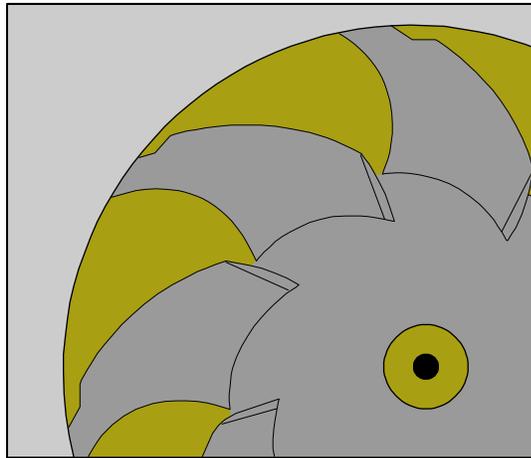
OM17087

A significant difference in the Type II Thermal Module volumetric and the Type I is the shape of the Thermal Module inlet area. As noted in the *BTX Interface Specification*, the Type I Thermal Module Interface is nearly square whereas the Type II Thermal Module Interface is not as tall as it is wide. Therefore, the duct and housing design for a Type II Thermal Module must consider the transition from a rectangular Thermal Module Interface through a circular fan, and then back out to the allowable width of the heatsink. Due to this transition, the fan housing was developed to minimize pressure losses and flow restrictions into and out of the fan. During development, it was necessary to evaluate different housing geometries to address these issues.

Two different housing geometries, at the exit of the impeller, were considered. When looking from the side of the fan assembly, these geometries can be described as concave and convex. Through fluid numerical analysis, the geometries were evaluated based on impact to pressure loss for a fixed flow rate. It was found that a convex housing geometry provided the smallest impact to pressure loss when used in conjunction with the integrated Type II reference design stator blades. Thus, convex housing is used in the thermal module.

Similar to the Type I design, the Type II Thermal Module incorporates a rubber seal at the Thermal Module Interface to restrict flow recirculation at the fan inlet. Unlike the Type I reference design, the Type II reference design seal is incorporated into the fan housing. By attaching the seal directly to the fan housing, it was possible to maximize the impeller diameter within the Type II Zone A and Thermal Module Interface height limitation. The Type II reference design impeller and fan housing are shown in Figure 99.

Figure 99: Type II Reference Design Impeller and Housing
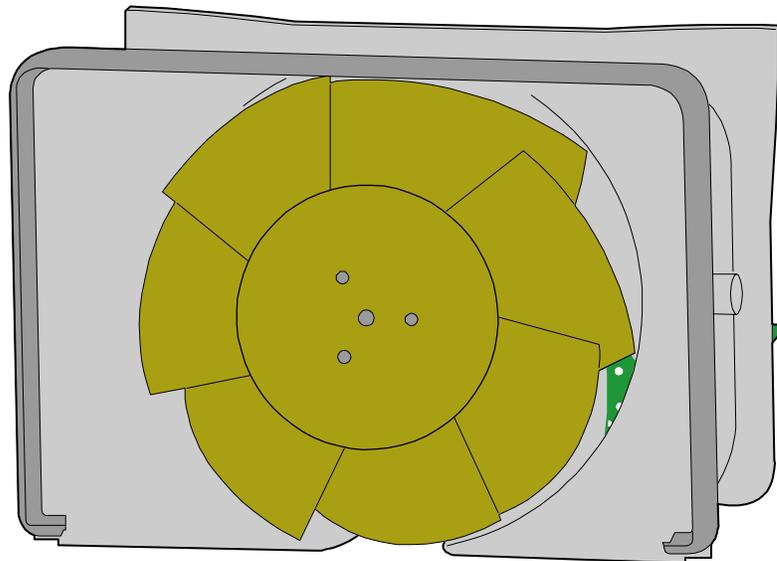
OM17088

Another significant difference between the Type I and Type II reference designs was the need for a flow partitioning device within the Type II due in large part to the difference in ratio of heatsink-to-Voltage Regulation inlet area. However, it should be noted that not every Type II Thermal Module may need to incorporate a flow partitioning device, nor will it be the case that a Type I Thermal Module will never benefit from a flow partition device. Therefore, it is necessary in each design to assess the flow rate requirements and impedance of the heatsink versus that of the below and above board flow region.

In order to provide the necessary flow to the reference design Type II heatsink while maintaining the minimum airflow requirements to the motherboard components, a flow partition device was inserted to restrict the airflow to the above and below board airflow path. The flow partition device increases the resistance of that path, essentially forcing more airflow into the heatsink. The flow partition device design is effectively a porous sheet metal screen located immediately downstream of the stator exit, as shown in Figure 100.

Figure 100: Type II Thermal Module Flow Partitioning Device
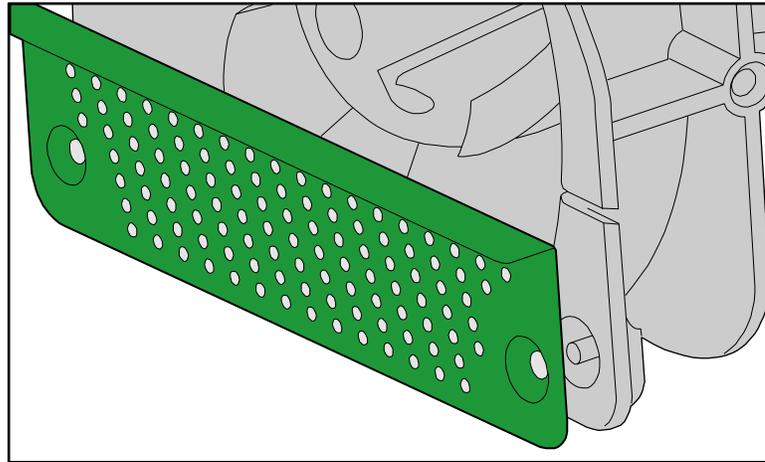
OM17089

Without flow partitioning, the Type II reference design resulted in an airflow split to the heatsink and above-below board paths of 59% and 41%, respectively. However, in order to meet Platform Compatibility Guide '04 B thermal requirement the heatsink required significantly greater airflow. Instead of increasing the Thermal Module fan operating point to achieve the required flow, the flow partition device was used to increase the airflow to the heatsink; effectively allowing a lower fan operating speed in this configuration. The perforations in the flow partitioning device were adjusted until the airflow split was 85% to the heatsink and 15% to above-below board path. The Type II reference design flow partitioning device has a Free Area Ratio FAR of 0.15. The benefit of the split is easily realized by applying the system impedance curves to the Type II fan curve as summarized in the following table. While the system flow rate is diminished by approximately 2.5 CFM it is shown that the heatsink flow rate increases by approximately 7.5 CFM. The advantage of flow partitioning can be seen from Figure 101. The high flow impedance of flow partitioning device forces the fan airflow to go through heat sink.

Table 22: System Flow Rates With and Without Flow Partitioning Utilizing the Same Fan Curve

Flow Rates	Without Flow Partitioning	With Flow Partitioning
System Flow (CFM)	36.5	34
Heatsink Flow (CFM)	21.5	29
VR Flow (CFM)	15	5

Figure 101: Airflow Path with Flow Partitioning Device (Cross-Section Side View of Type II Reference Design Thermal Module)
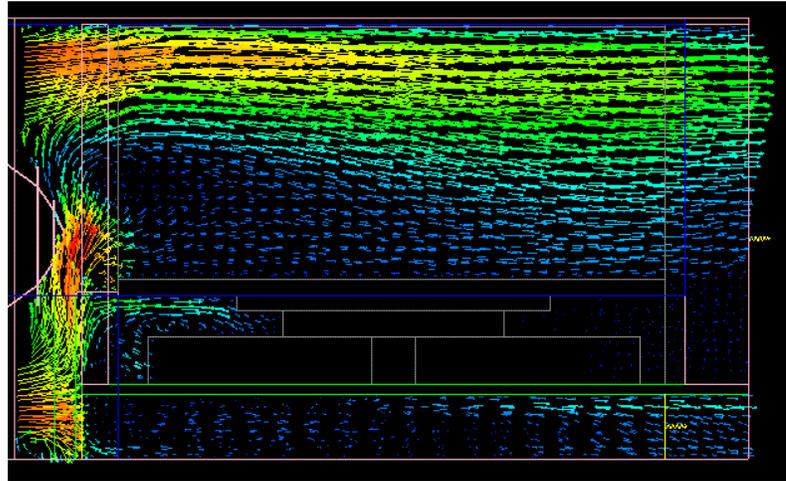
2.13.2.2 Heatsink Design

A Type II Thermal Module has less available height, therefore it is desirable to maximize the volume available within the keep-in while maintaining high fin efficiency. For this reason, the Type II heatsink utilizes copper fins, stacked along a horizontal copper base, that occupy the entire available volume (see Figure 102).

Figure 102: Type II Thermal Module Heatsink
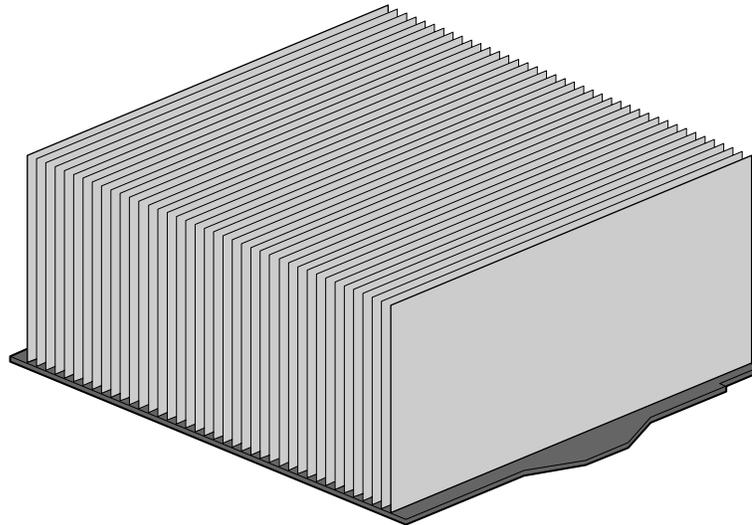
OM17090

Use of high thermal conductivity solder as interface between the fins and heat sink base minimizes the contact resistance (see Figure 103). The fin thickness of 0.4 mm and pitch of 2.31 mm was optimized through the use of 1^{st} order hand calculations and numerical conductive and convective detailed thermal analysis. Similar to the Type I heatsink analysis, the approach utilized a velocity profile that represented the detailed vector flow from the reference design fan and stator. The airflow pattern in heat sink is depicted in Figure 104.

Figure 103: Copper Stacked Fin Technology

Figure 104: Airflow through Type II Heatsink (Cross Section Side View)
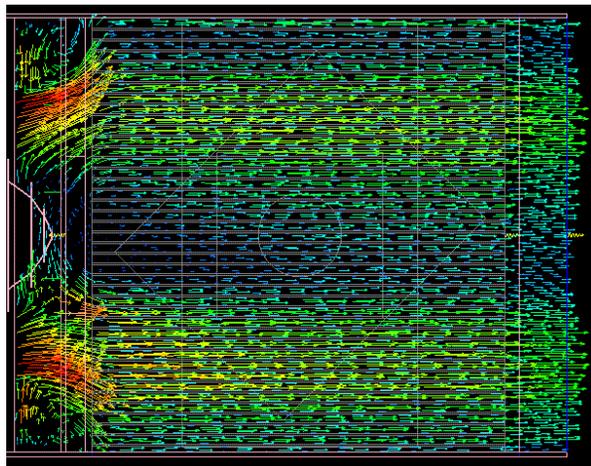
The base of the Type II reference design heatsink has a unique shape and was designed minimize spreading losses and conduction resistance while minimizing its weight. A view of the bottom surface of the heatsink base is shown in Figure 105. This optimized shape is manufactured using a microforging manufacturing technology. After the forging process, secondary operations are performed to finish the IHS pedestal surface and the fin attach surface to ensure reasonable flatness.

Figure 105: Type II Reference Design Heatsink Base (Bottom View)
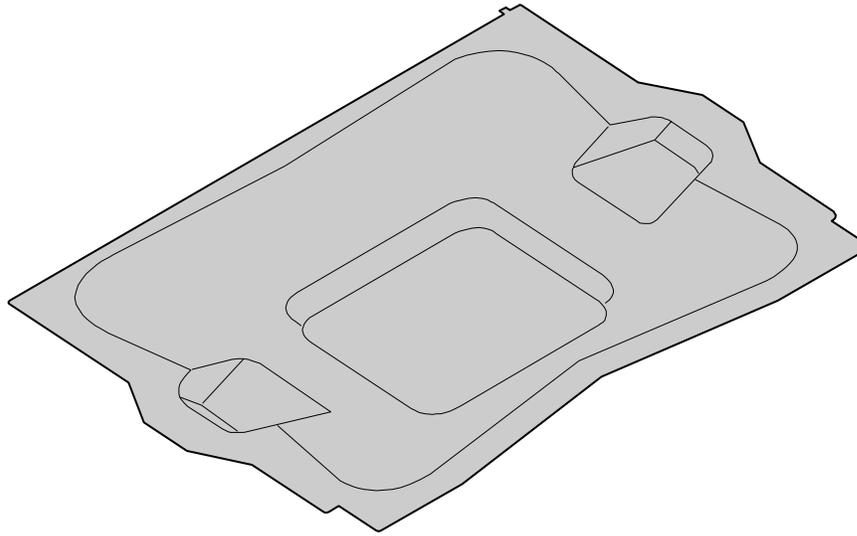
OM17091

The heatsink base also incorporates structural and mechanical interface features at the edges that mate with the Thermal Module duct. With this type of interface to the duct, a heatsink clip is unnecessary. This type of heatsink base design will be considerably more stiff than one that employs a clip to attach the heatsink to the duct. This will be important when designing the duct to apply the appropriate preload, since the preload requirement is a function of the Thermal Module stiffness and the Thermal Module stiffness will change significantly between these two configurations.

The bar chart in Figure 106 provides the percentage breakdown of the four key resistance components of the Type II reference design heatsink.

Figure 106: Resistance Stack-up of the Type II Heatsink
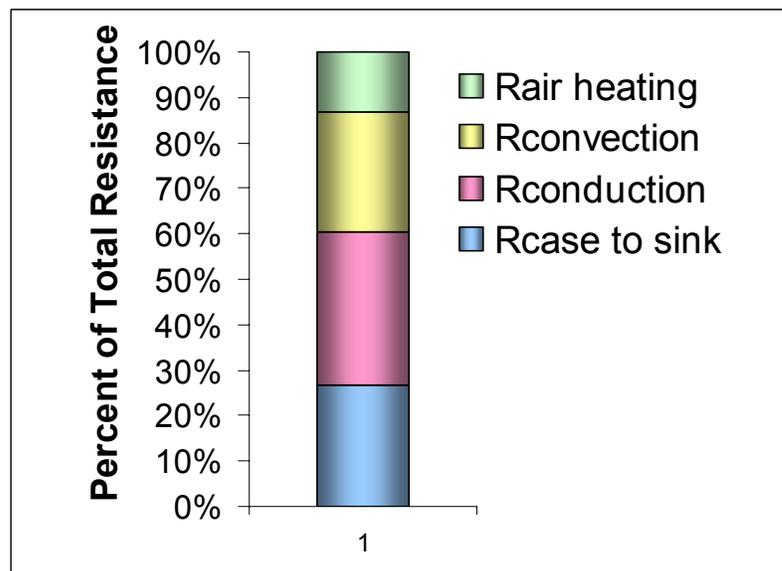
2.13.2.3 Duct Design

Like the Type I Thermal Module reference design duct, the Type II design was also developed to minimize airflow bypass around the heatsink, as well as house the various Thermal Module components. The Type II reference design duct also provides the structural interface to the chassis. The exploded assembly view of the Type II reference design in Figure 107 provides a good understanding of the duct design.

Figure 107: Exploded View of Type II Thermal Module Assembly
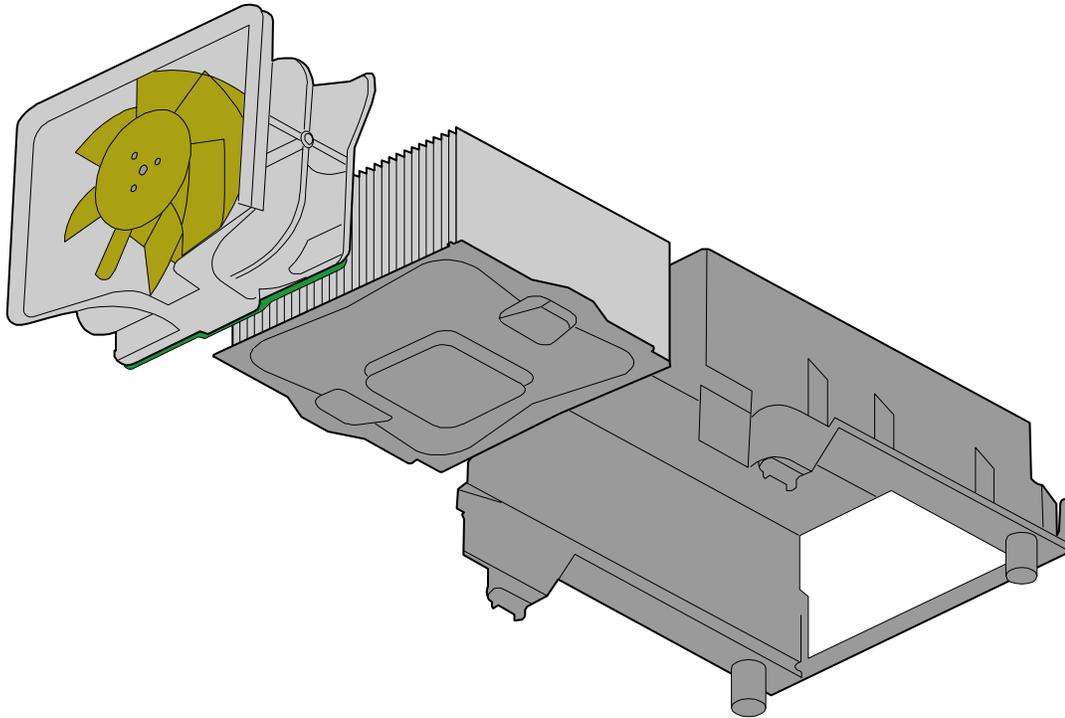
OM17093

2.13.2.4 Performance Validation

The performance of the Type II Thermal Module was measured experimentally within a reference design system. Performance was based on the ψ_{ca} performance as measured by a thermocouple attached to the IHS, measuring Tcase, and four ambient thermocouples at inlet to the Thermal Module Interface.

At the Thermal Module maximum operating point and a Design Ambient condition of 35C, the reference design meets the Platform Compatibility Guide '04 B requirement.

2.13.3 Type I and Type II Thermal Module Designs for 2004 Mainstream Targets

Intel segments the processor guidelines into Performance and Mainstream to offer maximum flexibility in system and motherboard design. The platform requirements for Platform Compatibility Guide '04 A are typically easier and less expensive to implement than those that meet the Platform Compatibility Guide '04 B platforms, since they scale with the required power delivery and dissipation requirement. It is Intel's intention to provide new Performance and Mainstream guidance each year, and to have the requirements

for the next generation Mainstream targets align with the platform capabilities required for the previous year's Platform Compatibility Guide.

The Thermal Module case study descriptions provided in previous sections target the Platform Compatibility Guide '04 B. This section will provide coarse guidance on Thermal Module design options for Platform Compatibility Guide '04 A. In outlying the Thermal Module changes that the reduction in required performance allows, Intel explored cost and acoustic reduction options (see Figure 108 and Figure 113).

Figure 108: Platform Compatibility Guide '04 A Thermal Module Design Options
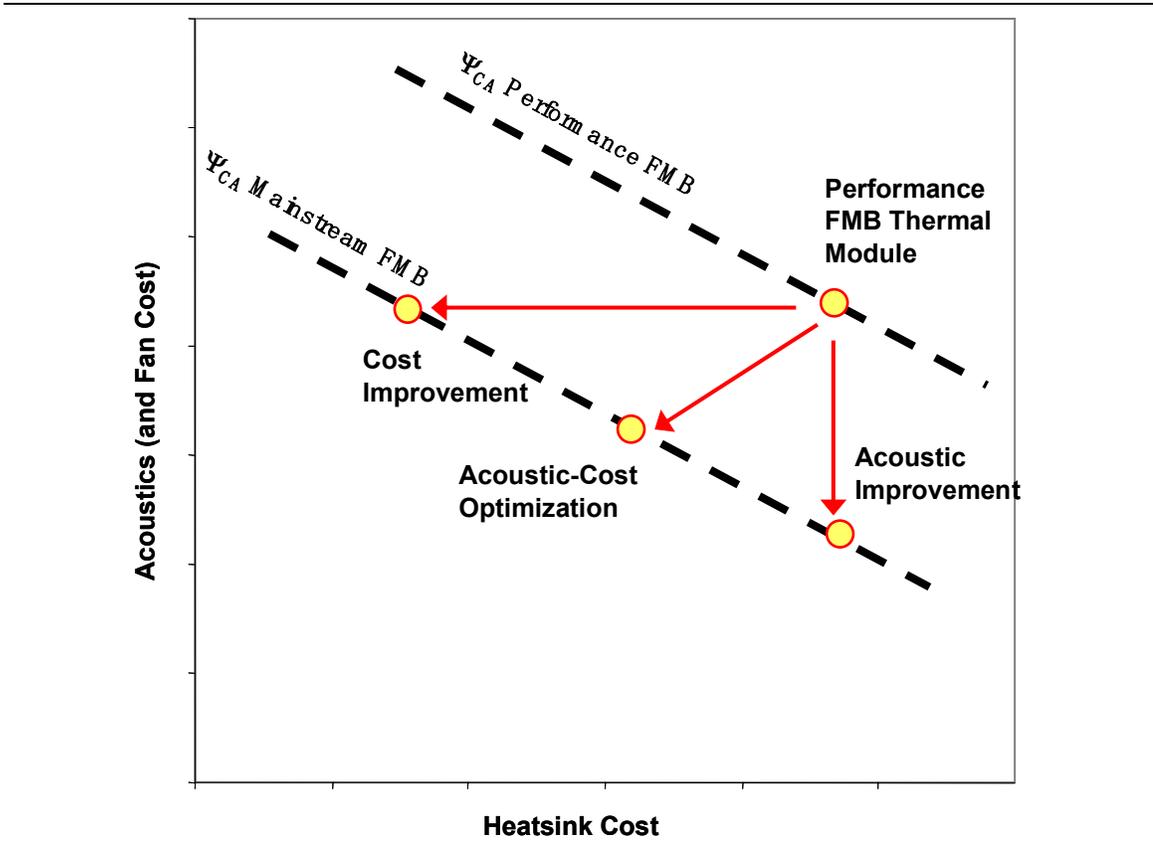
2.13.3.1 Mainstream Thermal Module - Cost Reduction

One option that Intel considered was to modify the Platform Compatibility Guide '04 B Type I Reference Design Thermal Module (see the embedded fin description in Section 2.13.1.2). A reduction in the fin count and copper core height would reduce the Thermal Module performance and cost. In fact, reducing the Type I Reference Design core height until the predicted heatsink performance met the required Platform Compatibility Guide '04 A thermal performance allows this reduced-height heatsink to be used in a Type II Platform Compatibility Guide '04 A Thermal Module.

The Type I Platform Compatibility Guide '04 A reduced core height embedded fin heatsink is approximately $2 cheaper than both the Type I Platform Compatibility Guide '04 B reference design (embedded fin heatsink) and the Type II Platform Compatibility Guide '04 B reference design (copper base – copper stacked fin heatsink).

Figure 109 shows a cross-section view of a reduced core height embedded fin heatsink in a Type I Thermal Module assembly, including arrows representing the three airflow paths: through the heatsink, through the Voltage Regulation (VR), and above the heatsink. Notice that the airflow path above the heatsink is created because the heatsink does not use the full height available inside the Thermal Module duct. As discussed in the Type II Thermal Module Case Study (Section 2.13.1), the fan airflow will bypass the heatsink if there is large open channel area available. In the Type II Reference Design Thermal Module, a Flow Partition Device (FPD) was inserted into the Voltage Regulation flow channel to force more airflow through the heatsink. For the Type I reduced core height heatsink illustrated in Figure 109 (side view cross-section) and Figure 110 (front view), a FPD inserted into the open channel between the top of the heatsink and the duct improves the thermal performance by forcing more air through the heatsink.

**Figure 109: Mainstream – Type I Reduced Height Embedded Fin Heatsink
(Side View Cross Section)**
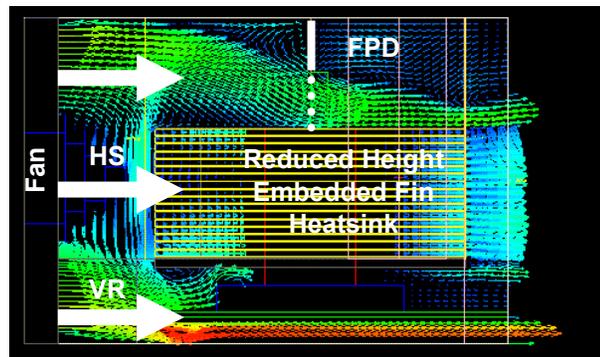
Figure 110: Mainstream - Type I Reduced Height Embedded Fin Heatsink (Front View)
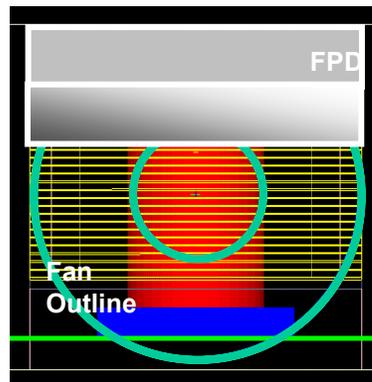
It may be beneficial to have a portion of the Type I FPD be completely solid, while leaving a portion of it perforated, as coarsely illustrated in Figure 110.

Intel's evaluation of the same heatsink design in a Type II Thermal Module indicates that it meets Platform Compatibility Guide '04 A requirements, but requires a smaller FPD (Figure 111).

Figure 111: Mainstream - Type II Reduced Height Embedded Fin Heatsink (Front View)
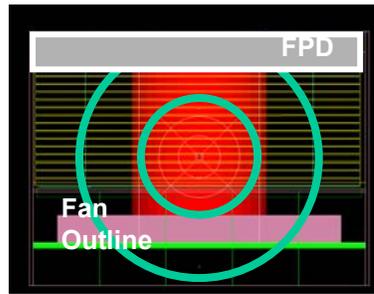
Figure 112 shows that Intel's performance predictions from detailed numerical thermal analysis meets the Platform Compatibility Guide '04 A requirement. The system operating points show the expected airflow through the heatsink when the Thermal Module is inserted into 12.9 liter (Type I Thermal Module) and 6.9 liter (Type II Thermal Module) system level numerical model.
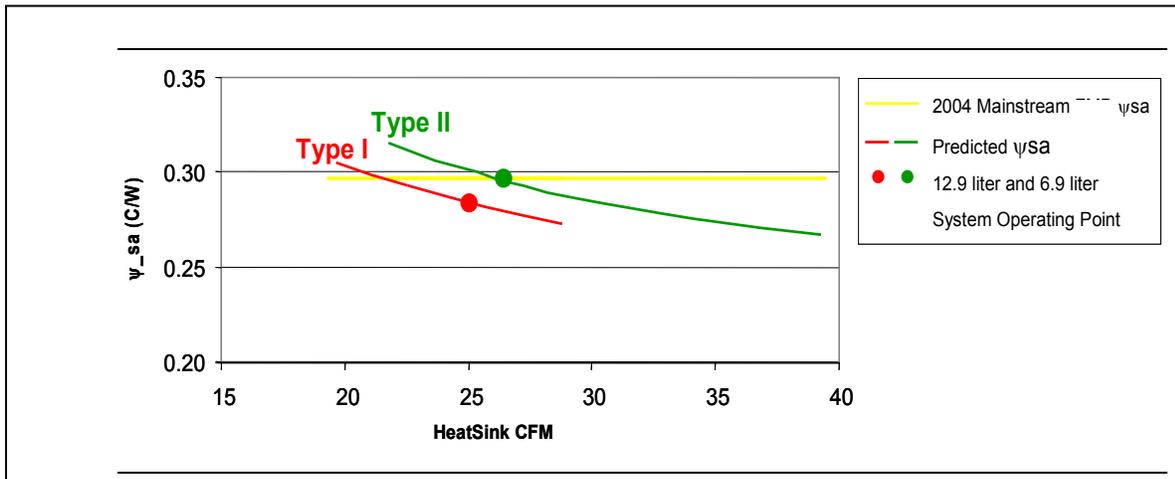
Figure 112: Mainstream - Reduced Height Embedded Fin Performance Predictions

2.13.3.2 Mainstream Thermal Module – Acoustic Improvement

A second option that Intel considered was to reduce the volume airflow through the Platform Compatibility Guide '04 B Reference Design heatsinks by reducing the fan speed. Again, Intel used detailed numerical thermal analysis to determine the heatsink volume airflow required to meet the Platform Compatibility Guide '04 A requirements and the system volume airflow required to meet the subsystem temperature requirements.

Table 23 outlines the fan speed reductions and the predicted reduction in system sound power and Thermal Module cost. The cost improvement is available because the fan motor torque requirement is lower for the lower Platform Compatibility Guide '04 A fan speeds, and a lower fan torque requirements allows the use of lower cost fan motor components.

Table 23: Mainstream - Fan Speed Reduction

Thermal Module		Fan Speed (RPM)	Idle Acoustic (BA)	Cost Improvement
Type I	Platform Compatibility Guide '04 B	4500 maximum 900 minimum	3.72	N/A
	Platform Compatibility Guide '04 A	3250 maximum 600 minimum	3.61	$0.55
Type II	Platform Compatibility Guide '04 B	6800 maximum 1550 minimum	4.1	N/A
	Platform Compatibility Guide '04 A	5500 maximum 1550 minimum	4.1	$0.55

Notice that the Type II minimum fan speed could not be reduced below the fan speed used for the Platform Compatibility Guide '04 B requirement. This is because the Thermal Module fan provides the airflow for the entire system, not just the processor. In system level numerical thermal analysis of the 6.9 liter Reference Design system, the PSU inlet temperature requirement could not be met at lower minimum fan speeds.

The implication from Table 23 is that fan speed reduction is constrained by the subsystem temperature requirements and not the processor requirements. Therefore, the airflow provided to the processor heatsink is more than needed when the Platform Compatibility Guide '04 B heatsink designs are used. This means that the heatsink performance at the minimum fan speed required by subsystem temperatures results in more heatsink thermal performance than is required. This is conceptually illustrated in Figure 113. See Section 2.13.3.3 for the results of the investigation into heatsink technology cost reduction at these subsystem-constrained fan speed settings.

Figure 113: Mainstream - Acoustic Improvement Constrained By Subsystem Temperature Requirements
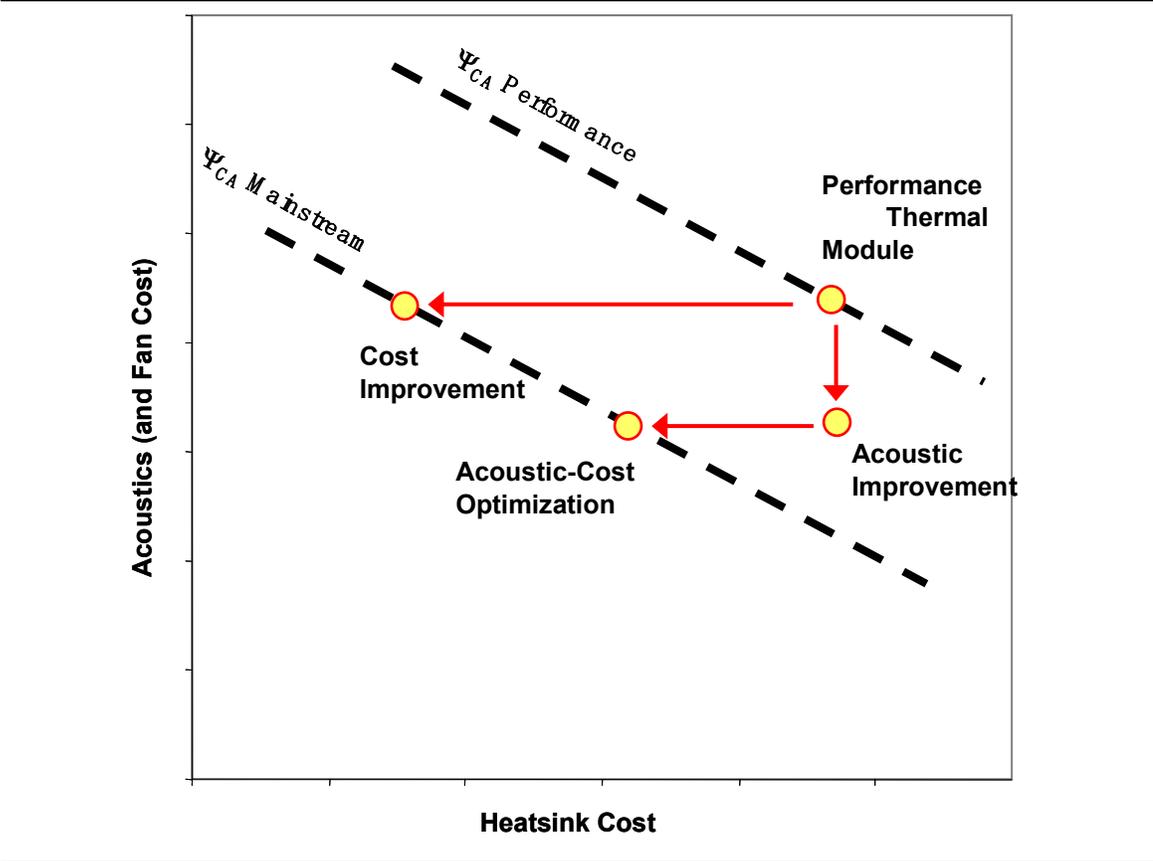
It is important to emphasize the interdependency between the Thermal Module fan and the subsystem temperature behavior. A reduction in Thermal Module fan speed will reduce the system operating point and the reduction in system airflow may lead to subsystem temperature violations. Of course, a lower operating point may be tolerable since a Mainstream system will have lower power loads than a Performance system. Since lower power loads result in lower temperature rise than high power loads, a lower system operating point may be sufficient. In system level numerical thermal analysis of the 12.9 liter and 6.9 liter Reference Design systems, lower processor power (compliant with the Platform Compatibility Guide '04 A) and lower fan speeds (per

Table 23 2004 Mainstream settings) were evaluated and found to be compliant with the subsystem temperature requirements.

2.13.3.3 Mainstream Thermal Module - Cost and Acoustic Optimization

The third option that Intel considered was one that used the minimum Thermal Module fan speed required to meet subsystem temperatures (see Section 2.13.3.2) and investigated heatsink technologies and design options that would reduce heatsink cost. This extensive evaluation included the following heatsink technologies:

- Embedded fin heatsink copper core height Vs heatsink airflow

- Copper base – aluminum stacked fin heatsink, varying fin count and fin height

- All aluminum heatsinks

    - Aluminum embedded fin core

    - Aluminum micro-forged base with aluminum stacked fin

    - Aluminum extrusions (horizontal base and vertical fins)

The outcome of that investigation resulted in the heatsink technology, fan speed settings, and estimated cost savings outlined in Table 24. The absence of all aluminum heatsink designs from the table is an indication that analysis showed these designs incapable of meeting the Platform Compatibility Guide '04 A thermal performance requirements.

Note: The cost improvement is relative to the Platform Compatibility Guide '04 B Type I and Type II Reference Design Thermal Modules identified in Section 2.13.1 and 2.13.2, respectively.

Note: The micro-forged copper base in the heatsink technology description is the same base used in the Platform Compatibility Guide '04 B Type II Thermal Module Reference Design identified in Section 2.13.1.

Table 24: Platform Compatibility Guide '04 A - Optimized Heatsink and Acoustic Performance

	Heatsink Technology /Design Description	Thermal Module Fan Speed (RPM)	Thermal Module Cost Improvement

Type I Thermal Module	• Micro-forged copper base
	• Aluminum stacked fins (soldered to copper base)
	    ○ 0.16 in. thick fins
	    ○ 27 fins
	    ○ Fins extend to Type I Thermal Module duct height
		3250 maximum
		
		600 minimum
			Approximately $3

Type II Thermal Module	• Micro-forged copper base
	• Aluminum stacked fins (soldered to copper base)
	    ○ 0.16 in. thick fins
	    ○ 32 fins
	    ○ Fins extend to Type II Thermal Module duct height
		5500 maximum
		
		1550 minimum
			Approximately $3

3. Acoustic Engineering

This chapter provides guidance on acoustic engineering principles as applied to Desktop Computer Systems. Specifically, issues and strategies related to the effective management of typical sources of noise in a Desktop System are outlined.

Special attention will be paid to integrated Fan Speed Control schemes.

3.1 Measures of Noise

Sound Power and Sound Pressure are two measures of the noise generated by any noise source and set of noise sources. Sound Power is, in essence, the average of noise in all directions, whereas Sound Pressure is noise in a particular direction. For Desktop computer systems, Sound Pressure is typically measured at a location directly in front of the system termed the Seated Operator Position. Since the directional noise characteristics of each Desktop system will vary with each design, Sound Pressure is not the most relevant characteristic to compare its acoustic performance, since some system designs may have very high Sound Pressure in directions other than the Seated Operator Position. The use of the Sound Power characteristic ensures a thorough, appropriate, and relevant comparison of Desktop systems.

Sound Power and Sound Pressure are two measures of the noise generated by any noise source and set of noise sources. Sound Power is, in essence, the average of noise in all directions, whereas Sound Pressure is noise in a particular direction. For Desktop computer systems, Sound Pressure is typically measured at a location directly in front of the system termed the Seated Operator Position. Since the directional noise characteristics of each Desktop system will vary with each design, Sound Pressure is not the most relevant characteristic to compare its acoustic performance, since some system designs may have very high Sound Pressure in directions other than the Seated Operator Position. The use of the Sound Power characteristic ensures a thorough, appropriate, and relevant comparison of Desktop systems.

3.2 Acoustic Performance Targets

Table 25: Acoustic Performance Targets

	Consumer Electronics			Standard PC		
	Outstanding	Target	Minimum	Outstanding	Target	Minimum
Acoustic Maximum	5.3	5.3	5.3	6.7	6.7	6.7
Acoustic Threshold	4.0	4.3	4.6	4.8	5.3	5.8
Acoustic Typical	3.3	3.8	4.3	4.3	4.8	5.3
Acoustic Floor	3.0	3.3	3.6	3.8	4.3	4.8

Note:
Acoustic Maximum – The maximum allowable system Sound Power at any operating condition.
Acoustic Threshold – The high end of acceptable system Sound Power in the majority of operating conditions.
Acoustic Typical – The acceptable system Sound Power in Idle operating condition.
Acoustic Floor – The minimum relevant system Sound Power in Idle operating condition.

3.3 Sources of Noise in a Desktop System

Noise is sound energy that is transmitted by structural or airborne vibrations. Typical sources of noise in a Desktop system include Hard Disk Drive (HDD) spindles and the various radial fans. Generally, the individual noise sources in a desktop system can be summed using Equation 47 to predict the system Sound Power:

Equation 47: Sound Power Summation across Multiple Sources of Noise

$$BA_{system} = \log10(10^{\,BA\;Source\;1} + 10^{\,BA\;Source\;2} + 10^{\,BA\;Source\;3} + \ldots)$$

From this equation, it should be obvious that one tactic for minimizing system sound power is to reduce the number of contributing individual noise sources. In BTX, a typical desktop system may satisfy all the subsystem temperature performance requirements using only two fans (the Thermal Module Fan and the PSU fan); thereby, reducing the typical sound power of a BTX system.

In Desktop Idle Mode per ISO 7779, a HDD is required to be spinning but not accessing. The Sound Power generated by an HDD in Idle Mode will largely be a function of the bearing type used in the HDD spindle.

Each operating fan in the system will contribute to the total system Sound Power. To minimize their contribution in Idle Mode, a Fan Speed Control scheme should be designed so that each fan operates at the minimum operating speed (RPM – revolutions per minute) necessary. Generally, fan speed should increase as ambient temperature and the power draw of various Desktop components increases, since these both contribute to component temperature increases. Fans are responsible for providing the airflow required to maintain component temperatures within their specifications. When fans provide the minimum airflow required to maintain component temperature compliance, the fan speed will be at its minimum, as will the Sound Power generated by fans. The relationship between Sound Power and fan speed is illustrated in Figure 114.

Figure 114: Sound Power as a Function of Fan RPM
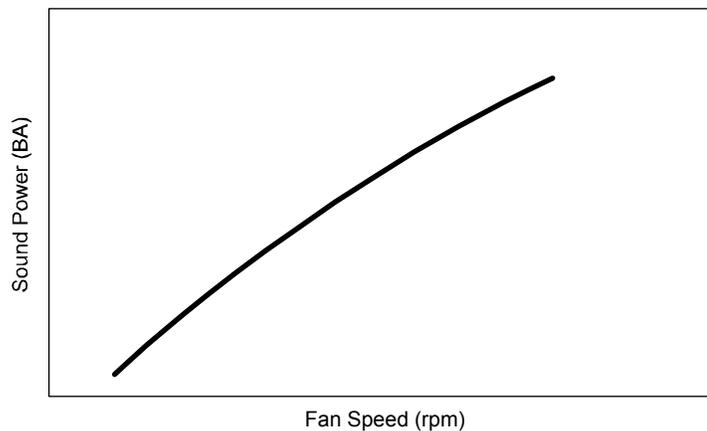
OM16753

3.4 Acoustic Engineering

3.4.1 Near Field Impedance Selection

Acoustic sound energy is often generated in the presence of near field impedance. When impedance is located very near the entrance or exit of a fan the acoustic sound energy often increases. This is an especially important consideration in the design of a Thermal Module and the Thermal Module Interface.

Within the Thermal Module, the location of near field exit impedance should be engineered to balance the airflow benefit against the acoustic detriment. Conditioning the airflow from a fan using a stator is discussed in detail in Section 2.8. If this type of engineering is pursued, characterization of the acoustic impact of the stator should be used to understand the integrated airflow and acoustic behavior. If the flow is not conditioned, the relative distance between the fan exit and heatsink entrance should be evaluated for its acoustic impact versus thermal performance benefit.

The Thermal Module Interface required by the *BTX Interface Specification* may be a ventilated opening or a fully open area. If ventilated, there may be acoustic benefit available by moving the vented surface away from the Thermal Module fan inlet. The *BTX Interface Specification* does not require that the vented area be on a concurrent plane with the Thermal Module Interface, so an offset vent design is completely compliant with the specification. Figure 26 in Section 2.4.5.4 shows an offset vent pattern integrated into the required Thermal Module Interface.

3.4.2 Fan Design for Acoustic Performance

Generally, the noise energy contribution from fan motor bearings is negligible. The type of bearing used is selected largely based on the long-term reliability performance required.

Whenever a fast-moving object passes closely by another object, the dynamic change in the regional air pressure and local air separation generates sound energy.

Most fan assembly housings have mechanical features connecting the fan motor to the outer portion of the fan housing, often termed struts. It is inevitable that the rotating blades of the fan's impeller will pass by these struts and generate this type of sound energy. The frequency at which this occurs is termed the Blade Pass Frequency (BPF) and it is a function of the operating speed and number of fan impeller blades.

Equation 48: Blade Pass Frequency

BPF = RPM • Number Blades / 60

BPF is a function of the fan's operating speed and blade count, as illustrated in Figure 115. That is, the BPF is higher at the fan's maximum operating speed than it is at its minimum operating speed. BPF decreases as the number of blades decreases. However, it is important to note that the sound energy created with each blade pass does not change with fan operating speed.

Figure 115: BPF as a Function of Fan RPM and Fan Blade Count

BPF (Hz)

Fan Speed (rpm)

BPF (Hz) N=8
BPF (Hz) N=7
BPF (Hz) N=6
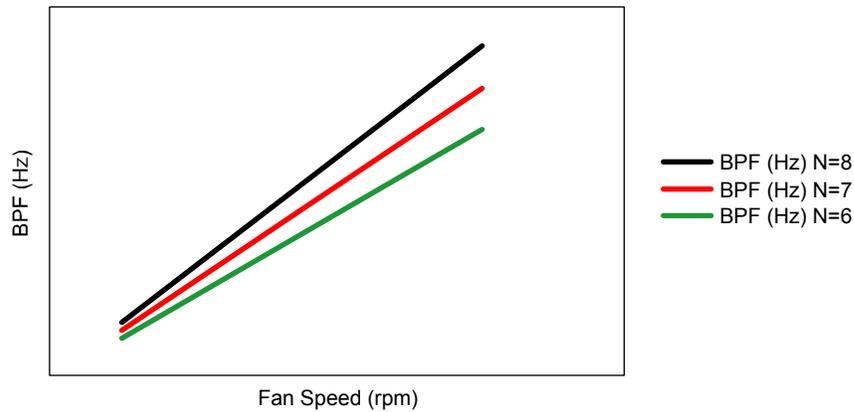
OM16754

If the ratio of the number of impeller blades to the number of fan struts is an integer, then a blade will pass each strut simultaneously. With this design, the sound energy generated at each strut location will occur at precisely the same time and at the BPF. These individual sound energy contributions will cumulate to increase the total sound energy at the BPF.

In a frequency spectrum graph of the sound energy generated by a Desktop system, the BPF and its second harmonic (BPF • 2) are often easy to identify because they are typically the frequencies with the peaks of high sound energy. A typical radial rotating fan frequency spectrum is shown in Figure 116 and the first and second harmonics of BPF are identified.

Figure 116: Typical Fan Sound Energy Frequency Spectrum Plot

Lpf (dB)

Frequency (Hz)
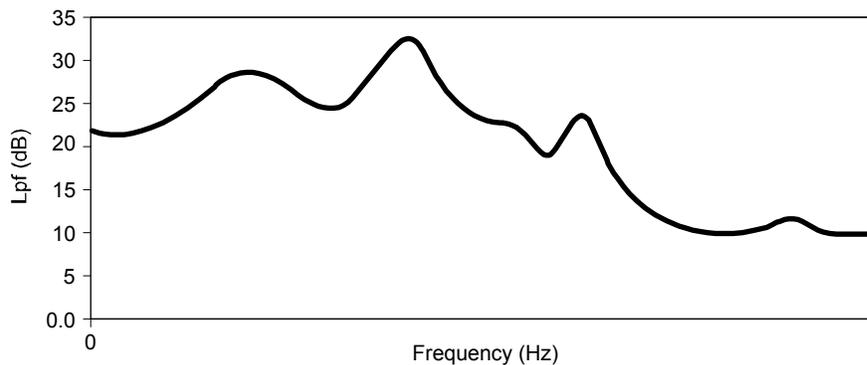
OM16755

A fan or Thermal Module supplier can eliminate blade pass sound energy by eliminating mechanical struts or moving them very far away from the rotating face(s) of the fan impeller. This is a difficult mechanical engineering task given the space constraints imposed on fan assemblies by small Desktop system designs.

However, even with fan struts, a fan or Thermal Module supplier can change the way the sound energy created by each blade pass cumulates. If the ratio of the number of fan blades to fan struts is not an integer, the blade pass will not occur simultaneously at each strut location. Although blade pass will occur at the BPF at each strut location, the individual blade pass sound energies will not cumulate because they do not occur at the same point in time.

Fluid (air) separation within the fan is another potential source of noise. There are several important fluid dynamic characteristics that describe these airflow and noise generation behaviors.

A blade creates airflow by creating changes in pressure. As a blade rotates and moves through the fluid (air), it creates pressure fields along and across both the front and rear surfaces of the blade. In the pressure field illustration in Figure 143, the black object is a blade profile and the pressure gradient uses dark blue to represent the lowest pressure and bright red the highest.

Figure 117: Pressure along a Fan Blade
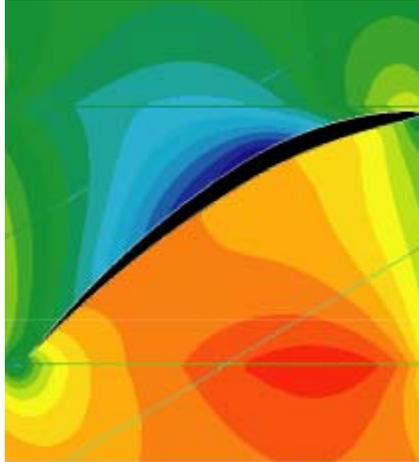
OM17115

These pressure fields create lift and drag forces on the blade and influence the way in which air is pulled into and pushed from the rotating impeller. Blade airfoil shapes designed to increase the lift force are desirable since they improve the volume airflow that a blade can pull in and exhaust. However, the drag force generally increases when the blade is designed to increase the lift force. Since the drag force decreases the volume airflow, the selection of an appropriate airfoil shape for a rotating impeller is a complex fluid dynamic engineering task that is not described in this Design Guide.

The blade airfoil shape also determines the amount of fluid separation that occurs along the blade. Separation occurs in fluid flow when shear forces overcome the fluid's momentum. These forces are highest in areas where the pressure gradient is highest. Since sound energy is created whenever separation occurs, the airfoil shape and fan speed determine the amount of sound energy generated due to separation.

It is also true that separation will occur when there is interaction between leading (in front of the blade) and trailing (behind the blade) pressure gradients. If the blade pitch is too narrow – that is, if the blade count is too high – there is likely to be an increase in the sound power due to flow separation. On the other hand, increasing the blade count increases the impeller's surface area, which can increase the impeller volumetric flow rate. It is likely that wind tunnel measurements of a fan's performance curve and acoustic chamber measurements of a fan's sound power will be required to optimize the blade count.

Fan supplier fluid dynamic engineers should understand and apply these general guidelines to select and engineer airfoil shapes appropriately.

A radial fan design creates a somewhat complex airflow velocity field at its exit. The exit airflow vector is a function of the airfoil shape and blade angle (see Figure 118 for a description of blade angle).

Figure 118: Blade Angle
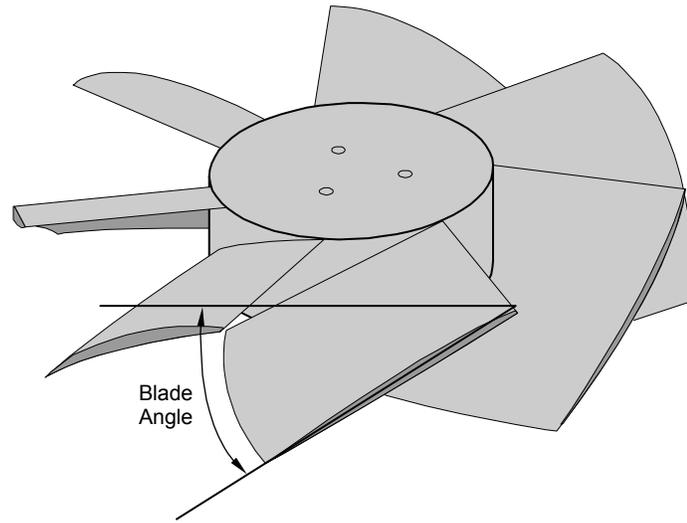
Blade
Angle

OM16756

The vector along which the airflow is directed off each fan blade changes along the length of the blade – from the blade hub to the blade tip (Figure 119 is a Top View of a rotating fan).

Figure 119: Airflow Vector as a Function of Blade Radius

Tip

Velocity
Vectors

Hub
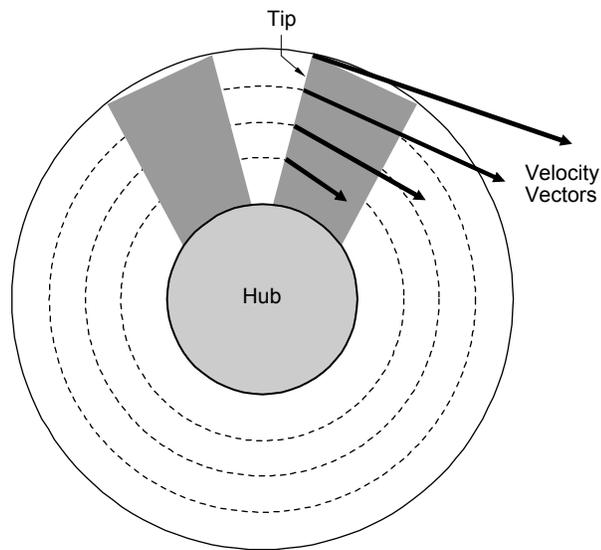
OM16757

Its component directions can define any vector. Blade airflow vector components are defined as axial and swirl, and are illustrated in Figure 120. Axial flow is flow in a direction parallel to the motherboard and orthogonal to the fan face.

Figure 120: Airflow Vector Components – Axial and Swirl
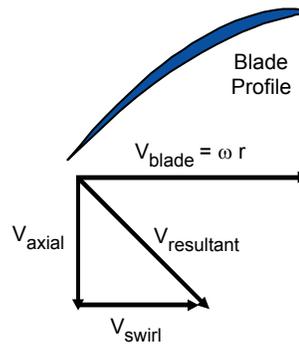
The axial component of the airflow generated by the Thermal Module is the most important component. Axial flow enters the CPU heatsink with less impedance than swirl flow and is, therefore, more efficient in creating heat transfer.

A fan impeller can be engineered to increase the axial component and decrease the swirl component of flow. Again, this is complex fluid dynamic engineering that is not described in this guide.

3.4.3 Additional Airflow Improvements for Acoustic Performance

A fan or Thermal Module supplier may also engineer an increase in the axial component of flow entering the CPU heatsink through the application of a stator. A stator is essentially a fixed (non-rotating) array of airfoil blades positioned near the exit of the rotating fan. Properly engineered, the stator converts a portion of the fan's swirl component to pressure. This increases the fan performance curve and, for a particular CPU heatsink design, will increase the total flow and axial flow entering the CPU heatsink. In Figure 121, this is reflected by the larger $V_{swirl, in}$ than $V_{swirl, out}$ and by the larger $V_{axial, out}$ and $V_{axial, in}$ when the stator is included in the fan's exit airflow path.

Figure 121: Stator Increase of the Airflow Axial Component
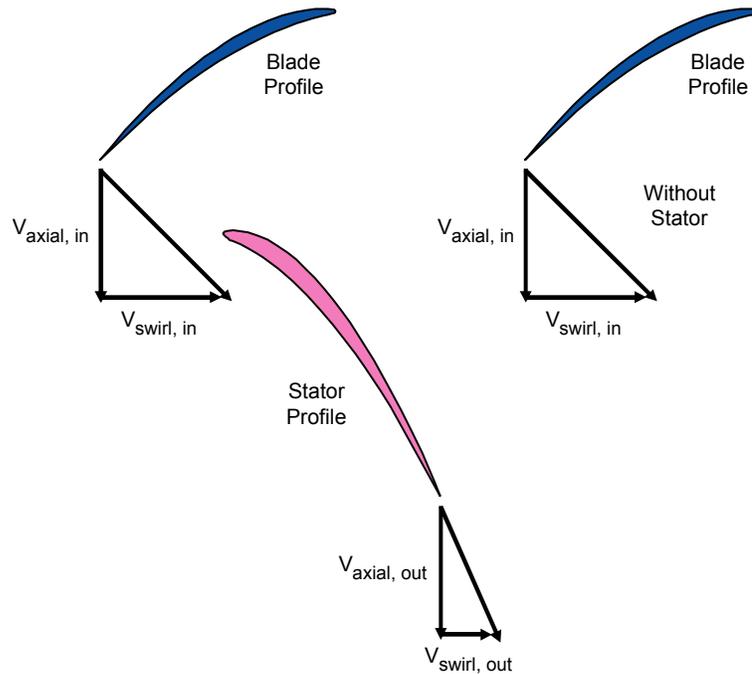
OM16759

The stator's reduction of the swirl component also decreases the flow impedance of the CPU heatsink and reduces the pressure in the rotating impeller portion of the fan. These impedance reductions lower the heatsink impedance curve and allow the fan to deliver even more airflow (see the fan performance curves in the Airflow section of this Design Guide).

The amount of axial flow that a fan can deliver to the CPU heatsink and system is determined by the fan performance curve. The amount of flow that a CPU heatsink or system requires is determined by the system thermal design and the various component heat powers and their temperature specifications.

When a fan is engineered to deliver more axial airflow – through appropriate impeller blade and /or stator design – it can deliver the airflow the system requires at lower fan speeds than a fan that delivers less axial flow.

In summary, a fan engineered to improve the axial component of flow reduces the impedance of the objects in the airflow path and the pressure in the fan. These allow the fan to deliver more useful axial airflow at any fan speed or the required airflow at lower fan speed. Since the fan's sound power is related to its operating speed, a system that uses a fan engineered to improve the axial component of flow will be quieter than one that does not.

The relative location of impedance in the airflow path influences the amount of volumetric flow that a fan can deliver. Impedance sources located very close to the fan's inlet or exit are termed near-field impedances – those farther away are termed far-field impedances. Near-field impedance sources reduce the volumetric airflow capability of a fan more than far-field.

Generally, the fan performance curve is an indication of its ability to delivery volumetric airflow against increasing far-field impedance sources. Since the impact of near-field sources may be greater, it will be relevant for Thermal Module and system integrators to conduct wind tunnel testing to determine the exact impact that a near-field impedance source may have.

Case Study engineering and test validation information for Intel reference design systems (6.9 liter and 13 liter system volumes) will be included in future revisions of this document.

Specific information on the engineering of the Intel reference design Type I and Type II Thermal Module fan impeller and stator will be included in Thermal Module Design Requirements Documents.

3.4.4 System Fan Speed Control

3.4.4.1 General Fan Speed Control Circuit Behavior

The objective of Fan Speed Control (FSC) in Desktop systems is to reduce the noise generated by a system when the system operating condition allows.

Generally, FSC circuits are designed to monitor critical system temperatures and adjust the speed of system fans accordingly. As described in detail in the Airflow section of this Design Guide, the heat transfer from any component or heatsink improves with increasing airflow. Increasing the airflow in a Desktop system is typically accomplished by increasing the operating speed of the system's fans.

The FSC circuit requires that a limit value be provided for each system temperature that it monitors. A circuit diagram for a typical FSC circuit is shown in Figure 122.

Figure 122: Typical Fan Speed Control Circuit Diagram
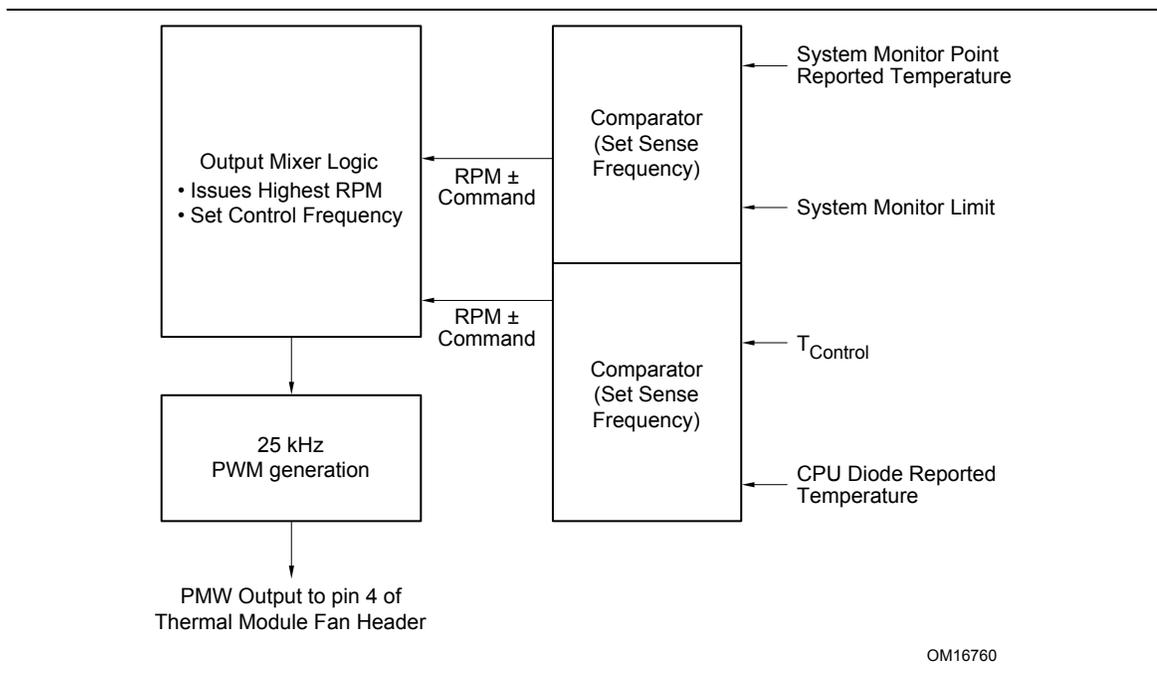
When the temperatures reported to the FSC circuit from each monitor location are below their limits, the FSC should allow the system's fans to operate at their minimum speeds.

When the temperature at any of the monitor locations exceeds its respective limit, the FSC must direct the fan speed to increase. FSC circuit requests to increase fan speed should continue until each reported monitor temperature is at or below its limit.

The relationship between monitor temperature and fan speed is illustrated by starting with Figure 123, which will be the basis for the remaining FSC figures and subsequent descriptions.

Figure 123: Fan Speed Response to Monitor Temperature
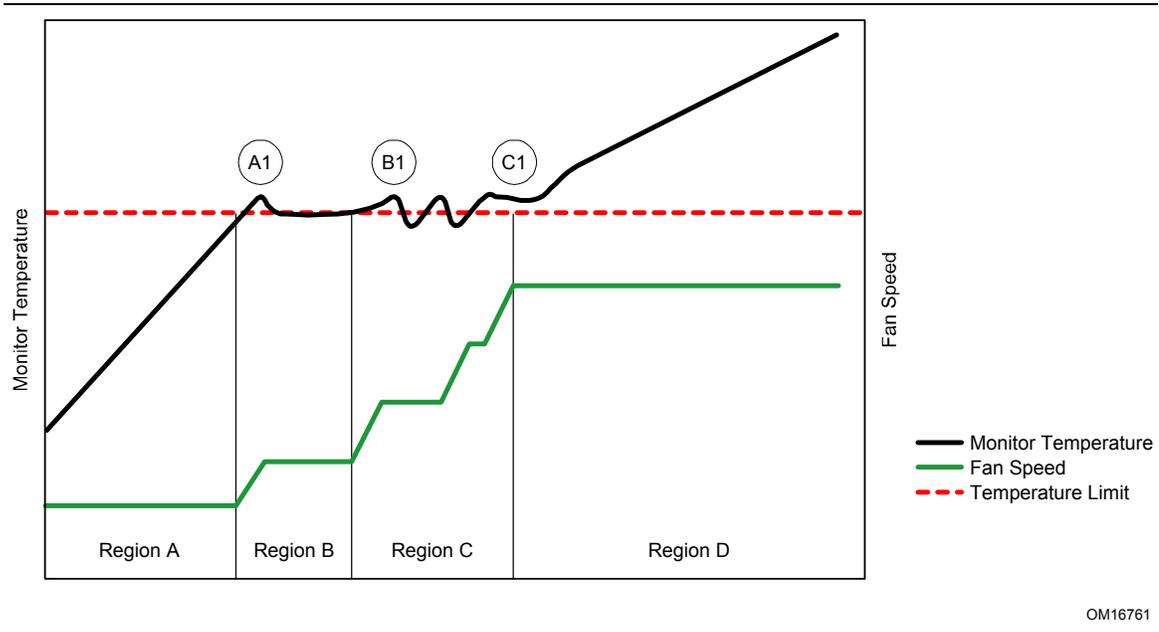
In region A, the fan speed is at its minimum. When the reported temperature at the monitor location is at or below its temperature limit, the FSC circuit will allow the fan speed to be at its minimum operating speed. Note that in Region A, the reported monitor temperature is increasing. This may be due to an increase in the external ambient temperature, an increase in the power applied to system components, or a combination of changes in external ambient temperature and applied power that results in increasing monitor temperature.

Figure 123, Point A1 is a point at which the reported monitor temperature exceeds its limit. The FSC circuit should immediately direct the fan speed to increment.

In Figure 123 Region B, the increase in fan speed is sufficient to maintain the reported monitor temperature at its limit and the FSC circuit will maintain this fan speed. Region B could represent on operating condition in which the external ambient temperature has stabilized at a higher value, the power applied to system components has stabilized, or a combination of changes in external ambient temperature or applied power that results in steady monitor temperature.

Figure 123, Point B1 is the point at which the reported monitor temperature again exceeds its limit. The FSC circuit should immediately direct the fan speed to increment.

In Figure 123 Region C, the FSC circuit continuously detects that the reported monitor temperature exceeds its limit and continuously directs the fan speed to increment. Region C could represent operating conditions similar to those described for Region A (see Figure 123); however, the fan speed increments

requested by the FSC circuit improve the heat transfer behavior to offset what would otherwise be an increase in monitor temperature. Therefore, the monitor temperature is maintained near its limit.

Figure 123, Point C1 is the point at which the fan is operating at its maximum speed. At this point, the fan cannot respond to FSC circuit requests to increment the fan speed despite the fact that the reported monitor temperature is above its limit.

In Figure 123 Region D, the fan speed is at its maximum operating point and the reported monitor temperature continues to increase. Region D could represent operating conditions similar to those described in Regions A and C; however, the fan speed cannot be incremented nor can the heat transfer behavior be improved. Therefore, the reported monitor temperature will continue to climb as dictated by the changing operating conditions.

In this circumstance, the maximum heat transfer capability of the system will determine not only the maximum reported monitor temperature but also the maximum temperature for all system components. If the external ambient temperature or applied component powers exceed the maximum values for which the system is designed, there may be component performance degradation, component failure, or a reduction in the system's long term reliability.

Note: Intel$^{®}$ processors have a Thermal Control Circuit that is designed to detect a processor over-temperature condition and temporarily lower the processor's operating speed to lower the dynamic CPU power. This control design attempts to detect and avoid processor operating temperatures that might cause catastrophic failure or erode the processor's long term reliability. More information on the Thermal Control Circuit is available in processor Datasheets.

If the reported monitor temperature drops below its limit, the FSC circuit should direct the fan speed to decrement. If the reported monitor temperature remains below its limit, fan speed reduction can continue until the fan is operating at its minimum operating speed.

Figure 114 illustrates the relationship between Sound Power and fan speed. It should be obvious that it is desirable to have all system fans operating at their minimum speed across a broad range of system operating conditions.

The temperature at system monitor locations is a function of the ambient temperature external to the Desktop system and the temperature rise introduced by heat transferred into the airflow by powered components. The temperature rise of the airflow is influenced by the number of powered components in the airflow path and as the power applied to these components – as both increase, the airflow temperature rise gets larger, as illustrated in Figure 124.

Figure 124: Airflow Temperature Rise as a Function of the Number of Powered Components and the Component Power Applied
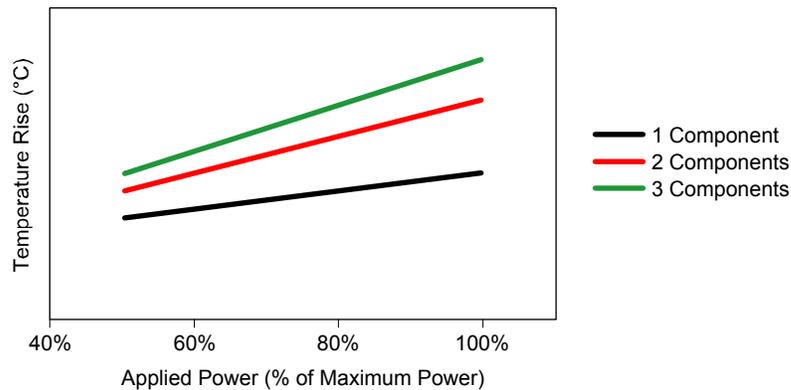
OM16762

3.4.4.2 Fan Speed Control Circuit Input Selection

An independent Fan Speed Control (FSC) circuit within the PSU typically manages the PSU fan speed. PSU fan speed is typically correlated to the power it is delivering to the system and the inlet airflow temperature. Although these two characteristics are often integrated by monitoring the temperature of a single PSU electrical component.

A BTX system will not require FSC circuit communication between the PSU and motherboard. That is, the Thermal Module and PSU fans can be regulated independently.

The FSC circuit for the Thermal Module fan should be integrated into the motherboard design. A control chip mounted to the board and connected to appropriate input and control motherboard traces forms the basis for Thermal Module fan regulation. The motherboard 4-pin fan header recommendations are outlined in the Thermal Module Design Requirements Document.

A Heceta-6e FSC circuit is limited to two monitor inputs, each of which includes a temperature signal from a monitor location and the temperature limit associated with that monitor location.

The Heceta-6e FSC circuit requires that the motherboard fan header for the Thermal Module fan be a 4-pin header that allows Pulse Width Modulation.

It is recommended that one of the FSC circuit input signals be acquired from the CPU diode (reported temperature) and $T_{Control}$ register (diode temperature limit).

There are anticipated system operating states in which the CPU power may be low but other system component power may be high. If the only FSC circuit input is from the CPU diode, the fan speed and related system airflow is likely to be too low to maintain component temperature specifications in these operating states. Therefore, it is recommended that a second input be acquired from a monitor location within the system.

The System Monitor Point is best determined through extensive system-level numerical thermal model execution or extensive prototype thermal testing. In either case, the temperature of critical components or the air temperature near critical components should be assessed for a range of system external temperature, component power, and fan speed operating conditions. The temperature at the selected location for the System Monitor Point should be well correlated to the temperatures at or near critical components.

For instance, in numerical model simulations or prototype testing it may be useful to monitor the temperature at the PSU airflow inlet, the MCH heatsink airflow inlet, the graphics add-in card heatsink airflow inlet, and the memory airflow inlet measure – since the component temperatures are typically directly proportional to their inlet airflow temperatures.

The simulation or empirical testing will need to be conducted across a range of operating conditions representative of the anticipated system use conditions. The following recommended operating conditions could be used to establish the correlation between a System Monitor Point temperature and critical component temperatures.

A system integrator will, of course, ensure compliance with the component temperature specifications at worst case power and external ambient temperature – a condition in which the Thermal Module and PSU fans must operate at their maximum operating speeds.

In the power and external ambient operating condition required by the system integrator's acoustic validation, all component temperature specifications must be met when the Thermal Module and PSU fans are at their minimum operating speeds.

Applications that force the graphics add-in card, memory, or chipset components to consume power near their maximum power limits but allow the CPU power consumption to be low should be evaluated. For instance, applications that decode or encode a CD or DVD (or simulate these application power loads and distributions). It is likely that this operating condition will require the Thermal Module and PSU fans to operate above their minimum operating speed in order to meet the component temperature specifications.

If the evaluation is conducted using numerical applications, it will be important to have the PSU fan curve and fan speed accurately represented, so that its influence on the system airflow pattern, velocity, and temperature is comprehended. The PSU fan speed will typically vary as a function of the PSU electrical load and the temperature of its inlet airflow.

In the evaluation of these operating conditions, a correlation should be established between the System Monitor Point and components' temperature. Different System Monitor Point locations should be evaluated and the location with the strongest correlation to the component closest to its temperature specification should be selected.

3.4.4.2.1   Selection of the System Monitor Point Location

A BTX system will often be designed such that the Thermal Module CPU heatsink exhaust provides the primary airflow stream for the MCH, graphics add-in card, memory, and PSU airflow inlet (see Figure 125). This implies that the System Monitor Point should be located in the Thermal Module exhaust (see Figure 25) and behind the MCH heatsink. Location at this position ensures that increases in the MCH heatsink or Thermal Module exhaust airflow temperature from MCH and CPU power are integrated with any increases in external ambient temperature rise.

Figure 125: Typical BTX System Airflow Illustration
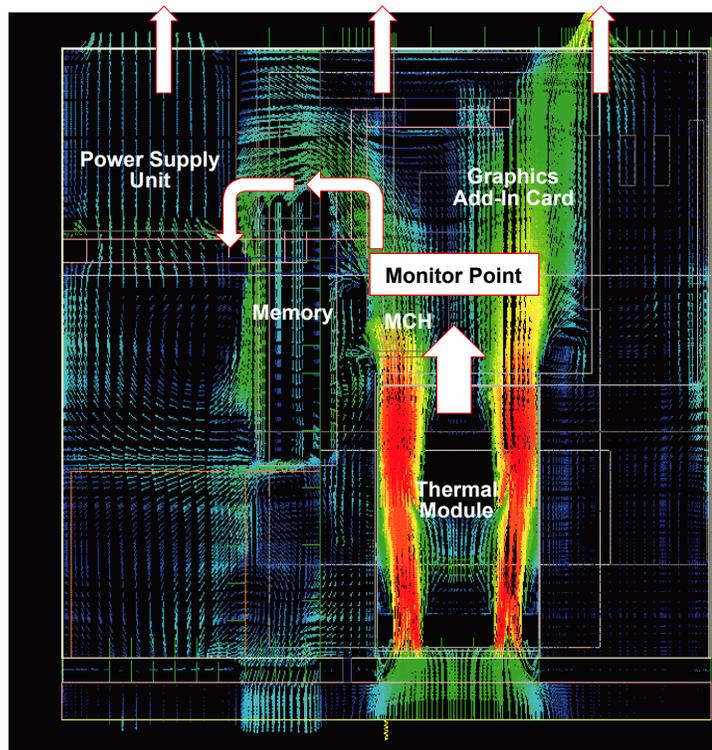
OM16791

The exact System Monitor Point location within the Thermal Module exhaust should be based on the location of the component whose temperature is closest to its specification in the operating condition evaluation outlined in Section 3.4.4.4.

3.4.4.3 Integrated and Optimized Fan Speed Control Operating Behavior

The Thermal Module fan is the primary source of airflow creating appropriate heat transfer behavior for CPU heatsink and the remainder of the system. Therefore, its speed regulation depends not only on the temperature response of the CPU, but also on the temperature response of the System Monitor Point.

The Thermal Module FSC circuit will regulate the Thermal Module fan speed in a way that attempts to keep both temperature inputs below their respective limits.

That is, if both the CPU diode and System Monitor Point are below their respective temperature limits, the Thermal Module fan speed may be maintained. If one or both of the reported input temperatures exceeds its limit, the fan speed will increment until both temperature inputs are below their respective limits or the fan reaches its maximum operating speed. This control logic is integrated in the Heceta-6e circuit.

A system integrator that desires optimal acoustic behavior in all operating conditions should select a Thermal Module whose fan includes a local thermistor. A thermistor mounted in the Thermal Module fan assembly will monitor the temperature of the air entering the system and use this information to augment the fan speed regulation direction provided by the FSC circuit.

If this augmentation is not included, the acoustic behavior of the system will be compromised in certain expected operating conditions. Figure 126 and Figure 127 illustrate how the Thermal Module fan

thermistor optimizes system acoustic behavior in an operating condition where the external ambient temperature is high (for instance, the maximum ambient temperature for which the system is designed and rated) and the steady state CPU power continually increases.

Figure 126: CPU Diode Temperature and Thermal Module Fan Speed as a Function of Component Power in Fixed High Ambient Temperature
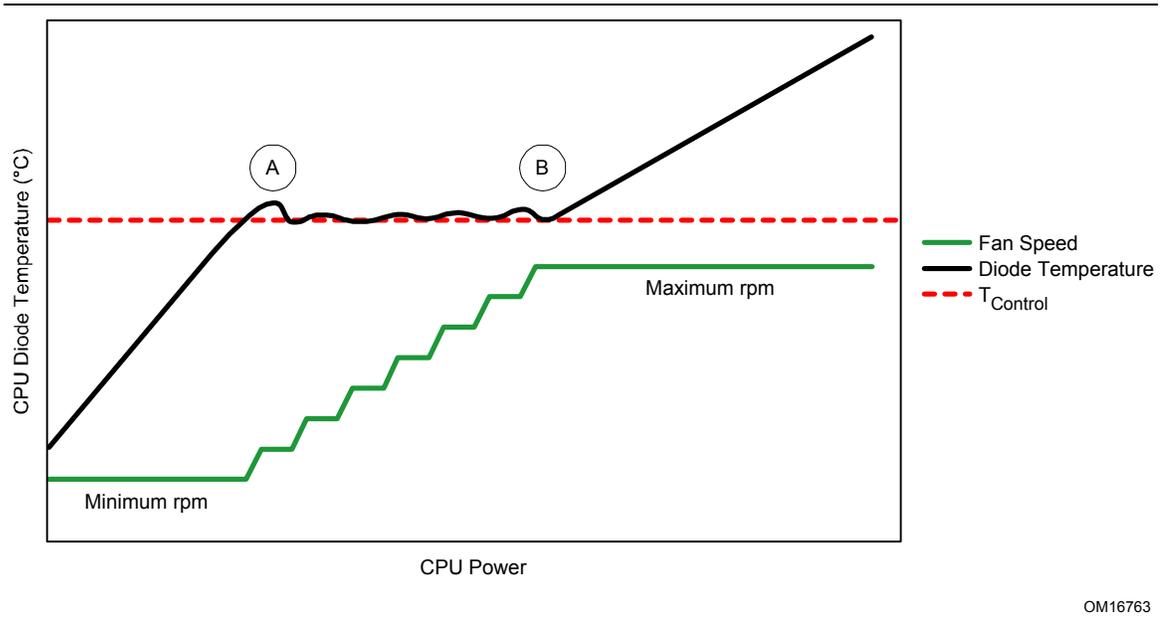
OM16763

Figure 127: CPU Case Temperature as a Function of Component Power in Fixed High Ambient Temperature
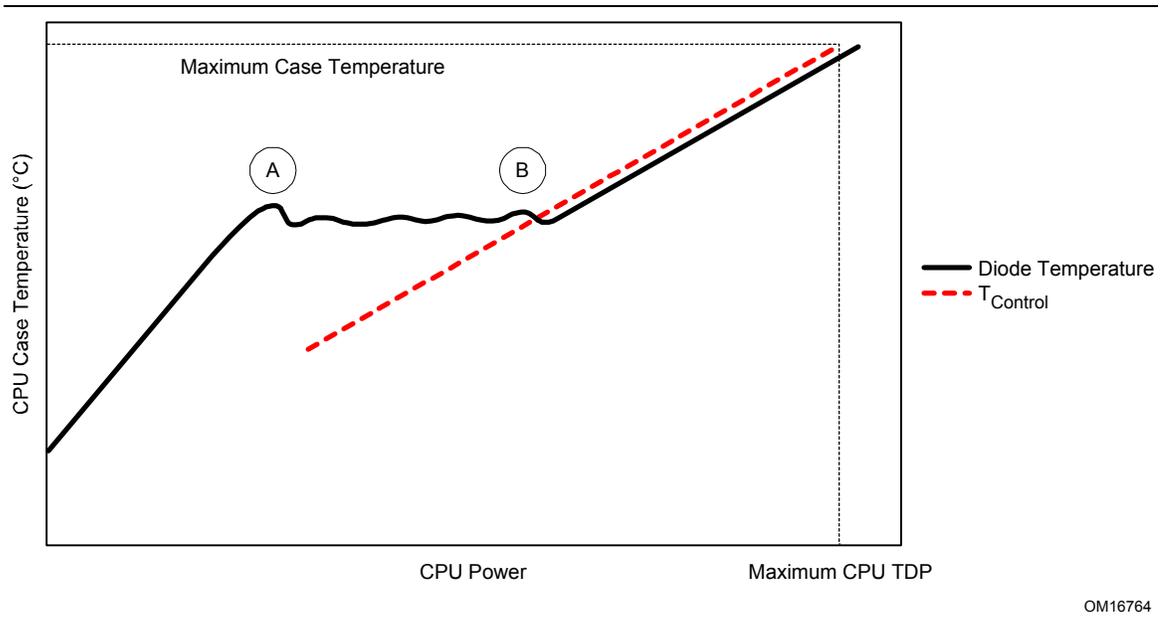
OM16764

Figure 126 and Figure 127 illustrate the CPU diode and case temperature response as CPU power increases (the reference points A and B are the same in each figure). Notice that the CPU case temperature will eventually follow the CPU Thermal Profile and reach the CPU maximum case temperature specification when the CPU power reaches the maximum Thermal Design Power (TDP). More information on the Thermal Profile Specification is available in CPU Datasheets.

This illustration assumes that a system designer has engineered the system's thermal behavior such that it matches the CPU Thermal Profile when the Thermal Module is at its maximum operating speed and the external ambient is at the system's maximum rated temperature.

As noted, the above illustrations show the system's various responses when the external ambient temperature is near its maximum rated value. However, Figure 128 through Figure 130 illustrate the CPU diode temperature, Thermal Module fan speed, and CPU case temperature at a low ambient temperature (for instance, an air conditioned room). The response when a Thermal Module fan thermistor augments the FSC circuit regulation is compared to FSC with thermistor augmentation. Notice that CPU diode temperature will reach the $T_{Control}$ limit at a higher power level when the external ambient is low (the inflection point labels in these Figures are at the same power as on the previous figures).

Figure 128: CPU Diode Temperature and Thermal Module Fan Speed as a Function of CPU Power in Fixed Low Ambient Temperature without a Thermal Module Fan Thermistor
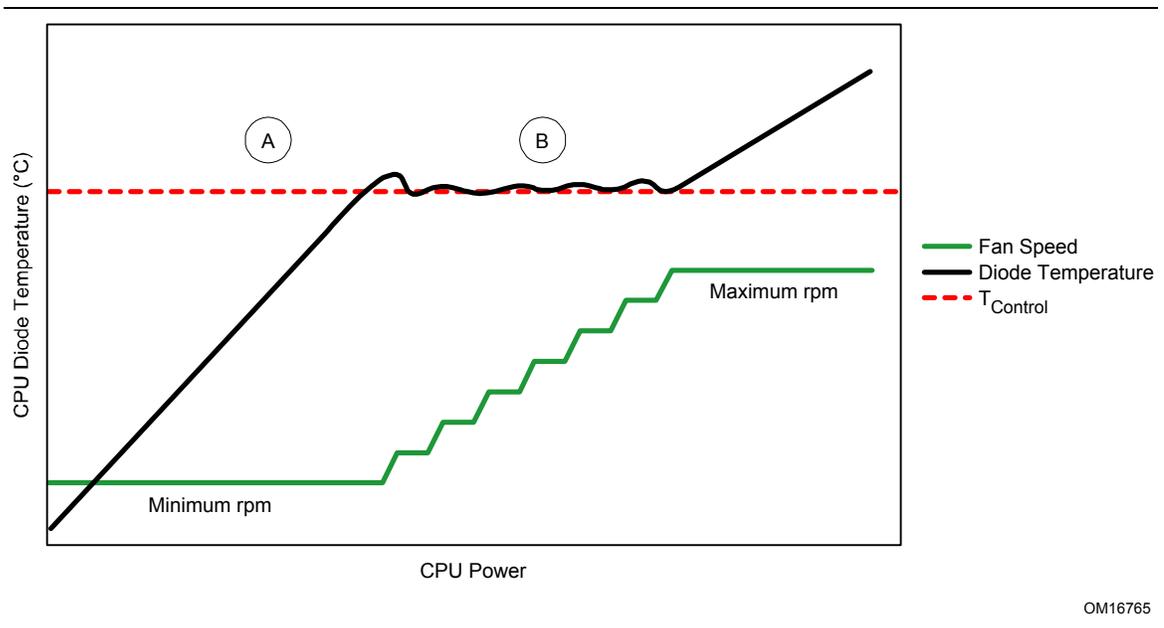
OM16765

Figure 129: CPU Diode Temperature and Thermal Module Fan Speed as a Function of CPU Power in Fixed Low Ambient Temperature with a Thermal Module Fan Thermistor
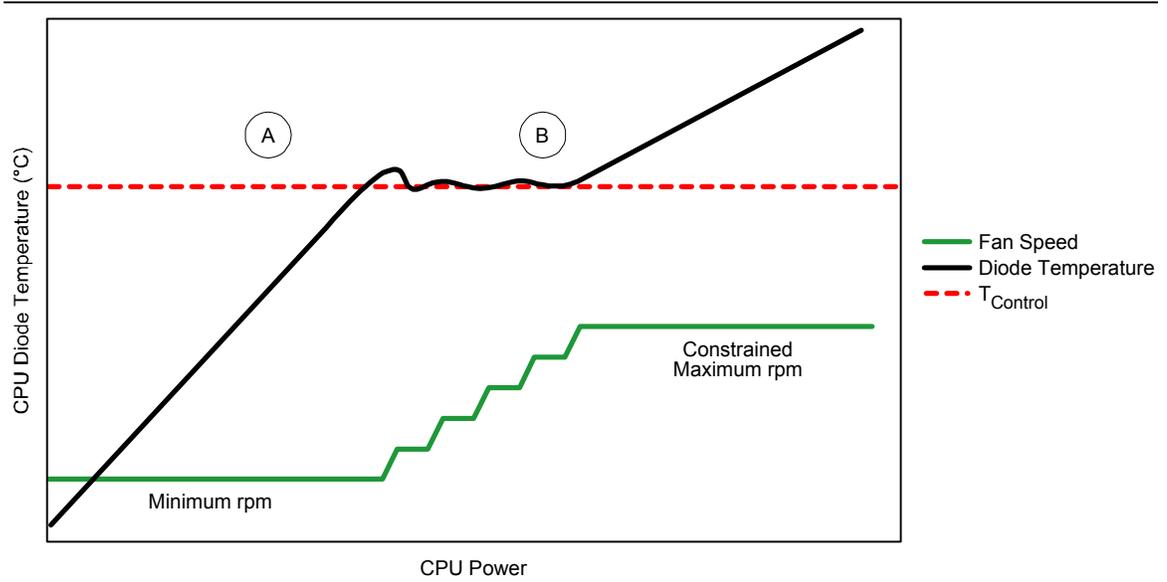
OM16766

Figure 130: CPU Case Temperature as a Function of CPU Power in Fixed Low Ambient Temperature With and Without a Thermal Module Fan Thermistor
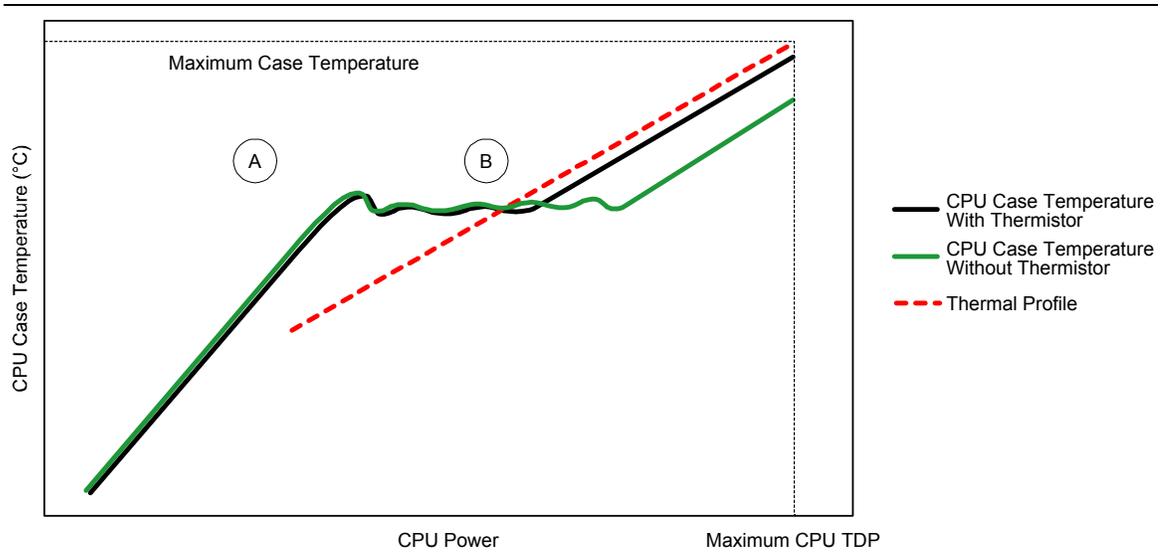
OM16767

Finally, Figure 131 illustrates the difference in system Sound Power when a Thermal Module fan thermistor is integrated into the FSC circuit. The illustration shows Sound Power as a function of CPU power in a low external ambient.

Figure 131: Sound Power as a Function of CPU Power in a Fixed Low Ambient Temperature With and Without Thermal Module Fan Thermistor
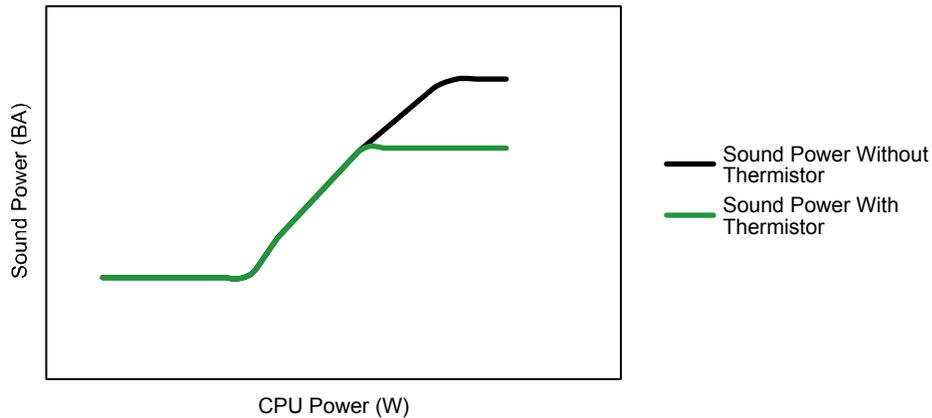
Fundamentally, the difference in fan speed and Sound Power response for these two FSC schemes can be explained as follows. When the CPU power has increased to the point where the CPU diode temperature is above its limit, the FSC circuit will repeatedly request that Thermal Module fan speed to increment. In the absence of any additional constraining regulation, the fan will eventually reach its maximum operating speed.

Notice in Figure 130 that, in the absence of Thermal Module fan thermistor augmentation to the FSC scheme, the CPU case temperature at CPU TDP will be many degrees below the CPU Thermal Profile and well below the maximum case temperature specification when CPU power is at TDP. The implication is that the Thermal Module is providing better heat transfer performance than is required because the fan is operating at a faster speed than needed.

When the Thermal Module fan thermistor's knowledge of the ambient temperature is introduced to the FSC scheme, that information can be used to constrain the speed at which the fan operates (termed the ad hoc maximum operating speed).

The FSC scheme will compare the actual ambient temperature reported by the thermistor to the maximum ambient temperature for which the system is designed and constrain the ad hoc maximum operating speed. This assures that the acoustic behavior of the system is optimized for a broader range of operating conditions.

This illustration assumes that the FSC scheme does not have knowledge of the CPU power. Therefore, the Thermal Module ad hoc maximum operating speed will be reached when there is a sustained reporting of CPU diode temperature greater than its limit. The ad hoc maximum operating speed must be set at a speed that ensures that the CPU maximum case temperature will be met at CPU TDP.

Section 3.4.4.2 outlines the selection of the System Monitor Point location. The temperature limit for this FSC input will determine the system's acoustic behavior in a broad range of operating conditions.

If the System Monitor Point temperature limit is set to a low value, then there will be a greater proportion of typical operating conditions where the Thermal Module fan will operate above its minimum operating speed. Whenever the fan speed is above its minimum operating speed, the system noise will increase. The higher this limit is set, the closer the component temperatures will be to their specifications but there will be a broader range of operating conditions in which the fan will be at its minimum operating speed.

The Thermal Module fan speed will be regulated based on both the CPU diode and the System Monitor point. In addition to describing the system responses if the CPU diode is used in FSC circuit input, Figure 126, Figure 128, and Figure 129 also describe the behavior of the System Monitor Point temperature when the power of nearby components (for instance, the MCH and graphics add-in card) increases. The Thermal Module fan speed and Sound Power behaviors illustrated in Figure 131 are identical when the System Monitor Point temperature instead of CPU diode temperature is the driving input to the FSC scheme.

3.4.4.4 Fan Speed Control Programming

To achieve improved acoustics, a BTX system must use a fan speed control IC. Typically this IC is integrated on the motherboard. The IC must be configured to control the fans to meet thermal requirements as well as to provide improved acoustics. The motherboard layout and schematic and the bios code depend on the fan speed control IC implementation.

3.4.4.4.1 Component Selection

In order to select a fan speed control IC for use in a BTX system, the required capabilities should be understood. The number and type of temperature monitor inputs, the number and type of fan speed control outputs, and automatic fan speed control capability are some of the important features.

The BTX Reference Design Systems use two diode-based temperature monitor inputs and one 25 kHz 4-wire fan PWM output. The temperature monitor points are used to sense the processor diode and the system monitor point. The 4-pin PWM-controlled fan is integrated into the Thermal Module. After it is configured, automatic fan speed control mode should be used. The programmed algorithm runs within the fan speed control IC.

The following (Heceta 6e) fan control ICs offer two external temperature sensor inputs, three 25 kHz PWM fan control outputs, and automatic fan speed control mode: Analog Devices ADT7474, National Semiconductor LM96000, and SMSC EMC6D103. The fan speed control IC used in the BTX Reference Design Systems is the SMSC EMC6D103.

3.4.4.4.2 Schematic

Figure 132 shows a schematic diagram of the components used in the BTX Reference Design fan speed control system. Figure 133 and Figure 134 show the tachometer and processor diode sense schematics, respectively.

Figure 132: BTX Reference Design System FSC Schematic
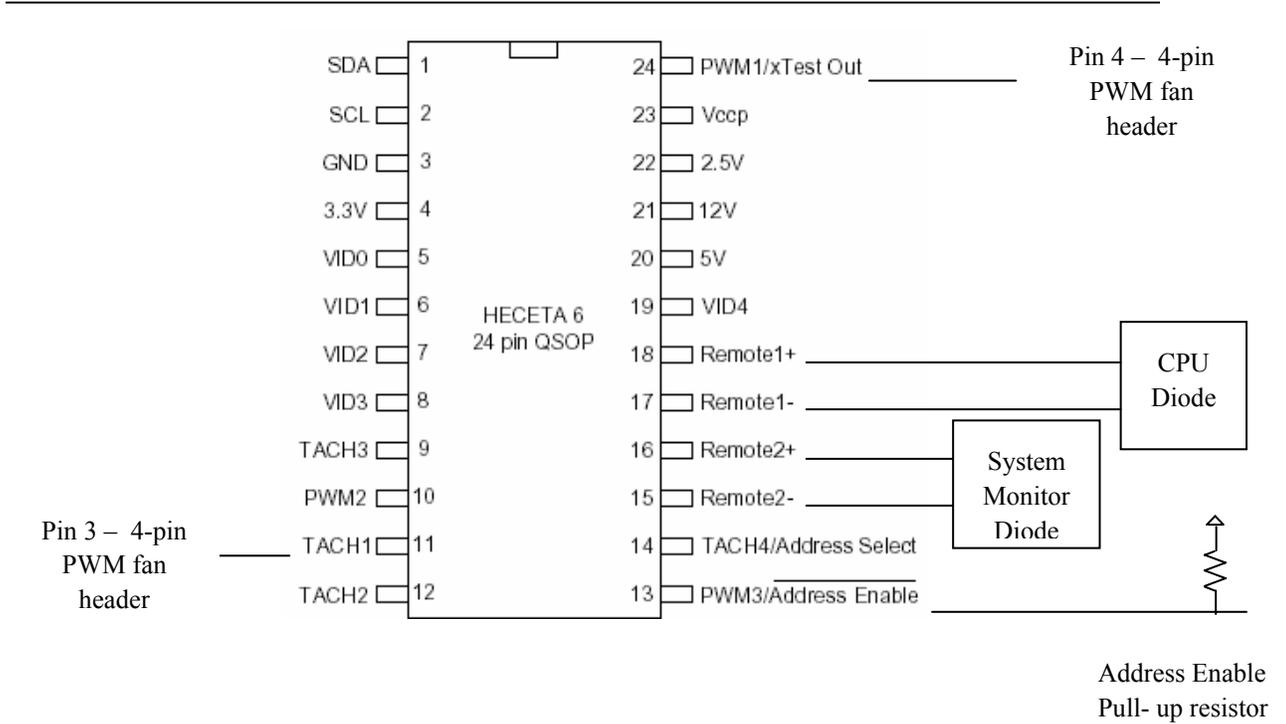
Figure 133: Tachometer Interface Circuit
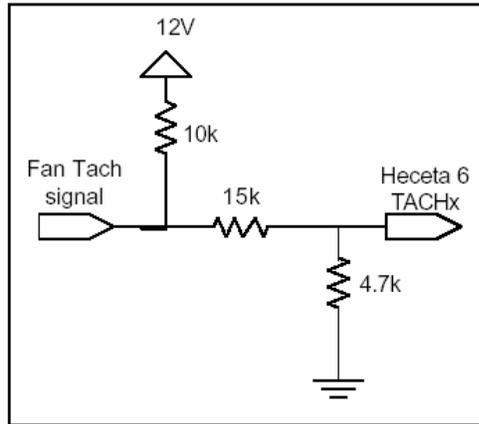
Figure 134: Diode Temperature Sensor Interface Circuit
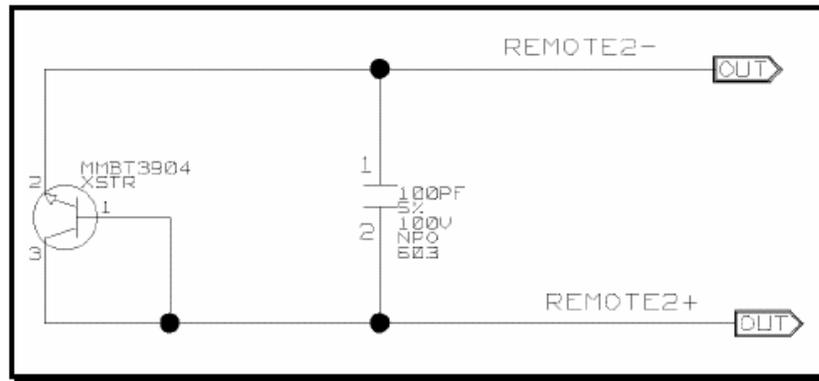
3.4.4.4.3   System Thermal Evaluation (Numerical and Empirical descriptions)

System level thermal numerical analysis was used to determine the optimal location for the system monitor point and also to predict its required temperature trip (limit) value. During system thermal investigation, which was conducted prior to system thermal validation, the temperature behavior at the system monitor point was measured to ensure that it was correlated with the numerical model predictions. The numerical model simulations and investigative thermal testing were conducted for a number of system power loads, ambient, and fan speed combinations. From these evaluations, temperature trip (limit) values for the system monitor point can be determined that ensure that the system fans operate in a manner that ensures that all subsystem temperature requirements are met across a wide range of operating conditions.

3.4.4.4.4   Heceta-6e BIOS Code

The motherboard BIOS should be programmed to initialize and set the fan control IC. The program and IC settings will likely be different for the different fan control ICs. The configuration of the fan control IC is typically done over the SMBus or LPC Bus during POST. The Heceta 6e is an SMBus device and uses a set of read/write registers for configuring the device.

3.4.4.4.4.1 Heceta initialization on SMBus

The Heceta-6e has an option for 3 different SMBus addresses. This requires the use of the PWM/Address Enable pin as well as the Address Select pin. Please refer to the Heceta-6e component specification for details. On the BTX Reference Design RVP, a pull-up resistor was connected to the PWM/Address Enable pin to configure the component for SMBus address 0101110b (5Ch). During boot up, the Heceta 6e starts in address enable mode. After the first valid SMBus transaction occurs (one that starts with the first five bits: 01011b), the Heceta 6e will latch its address.

3.4.4.4.4.2 Minimum fan speed program

The following assumes the use of an Intel Pentium 4 processor with a diode temperature sensor and support for Tcontrol. Refer to the processor Datasheet for details on Tcontrol.

Bits [15:8] in the 01A2h MSR register in the processor are used to store the Tcontrol offset value. This offset value should be read and added to the appropriate Tcontrol base value by the BIOS. The total of the offset and base is the Tcontrol value for the specific processor in use.

The Heceta 6e has register settings for Tmin and Trange. Tcontrol is equivalent to the Tmax for the processor diode temperature. Tmax can be computed by adding Tmin and Trange. In order to program the

Tmin value for the Heceta 6e, Tmin must be computed. Tmin is determined by subtracting Trange from Tcontrol (Tmin =Tcontrol – Trange). The BIOS needs determine the correct Tmin value based on the Tcontrol offset value from the MSR register, the Tcontrol base value, and the Trange value that is also programmed into the Heceta 6e. See the illustration in Figure 135

Figure 135: Tcontrol Heceta 6e Register Description
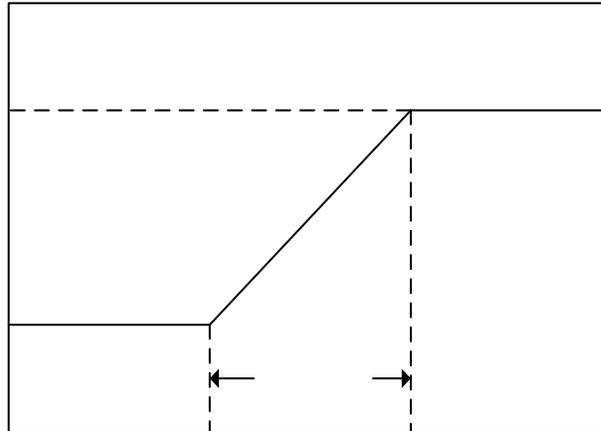
3.4.4.4.5   Heceta-6e Settings

The fan control IC uses a register set to govern its operation. When initially powered up, these register settings are in their default settings and must be configured by the BIOS to make the IC operate as required. Refer to the fan control IC component specification for a complete register set description. The registers corresponding to each function are listed below.

The BIOS should follow the steps outlined below to configure the fan registers on the Heceta 6e. All steps may not be necessary if default values are acceptable. Regardless of changes made by the BIOS to the limit and parameter registers during configuration, the Heceta 6e must continue to operate based on default values until the Start bit, in the Start Monitoring register, is set. Once a fan is put into auto fan mode, by settings its configuration register and setting the start bit to 1, the Heceta 6e will operate per the values set by BIOS in the limit and parameter registers.

- Set limits and parameters

    [5C-5Eh] Set the configuration register.

    [5C-5Eh] Set the fan spin-up delays.

    [5F-61h] Set PWM frequencies and auto fan control range.

    [62-63h] Set fan ramp rate and min/off

    [64-66h] Set the PWM minimum duty cycle.

    [67-69h] Set the fan temperature limits.

    [6A-6Ch] Set the temperature absolute limits.

- 2) [40h] Set bit 0 (Start) to start monitoring.

- 3) [40h] Set bit 1 (Lock) to lock the limit and parameter registers (optional)

Table 26 outlines the settings used in the BTX Reference Design Systems.

Table 26: BTX Reference Design System Heceta 6e Register Settings

Register Address	Register Name	Bit 7 (MSb)	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0 (LSb)	Default Value	BTX S1 Setting
40h	Configuration Register 1	Vcc	TODIS	FSPDIS	V X I	Override	Ready	Lock	Start	00h	01h
5Ch	Fan1 Configuration	Zone bit 2	Zone bit 1	Zone bit 0	PWM Invert	RESERVED	Spin bit 2	Spin bit 1	Spin bit 0	62h	C7h
5Dh	Fan2 Configuration	Zone bit 2	Zone bit 1	Zone bit 0	PWM Invert	RESERVED	Spin bit 2	Spin bit 1	Spin bit 0	62h	C7h
5Eh	Fan3 Configuration	Zone bit 2	Zone bit 1	Zone bit 0	PWM Invert	RESERVED	Spin bit 2	Spin bit 1	Spin bit 0	62h	C7h
5Fh	Zone1 Range/ Fan1 Freq	Range bit 3	Range bit 2	Range bit 1	Range bit 0	Freq bit 3	Freq bit 2	Freq bit 1	Freq bit 0	C3h	4Ah
60h	Zone2 Range/ Fan2 Freq	Range bit 3	Range bit 2	Range bit 1	Range bit 0	Freq bit 3	Freq bit 2	Freq bit 1	Freq bit 0	C3h	4Ah
61h	Zone3 Range/ Fan3 Freq	Range bit 3	Range bit 2	Range bit 1	Range bit 0	Freq bit 3	Freq bit 2	Freq bit 1	Freq bit 0	C3h	4Ah
62h	Min/Off, Fan1 Ramp Rate	Off bit 2	Off bit 1	Off bit 0	RESERVED	F1 RR Enable	F1 RR bit 2	F1 RR bit 1	F1 RR bit 0	00h	E8h
63h	Fan2, Fan3 Ramp Rate	F3 RR Enable	F2 RR bit 2	F2 RR bit 1	F2 RR bit 0	F3 RR Enable	F3 RR bit 2	F3 RR bit 1	F3 RR bit 0	00h	88h
64h	Fan1 PWM Minimum	bit 7	bit 6	bit 5	bit 4	bit 3	bit 2	bit 1	bit 0	80h	2Dh
65h	Fan2 PWM Minimum	bit 7	bit 6	bit 5	bit 4	bit 3	bit 2	bit 1	bit 0	80h	2Dh
66h	Fan3 PWM Minimum	bit 7	bit 6	bit 5	bit 4	bit 3	bit 2	bit 1	bit 0	80h	2Dh
67h	Zone1 Fan Temp Limit	bit 7	bit 6	bit 5	bit 4	bit 3	bit 2	bit 1	bit 0	5Ah	XXh*
68h	Zone2 Fan Temp Limit	bit 7	bit 6	bit 5	bit 4	bit 3	bit 2	bit 1	bit 0	5Ah	78h
69h	Zone3 Fan Temp Limit	bit 7	bit 6	bit 5	bit 4	bit 3	bit 2	bit 1	bit 0	5Ah	30h
6Ah	Zone1 Temp Absolute Limit	bit 7	bit 6	bit 5	bit 4	bit 3	bit 2	bit 1	bit 0	64h	64h
6Bh	Zone2 Temp Absolute Limit	bit 7	bit 6	bit 5	bit 4	bit 3	bit 2	bit 1	bit 0	64h	7Fh
6Ch	Zone3 Temp Absolute Limit	bit 7	bit 6	bit 5	bit 4	bit 3	bit 2	bit 1	bit 0	64h	4Bh

* The Zone 1 Fan Temperature Limit (Tmin) setting depends on Tcontrol. (Tmin = Tcontrol-Trange) (Tcontrol = Tcontrol base + Tcontrol offset)

Refer to the SMSC EMC6D103 specification for details on register settings. The following is an illustration of how to program a setting:

- Perform an SMBus write command to SMBus address 5Ch, register 5Fh, value 4Ah (01001010b)

- This command will set the following settings:

  o Trange setting of 0100b = 5 degrees C

  o Frequency setting of 1010b = 25 kHz

- Other register settings can be set similarly.

3.4.4.4.5.1 Automatic mode algorithm

Figure 136: Automatic Mode Algorithm Description
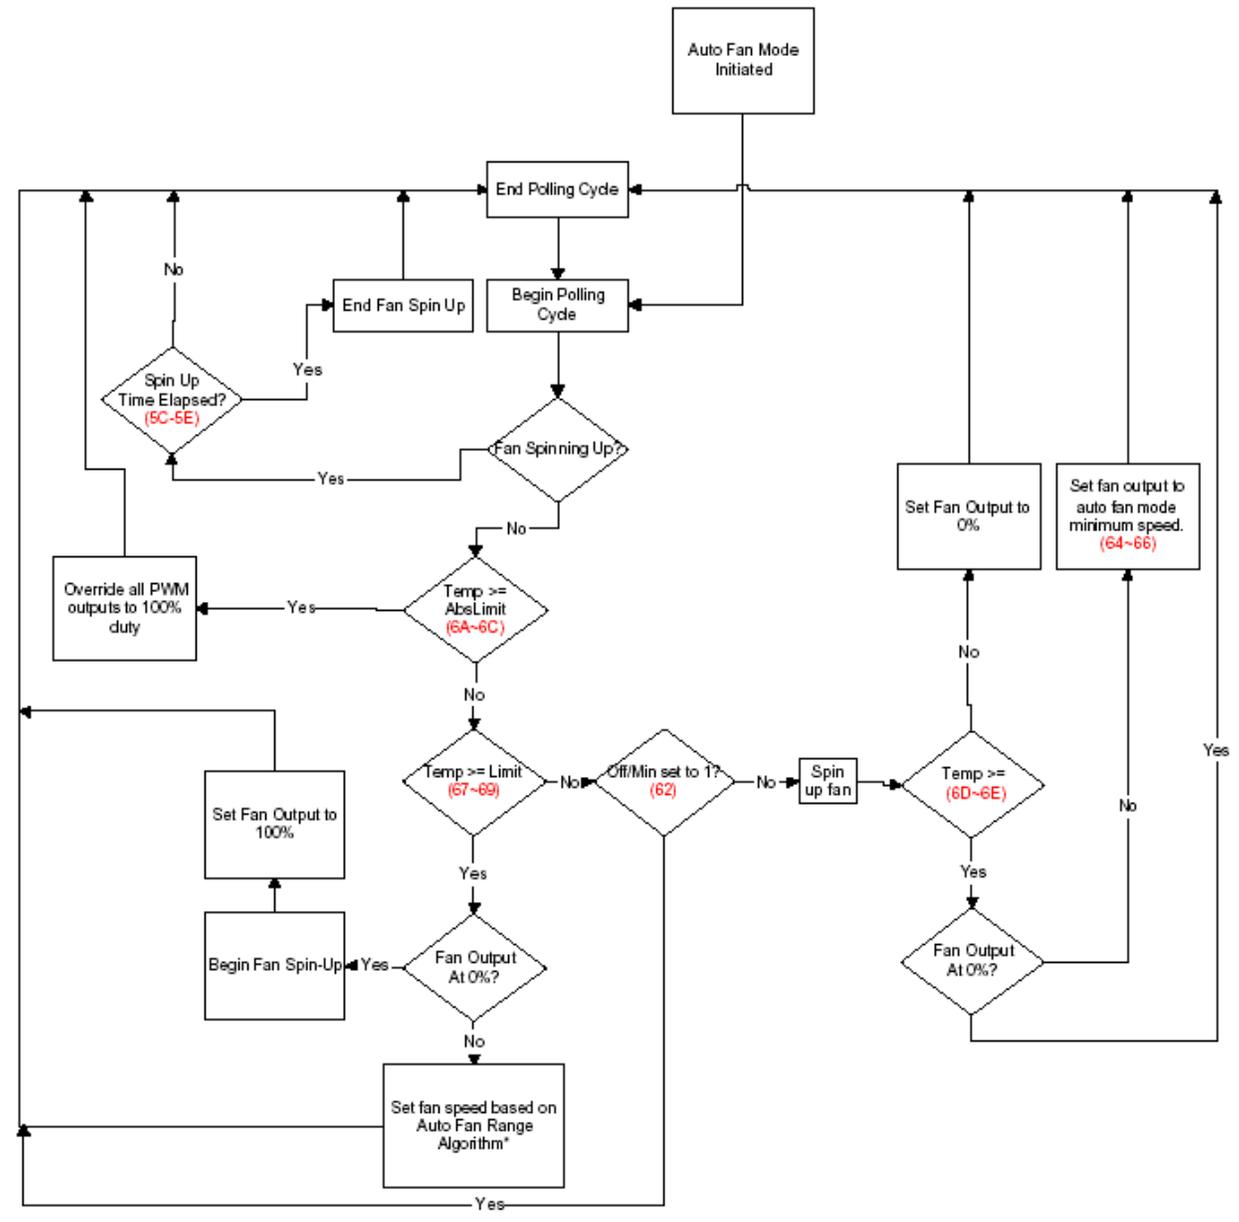
3.4.5 Case Study – Intel Reference Design Acoustic Engineering and Performance

3.4.5.1 Noise Source Summation

The 6.9 liter and 12.9 liter Intel desktop reference designs use only two fans – a Thermal Module fan and a PSU fan – thereby, minimizing the number of noise sources. Intel system level thermal CFD analysis has been conducted on a wide variety of system sizes and system profiles; in all cases it has been demonstrated that adequate system cooling is available for all subsystems with the use of only two fans. In fact, the

Type I Intel reference design Thermal Module, using the Intel reference design 90-mm fan and stator, provides sufficient system airflow in system sizes up to a standard full-tower.

The 12.9 liter reference design uses Intel patented impeller and stator technology applied to a 90-mm Type I Thermal Module fan, which delivers a substantial fan curve improvement at all operating speeds in the region of typical BTX system impedance. For a given airflow requirement and system impedance curve, the improved fan curve allows the fan to operate at a lower RPM.

The Intel reference design Type I Thermal Module 90-mm fan operates at 4500 RPM at its maximum operating speed in free stream and approximately 4300 RPM when installed in the Intel 12.9 liter reference design. The FSC circuit allows the fan a minimum operating speed of 900 RPM, which provides sufficient airflow to all subsystems in the idle operating condition. The Type I Intel reference design Thermal Module measured Sound Power as a function of fan speed is shown in Figure 137.

Figure 137: Intel Type I Thermal Module Sound Power as a Function of Fan Speed

BA = -8.4 + 3.9*LOG (rpm)

OM17094
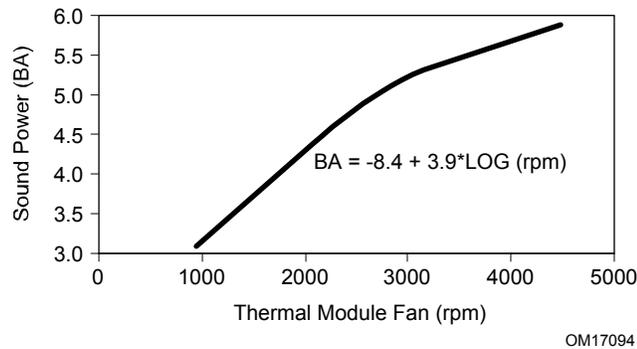
This same impeller and stator design technology was applied to the 70-mm Type II Thermal Module fan used in the 6.9 liter reference design. Its maximum operating speed is 6800 RPM free stream and its minimum speed is 1500 RPM, in the 6.9 liter Intel reference design. The Type II Intel reference design Thermal Module measured Sound Power as a function of fan speed is shown in Figure 138.

Figure 138: Intel® Type II Thermal Module Sound Power as a Function of Fan Speed

BA = -6.6 + 3.3*LOG (rpm)

OM17095
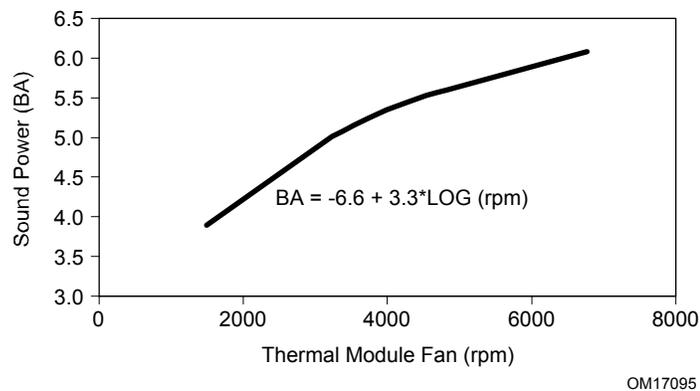
Intel also integrated a MiTAC CFX12V PSU into the 12.9 liter reference design system. The CFX12V reference design PSU has minimum and maximum operating speeds of 1550 and 3600 RPM, respectively. A HiPro LFX12V PSU was integrated into the 6.9 liter reference design, with fan minimum and maximum

operating speeds of 1600 and 3825 RPM, respectively. The measured Sound Power as a function of fan speed for both reference design power supplies are shown in Figure 139 and Figure 140.

Figure 139: CFX12V Reference Design PSU Sound Power as a Function of Fan Speed

BA = -18 + 6.64*LOG (rpm)

OM17096
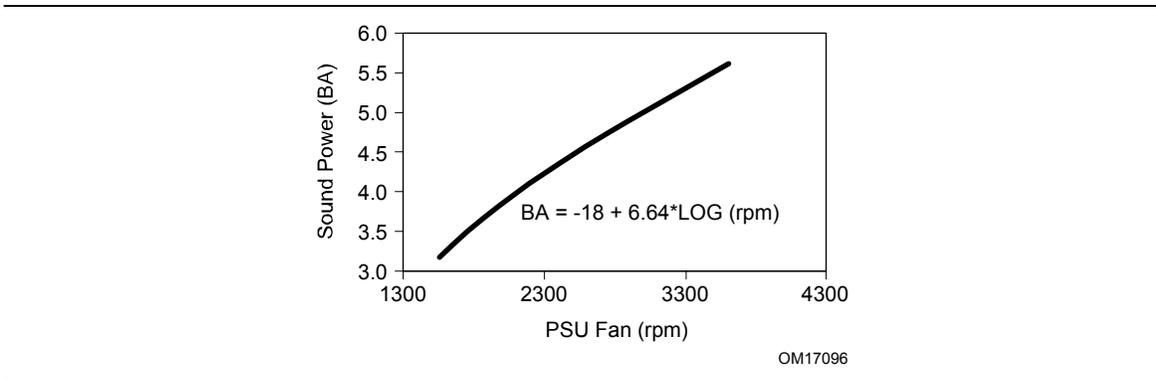
Figure 140: LFX12V Reference Design PSU Sound Power as a Function of Fan Speed

BA = -15.5 + 5.88*LOG (rpm)

OM17097
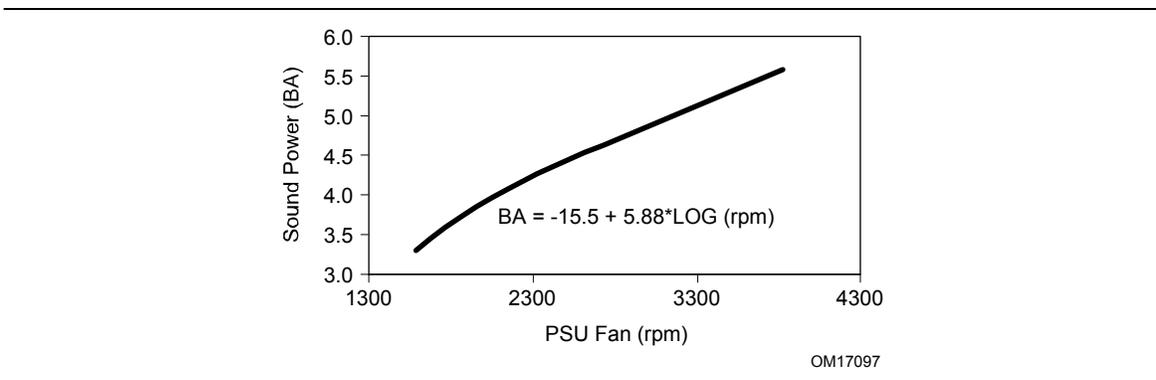
Intel also measured the Sound Power of the HDD installed into the reference designs. In its idle operating condition (spinning bit not accessing), its measured sound power is 3.2 BA. With all the individual noise sources in the system characterized, Equation 47 can be used to predict the Acoustic Typical (e.g., idle operating condition) Sound Power for the Intel system reference designs is shown in Table 27.

Table 27: Source Sound Power Estimates

	Thermal Module BA	PSU BA	HDD BA	System BA
12.9 liter reference design	3.5	3.2	3.2	3.8
6.9 liter reference design	3.6	3.3	3.2	3.9

3.4.5.2 Bezel Airflow Impedance Impact

Please refer to Section 2.7.3.1.2, Table 10 for the acoustic impact of increasing the front bezel or Thermal Module interface ventilation impedance.

4. Structural Design

4.1 Definition of Structural Terms

Term	Symbol	Definition
Support and Retention Module	SRM	Structural enabling component intended to distribute processor heatsink inertial loads to the edges of the chassis.
Ball Grid Array	BGA	
Preload		
System response		
Natural frequency		
Material Yield Strength	FtY	
Material Ultimate Strength	FtU	
Material Ultimate Strain		
Material Elongation Limit		Same as Material Ultimate Strain
Cross-Sectional Moment of Inertia	I	Area moment of inertia through the cross-section of a given structural component

4.2 Principles of Platform Structural Design

Platform structural design is centered on management of inertial loads under shock and vibration conditions. It is under these load conditions that the motherboard and its components are most susceptible to damage. Inertial loads will develop from massive components mounted to the motherboard (such as the CPU heatsink) or the chassis (such as the PSU). When the system is exposed to an acceleration profile associated with a shock or vibration event, those massive components exert inertial loads onto the motherboard and surrounding support structure.

Platform failure can occur if inertial loads are not managed properly. The inertial loads will cause the motherboard to flex, thereby inducing strain in the motherboard and the board-mounted components. Excessive board flexure can cause damage to such components. Of specific concern are the surface mounted components, including the socket and MCH. Special care must be taken to ensure adequate protection for these components whenever the motherboard is mechanically loaded.

Several structural engineering strategies are available to help manage the inertial load conditions cited above. Some strategies are obvious and straightforward such as limiting mass levels of board-mounted component or reducing shock acceleration profiles. Others are more complicated and require a detailed explanation. For example, preloading, use of alternate load paths, and topside stiffening are three effective techniques for managing board flexure under inertial load conditions. Each of these strategies will be described later in this section.

4.3 Typical Load Conditions

Several load conditions must be considered in platform structural design. These include handling loads, assembly loads, and shipping-induced shock and vibration loads. In addition, the clip preload (used to hold the heatsink tight against the top of the processor) may also need to be accounted for as a relevant design load condition.

The predominant loads that drive platform structural design are the shock and vibration loads. These load conditions are derived from shipping environments in which package drops generate shock loads and shipping environment generates vibration loads. These two load conditions are generally far more severe than handling or assembly loads, and therefore drive the majority of the structural design.

4.3.1 Shock Load Conditions

The shock load condition is typically represented by an acceleration impulse applied to the system to represent a package drop in a shipping environment. The acceleration impulse may reach levels of 25 g's to 30 g's with duration on the order of 10 ms. (Actual system shock requirements will vary and depend on the actual shipping environment or the system drop validation requirements and the level of protection provided by system packaging.) The response of the system under this load condition will generate inertial loads on all massive system components.

Massive components in the system may include the processor fan/heatsink assembly, other component heatsinks (MCH, ICH, etc.), the motherboard, add-in cards, drives, power supply, etc. Each of these components will develop an inertial load under the shock condition. Those inertial loads must be managed by the system structural design. The loads are managed by designing structural members to carry the loads from their points of origin out to the edges or corners of the system (where the system is typically supported by its packaging and is most stiff).

Of particular interest is the role the motherboard plays in distribution of inertial loads. Traditionally, the motherboard has been expected to bridge the inertial load from the processor heatsink to the board mounting points where the load is reacted into the chassis. As heatsink mass levels increase (typically with increasing power dissipation requirements), the heatsink inertial loads eventually exceed the load carrying capability of the motherboard.

4.3.1.1 Calculating Inertial Loads

Inertial loads are dependent on three primary factors: mass, acceleration, and DAF (dynamic amplification factor). The equation is given by:

Equation 49: Inertial Load

$P = maf$

Where:

P = peak inertial load

m = component mass

a = peak input acceleration

f = DAF

The DAF is defined as the ratio of the peak acceleration of the component to the peak acceleration of the input. The DAF is dependent on the system response and the location of the mass in the system. The response of the entire system can be analyzed through numerical methods. However, many of the system components can be represented reasonably well with simpler spring-mass models.

4.3.1.1.1 Single Degree-of-Freedom System

A common representation of a board-level shock condition is given by single DOF (degree of freedom) spring-mass system (see Figure 141). In this case, the mass represents the heatsink, the spring represents the board bending stiffness, and the ground point represents the motherboard support points. The acceleration input is imparted to the motherboard supports. The acceleration is transferred to the heatsink through the motherboard. And the heatsink responds with an acceleration profile of its own. The response of the heatsink can be calculated with using the spring-mass model.

Figure 141: Spring-Mass Representation for Single DOF System
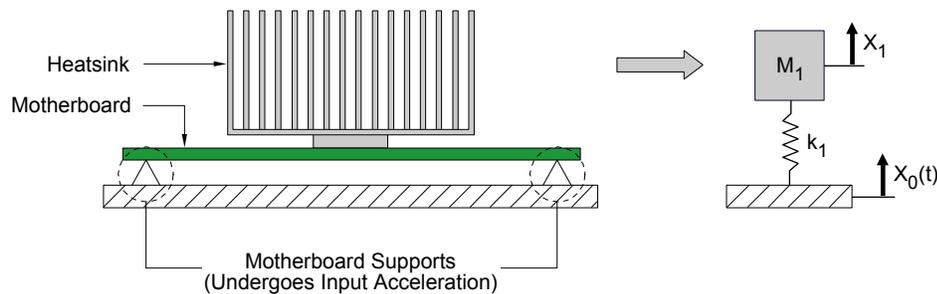
The governing equation for this system is given by:

Equation 50: Spring Mass for Single DOF

$$m_1 \ddot{x}_1 + k_1 x_1 = k_1 x_0(t)$$

Where:

x1 = deflection of the heatsink

x0(t) = movement of the support as a function of time (the location of the input acceleration)

k1 = stiffness of motherboard

m1 = mass of the heatsink

For a given acceleration input, the equation above can be solved to obtain the response of the heatsink. For a typical motherboard shock profile (square wave or trapezoidal wave), the maximum acceleration seen by the heatsink will range between 0 and 2 times the input acceleration, thereby yielding a DAF between 0 and 2. The actual DAF is dependent on the duration of the impulse. If the acceleration impulse duration is brief relative to the natural frequency of the system, then the DAF will tend toward 0. If the duration is long, then the DAF will reach 2. Figure 142 shows the relation between DAF and impulse duration for a single DOF mass-spring system. Many platforms in the industry today will generate a DAF close to 2 based on typical values for motherboard stiffness and processor heatsink mass.

Figure 142: Dynamic Amplification Factor
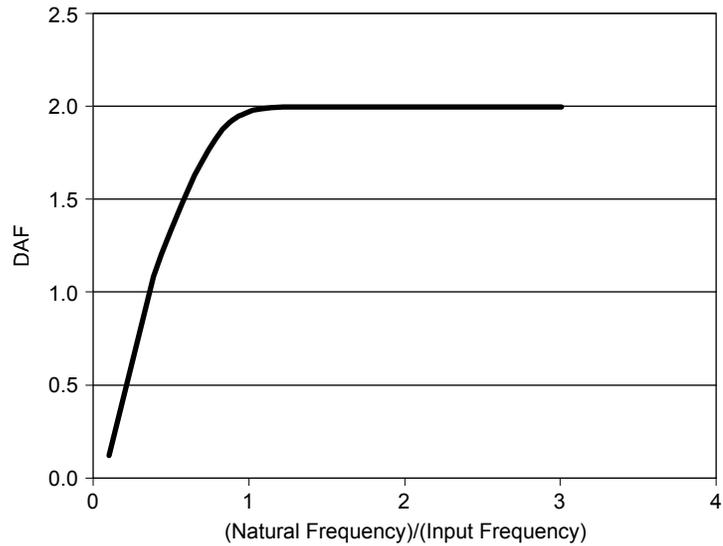
OM16770

4.3.1.1.2 Two Degree-of-Freedom System

For more complicated systems, a single DOF spring-mass model is not sufficient to predict the component inertial load response. For example, when the motherboard is installed in a system, the board is mounted to the chassis pan, which is itself a spring. This spring element must now be considered in the mass-spring model. In addition, the motherboard mass becomes significant, as its own inertial load will contribute substantially to the chassis pan deflection. As such, a second mass element must be added to the model. This results in a two-DOF spring-mass model representation of a system-level shock condition (see Figure 143).

Figure 143: Spring-Mass Representation for Two-DOF System
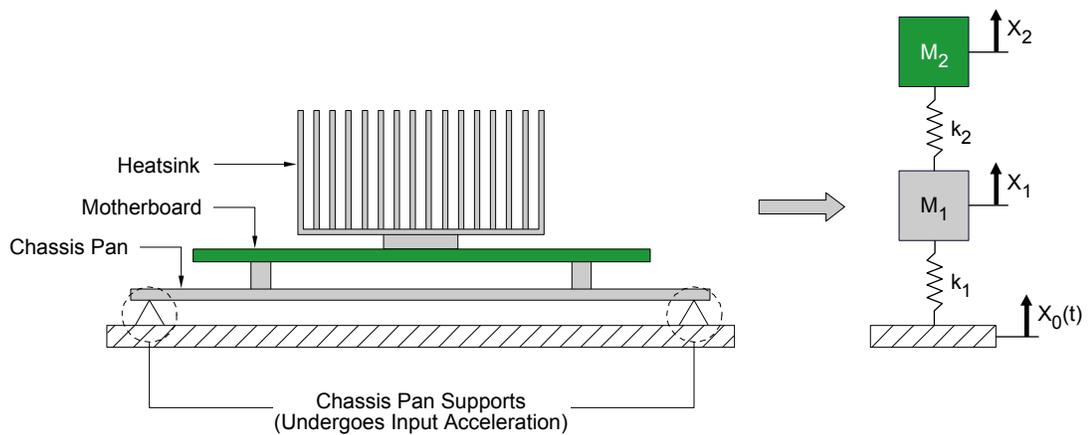
OM16771

The governing equation for this system is given by:

Equation 51: Spring Mass for Two DOF

$$\begin{bmatrix} m_1 & 0 \\ 0 & m_2 \end{bmatrix} \begin{Bmatrix} \ddot{x}_1 \\ \ddot{x}_2 \end{Bmatrix} + \begin{bmatrix} k_1 + k_2 & -k_2 \\ -k_2 & k_2 \end{bmatrix} \begin{Bmatrix} x_1 \\ x_2 \end{Bmatrix} = \begin{Bmatrix} k_1 x_0(t) \\ 0 \end{Bmatrix}$$

The equation can be generalized for multi-DOF systems in matrix form as,

Equation 52: Matrix of Spring Mass for Two DOF

$$[m]\{\ddot{x}\} + [k]\{x\} = \{F(t)\}$$

For a given acceleration input, the equation above can be solved to obtain the response of the heatsink. Again, the resulting DAF is dependent on the duration of the impulse relative to the fundamental frequencies of the system. For a square wave acceleration profile, the DAF will range between 0 and 4. The theoretical maximum of 4 cannot be reached without the proper ratio of natural frequencies between the two subsystems. Many platforms in the industry today do not have the stiffness and mass ratios needed to approach this maximum value. DAF's close to 3 are much more common based on typical values for motherboard stiffness, chassis pan stiffness, and component mass levels.

4.3.1.1.3   Complex System Modeling

Systems that have even more complexity in their structural layout may not lend themselves to simple representation through spring-mass models. This added complexity might come from inclusion of additional masses attached to the chassis pan (such as drives, PSU, etc.), additional masses attached to the motherboard (such as add-in cards), or a more complicated connectivity between the heatsink, board and chassis pan (such as when using an SRM).

In these cases, finite element analysis can be used to determine dynamic amplification factors. It is important to identify and represent all significant mass elements in the model. The more detailed the mass distribution, the more accurate the model prediction will be. It is equally important to represent all stiffness elements in the load path. Again, the more precise the representation of the element stiffness, the more accurate the model prediction will be.

A transient finite element analysis must be run to capture the DAF. This analysis will require input of the acceleration profile. This is typically a trapezoidal wave or sine wave. It can be applied either as a time-varying gravity profile, or as an acceleration profile applied to the system support points.

As with any finite element analysis, special care must be taken to ensure that the boundary conditions are modeled appropriately. This applies to system support boundary conditions as well as component-to-component interface conditions. Unrepresentative boundary or interface conditions can substantially alter the effective stiffness of a component or the system as a whole. This will yield inaccurate results for all analysis output including displacements, stresses and acceleration factors.

4.3.1.2 Applying Inertial Loads

Once the inertial loads have been determined for each of the massive components of interest, they can be applied in the system to evaluate the stress levels and deflections of system components.

4.4 Failure Modes and Methods of Evaluation

Typical failure modes associated with the shock condition fall into two categories: deflection-related failures, and stress-related failures. These failure modes will be described in the following sections.

4.4.1 Deflection Failure Modes

Deflection-related failures can occur when components deflect to the point of contacting nearby components or disengaging from their attachment mechanisms.

When massive components deflect under their own inertial load to the point of contacting nearby components an impact load is generated between the two components. If the load is large enough, or if the components are fragile, damage to either component can occur.

The MCH and other bare-die components may be susceptible to this failure mode. Bare die components, are not designed to sustain substantial impact loads. Therefore, component impact must be avoided through system structural design. Deflections of the motherboard and surrounding components must be controlled to prevent over-deflection and impact during shock conditions.

Even those bare-die components covered with heatsinks can be susceptible depending on the heatsink attach mechanism. Heatsinks with clip attach mechanisms will transfer impact loads directly to the die, and therefore do not offer substantial protection to the component. In these cases, deflections must be controlled to prevent system components from impacting the top or side of the MCH heatsink.

Strategies for managing system component deflections are given in Section 4.5.

Disengagement is a second failure mode associated with over-deflection. System components such as brackets or clips that rely on tab-and-slot or snap-on attach mechanisms may be subject to disengagement during shock conditions. If the attach mechanism is deflected far enough, the mating features may be spread apart or opened, thereby allowing disengagement. Attach mechanisms should be sized according to anticipated loading to prevent such disengagement.

4.4.1.1 Calculating Component Deflections

Component deflections are dependent on three primary factors: inertial loading, stiffness of the structural elements connecting the component to the chassis ground points, and support conditions of the structural elements.

4.4.1.1.1 Beam Representation of Structural Components

A common representation for certain brackets and clips is achieved with a uniform cross-section beam. This simplified representation offers a quick and effective method of evaluating component deflections. Beam bending relationships are well established through beam bending theory. The basic load-deflection relationship is provided below for an example set of boundary conditions. For a more thorough presentation of beam bending formulas and theory, reference *Roark's Formulas for Stress and Strain*.

The example below shows how a simple beam can be used to represent a heatsink clip (see Figure 144). In this case the clip is attached to the motherboard at its two ends with pinned connections and the heatsink is attached to the center of the clip. The heatsink applies inertial load to the clip in a shock condition. The load-deflection relationship for the beam is given by:

Equation 53: Beam Load Deflection

$$\Delta = \frac{PL^3}{48EI}$$

Where:

Δ = Deflection at center of beam

P = Heatsink inertial load

L = Distance between supports

E = Material modulus

I = Cross-sectional moment of inertia

Figure 144: Beam Representation of Heatsink Clip in Bending
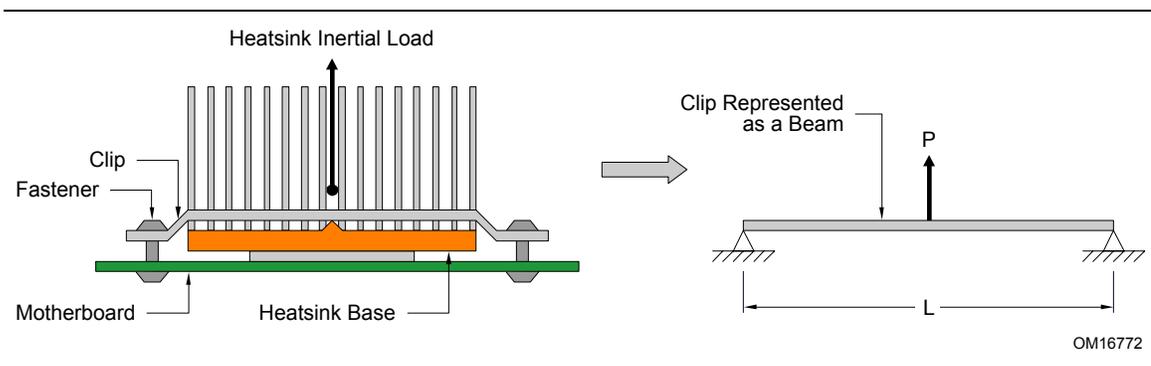
If instead, the clip is fastened to the motherboard in such a way as to restrain rotation at its ends, then the beam representation changes to clamped end supports (see Figure 145). This drastically changes the effective stiffness of the beam as seen from the new load-deflection equation shown below. The deflection drops by a factor of two as a result of the boundary condition change.

Equation 54: Load-Deflection Equation

$$\Delta = \frac{PL^3}{96EI}$$

Figure 145: Beam Representation of Heatsink Clip with Alternate Boundary Conditions
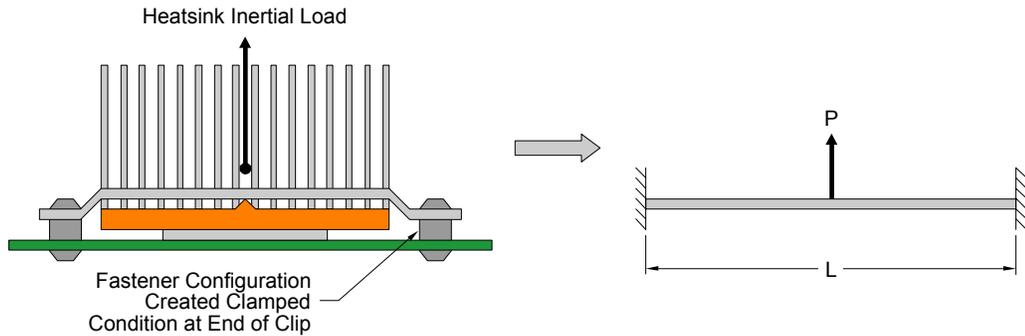
4.4.1.1.2 Plate Representation of Structural Components

A uniform thickness plate can be used to represent certain platform structural components such as the chassis pan and motherboard. Again, this simplified representation offers a quick and effective method of evaluating component deflections. Plate bending relationships are well established through plate bending theory. The basic load-deflection relationship is provided below for an example set of boundary conditions. For a more thorough presentation of plate bending formulas and theory, reference *Roark's Formulas for Stress and Strain.*

The example below shows how a uniform thickness plate can be used to represent a motherboard under load (see Figure 146). In this case, the board has a heatsink mounted at its center and the board is supported at its corners. When the board is subjected to a shock condition, the heatsink will exert an inertial load at the center of the board. The board can be represented as a plate in bending. Several simplifying assumptions must be made. The board support condition can be grossly represented as pinned supports (or simple supports) around its edges. The heatsink inertial load can be represented as a uniform pressure load over a small circular region at its center. These simplifications allow a coarse representation of the problem that matches configurations contained in structural handbooks. The load-deflection relationship for this simplified plate representation is given by:

Equation 55: Plate Deflection Equation

$$\Delta = \frac{\alpha P b^2}{E t^3}$$

Where:

Δ = Deflection at center of plate

P = Heatsink inertial load

b = Plate width

E = Material modulus

t = Plate thickness

α = numerically derived fit parameter related to plate aspect ratio (α = 0.1267 for a/b = 1.0; reference *Roark's Formulas for Stress and Strain* for other a/b values)

Figure 146: Plate Representation of Motherboard in Bending
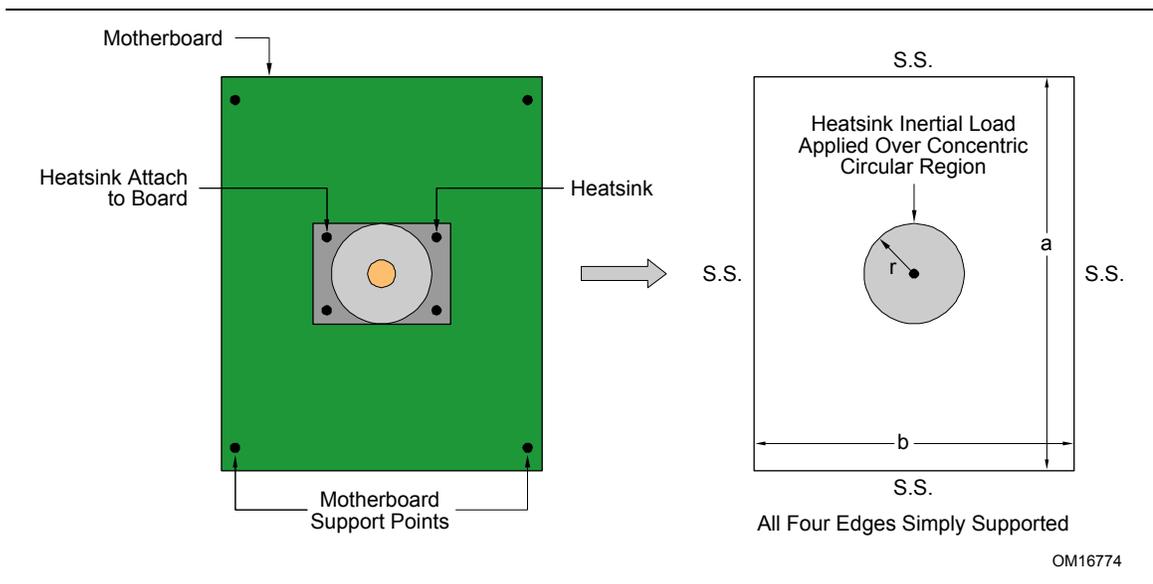
Again, the sensitivity to the support condition can be investigated by considering a case where the edges of the plate are clamped rather than simply supported. In this case, the displacement still given by Equation 55. But in this case, the fit parameter, α, drops to $\alpha = 0.0611$ (for plate aspect ratio a/b = 1.0). So the plate deflection is cut in half with the addition of clamped boundary supports.

4.4.1.1.3   Representation of Complex Structural Components

Beam and plate representations can be effective tools in developing first order analysis of system component deflections. However, their accuracy is limited due to the simplifying set of assumptions associated with their use. As component geometry, load distribution, or boundary conditions become more complicated, finite element analysis (FEA) is required to obtain reasonably accurate prediction of component deflections.

As with any numerical analysis, special care should be taken to ensure that the model geometry and boundary conditions are input properly. It is always a good practice to check the output from finite element analysis against hand calculations like those examples shown in the sections above. The analyst would not expect the predictions to match precisely, but should be able to check that the component deflections generated from the FEA are on the same order as those predicted through hand calculations. The analyst should also be able to use the hand calculations to understand trends (the sensitivity of the results to certain parameters), and then confirm that the FEA results display the same trends. System-level FEA modeling can be very complicated, and model set-up is prone to mistakes. Therefore, checking model results against hand calculations is absolutely critical in ensuring model accuracy.

4.4.1.2 Evaluating Component Deflections Against Failure Thresholds

Once component deflections have been determined, they can be evaluated against the failure threshold to determine whether the level of deflection is acceptable. For example, as a motherboard deflects upward during shock, the deflection in the vicinity of the MCH heatsink can be compared to the gap between the top of that heatsink and the nearest system component above it (such as a peripheral drive). If the predicted upward deflection were less than the design gap, then no contact between the components would be expected during shock. However, if the deflection were greater than the design gap, then the heatsink would be expected to strike the peripheral drive, placing the MCH bare die component (as well as the

drive) at risk of failure. In this case, the design gap between the two components may need to be adjusted, or the board support points may need to be reinforced to reduce the board deflection.

4.4.2 Stress Failure Modes

Stress-related failures can occur when load-carrying components are flexed or strained to the point of yielding or fracturing. Excessive yielding can generate permanent deformation or distortion of the component from its original shape. Mechanical retention capability, preload capability or other mechanical functionality of the component can be lost or degraded in this case. In its limit, yielding can lead to component fracture.

Typically, components subjected to bending loads are susceptible to this failure mode. These components include the chassis pan, chassis side panels, peripheral support brackets, motherboard, heatsink retention components, etc. Stress-related failures can be avoided with good system structural design. Strategies for managing component stress levels are given in Section 4.5.

Solder ball fracture is a second failure mode associated with over-stress. Surface-mounted components (such as the processor socket or MCH) can experience this type of failure. In conditions where the motherboard is subjected to severe bending deflection, these surface-mount components develop tension and shear stresses within each solder ball in the BGA. Such stress conditions can generate ball fracture. This failure mode is addressed separately in Sections 4.6 and 4.6.3.

4.4.2.1 Calculating Component Bending Stress

Stress levels for the majority of system components are dominated by bending stresses – those stresses that result from component bending. (One notable exception is the stress condition for BGA loading, which is dominated by tension and shear stresses. Evaluation against this failure mode will be covered in Section 4.4.3). Component bending stress levels are dependent on three primary factors: inertial loading, cross-sectional geometry of the structural element, and support conditions of the structural element.

4.4.2.1.1   Beam Representation of Structural Components

Recalling Section 4.4.1.1.1, a common representation for certain brackets and clips is achieved with a uniform cross-section beam. That section showed how to determine component deflections for a beam in bending. The present section will show how to calculate component stresses for the same example.

The basic stress equations are provided below for this simple example. For a more thorough presentation of beam bending formulas and theory, reference *Roark's Formulas for Stress and Strain*.

The example from Section 4.4.1.1.1 showed how a simple beam could be used to represent a heatsink clip (see Figure 144). The clip is supported at its two ends and loaded at its center. This sets up a bending condition in the beam. The beam load condition can be represented with the free body diagram (FBD) shown in Figure 147.

Figure 147: Plate Representation of Motherboard in Bending
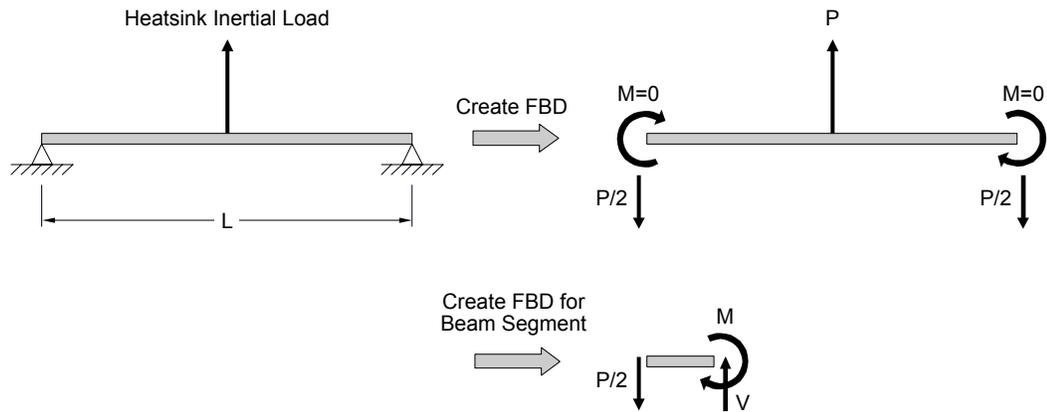
From the FBD, it is possible to determine the reaction loads at the end supports. Since this example is loaded at its center, the reaction at both ends is simply P/2. Figure 147 also shows the FBD for a segment of the beam. This beam segment can be used to derive the shear and moment diagrams for the beam. Shear and moment diagrams offer another method of visualizing the beam bending condition.

The shear can be calculated at any section along the length of the beam by integrating the applied load distribution along the length of the beam up to that section. The governing equation is given by:

Equation 56: Shear Load

$$V = \int_0^x P(x)$$

Where:

   V = Shear load acting on the face of the section

   P(x) = Distribution of applied load as a function of x-coordinate

   x = x-coordinate of the section

For the present example, evaluation of Equation 56 yields the following shear distribution:

Equation 57: Shear Distribution

$$V(x) = \begin{cases} P/2, & \text{for } 0 \le x \le L/2 \\ -P/2, & \text{for } L/2 \le x \le L \end{cases}$$

Likewise, integrating the shear distribution across the length of the beam can create the bending moment diagram. The governing equation is given by:

Equation 58: Bending Moment

$$M = \int_0^x V(x)$$

Where:

M = Bending moment acting on the face of the section

For the present example, evaluation of Equation 58 yields the following bending moment distribution:

Equation 59: Bending Moment Distribution

$$M(x) = \begin{cases} Px/2, & \text{for } 0 \leq x \leq L/2 \\ P(L-x)/2, & \text{for } L/2 \leq x \leq L \end{cases}$$

The bending stress can then be calculated at any point along the length of the beam with the following equation:

Equation 60: Bending Stress

$$\sigma = \frac{Mc}{I}$$

Where:

c = Distance from neutral bending axis to top or bottom of beam

I = Beam cross-sectional moment of inertia

For the present example, evaluation of Equation 60 yields the following bending stress distribution:

Equation 61: Bending Stress Distribution

$$\sigma(x) = \begin{cases} Ptx/4I, & \text{for } 0 \leq x \leq L/2 \\ Pt(L-x)/4I, & \text{for } L/2 \leq x \leq L \end{cases}$$

Where:

t = Thickness of beam

Applying these equations to the present example yields the shear and bending moment diagrams shown in Figure 148. The bending stress along the length of the beam is also shown in this figure. Notice the peak moment and bending stress occurring at the center of the beam.

Figure 148: Diagrams for Shear, Bending Moment, and Bending Stress in a Simply Supported Beam
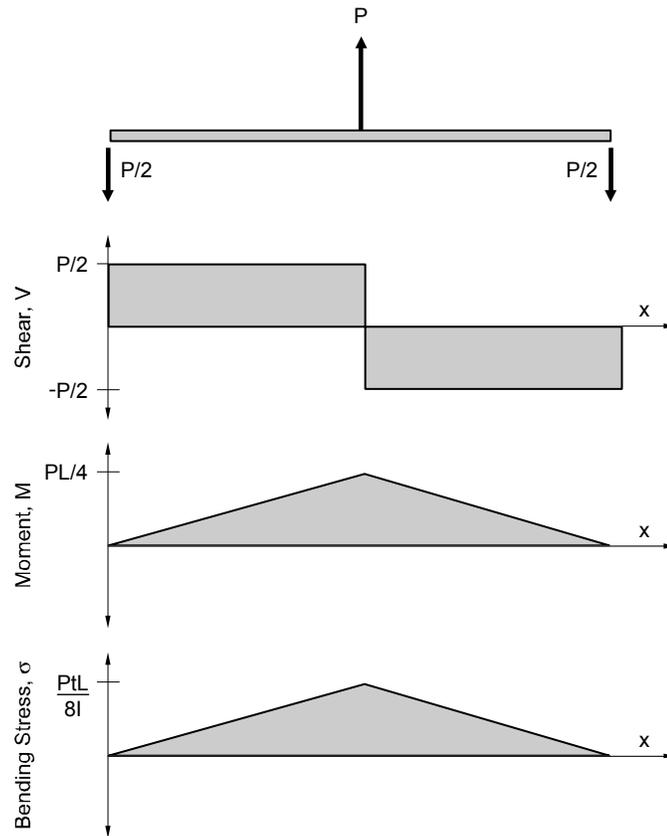
OM16776

Again, the boundary condition dependency can be investigated. If the clip is fastened to the motherboard in such a way as to restrain rotation at its ends, then the beam representation changes to clamped end supports (see Figure 145). The change in boundary conditions substantially changes the stress response of the beam. The equations for shear, bending moment, and bending stress distributions are given below:

Equation 62: Shear Distribution (Clamped End Supports)

$$V(x) = \begin{cases} P/2, & \text{for } 0 \le x \le L/2 \\ -P/2, & \text{for } L/2 \le x \le L \end{cases}$$

Equation 63: Bending Moment Distribution (Clamped End Supports)

$$M(x) = \begin{cases} P(L/8 - x/2), & \text{for } 0 \le x \le L/2 \\ P(-3L/8 + x/2), & \text{for } L/2 \le x \le L \end{cases}$$

Equation 64: Bending Stress Distribution (Clamped End Supports)

$$\sigma(x) = \begin{cases} Pt(L/4 - x)/4I, & \text{for } 0 \le x \le L/2 \\ Pt(-3L/4 + x)/4I, & \text{for } L/2 \le x \le L \end{cases}$$

The new shear and bending moment diagrams, as well as the stress distribution are shown in Figure 149.

Figure 149: Diagrams for Shear, Bending Moment, and Bending Stress in a Fixed Support Beam
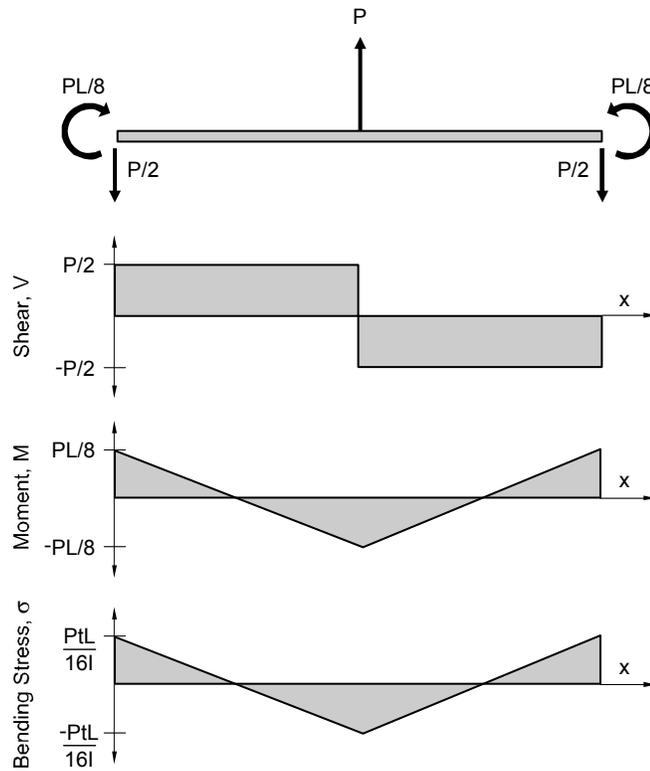
OM16777

4.4.2.1.2   Plate Representation of Structural Components

Recalling Section 4.4.1.1.2, a common representation for broad, flat structural components such as the chassis pan or motherboard is achieved with a uniform thickness plate. That section showed how to determine component deflections for a plate in bending. The present section will show how to calculate component stresses for the same example.

The basic stress equations are provided below for this simple example. For a more thorough presentation of plate bending formulas and theory, reference *Roark's Formulas for Stress and Strain.*

The example from Section 4.4.1.1.2 showed how a simple plate could be used to represent a motherboard under the inertial load of a heatsink (see Figure 146). The plate representation was supported with pinned constraints along its edges and was loaded at its center. The maximum bending stress occurs at the center of the plate and is given by:

Equation 65: Maximum Bending Stress

$$\sigma = \frac{3P}{2\pi t^2}\left[(1+\nu)\ln\frac{2b}{\pi r} + \beta\right]$$

Where:

σ = Plate bending stress

P = Heatsink inertial load

t = Thickness of plate

ν = Poisson's ratio for material

b = Plate width

r = Radius of circle over which load is applied

β = Numerically derived fit parameter related to plate aspect ratio (β = 0.435 for a/b = 1.0; reference *Roark's Formulas for Stress and Strain* for other a/b values)

Equation 65 provides the maximum principle stress in the plate. The actual stress state is 2 dimensional. That is, it has stress components in both x and y-directions. This necessitates use of a failure theory that can account for multi-dimensional stress states, such as Von Mises, Max Principle Stress, or Tresca failure theories (reference *Roark's Formulas for Stress and Strain* for explanation and use of these theories). The Von Mises failure criterion will be presented in Section 4.4.2.2.

4.4.2.1.3 Representation of Complex Structural Components

Beam and plate representations can be effective tools in developing first order analysis of system component stresses. However, their accuracy is limited due to the simplifying set of assumptions associated with their use. As component geometry, load distribution, or boundary conditions become more complicated, finite element analysis (FEA) is required to obtain reasonably accurate prediction of component deflections.

As cautioned in Section 4.4.1.1.3, special care should be taken to ensure that the model geometry and boundary conditions are input properly. Hand calculations should be used to confirm that the component stresses generated from the FEA are reasonable and following the same trends as predicted by beam bending and plate bending theories. Checking FEA results against hand calculations is absolutely critical in ensuring model accuracy.

4.4.2.2 Multi-Dimensional Stress States

Stress is inherently a three-dimensional entity, and is commonly represented in tensor formulation. A symmetric, three-by-three matrix as shown in Equation 66 represents the stress tensor.

Equation 66: Stress Tensor Matrix

$$[\sigma] = \begin{bmatrix} \sigma_x & \tau_{xy} & \tau_{xz} \\ \tau_{xy} & \sigma_y & \tau_{yz} \\ \tau_{zx} & \tau_{yz} & \sigma_z \end{bmatrix}$$

Beam bending theory deals with only a single component of the stress tensor (σ_x), since beam bending generates something very close to a uniaxial stress state. Likewise, plate bending theory assumes a state of plane stress, and therefore uses only a subset of the stress tensor. Finite element analysis will use appropriate subsets of the stress tensor depending on the element type being used. Beam elements will use a single element of the tensor, while plane stress elements (such as plate elements) will use only the in-plain tensor elements. Solid elements will typically use the full stress tensor.

Failure theories are generally developed to account for a full three-dimensional stress state. An example is the Von Mises failure theory, which compares the distortional stress (or Von Mises stress) to the yield stress:

Equation 67: Material Yield Condition

$$\text{Material yield occurs if } \sigma_{VM} \geq \sigma_y$$

Where:

Equation 68: Von Mises Stress

$$\sigma_{VM} = \left[\frac{1}{2} \left[(\sigma_x - \sigma_y)^2 + (\sigma_y - \sigma_z)^2 + (\sigma_z - \sigma_y)^2 \right] + 3(\tau_{xy}^2 + \tau_{yz}^2 + \tau_{zx}^2) \right]^{\frac{1}{2}}$$

σ_y = Material yield stress

In the case on plate bending, a plane stress state exists, in which $\sigma_z = \tau_{yz} = \tau_{zx} = 0$, and the equation for Von Mises stress can be simplified to:

Equation 69: Von Mises Stress (Plane Stress State)

$$\sigma_{VM} = \left[\sigma_x^2 + \sigma_y^2 - \sigma_x \sigma_y + 3\tau_{xy}^2 \right]^{\frac{1}{2}}$$

For the case of beam bending (uniaxial stress state), the equation for Von Mises stress can be simplified even further to:

Equation 70: Von Mises Stress (Uniaxial Stress State)

$$\sigma_{VM} = \sigma_x$$

Other failure theories such as Tresca of Max Principle Stress are also based three-dimensional stress states and can be simplified in the same manner based on the applicable stress state. In all cases, it is important to account for multi-dimensional stress components when present. Failure to account for all substantial stress components can result in under-predicting the effective stress state and lead to a non-conservative evaluation of component strength.

4.4.2.3 Evaluating Component Stresses Against Failure Thresholds

Once component stresses have been determined, they can be evaluated against the failure threshold to determine whether the level of stress is acceptable. The acceptable level of stress will depend primarily on the material strength. But will also depend on the application. In some cases, the component will be considered to have failed once the material yield strength is reached. In other cases, the component will be allowed to suffer local yielding until the yield zone becomes too large or the ultimate strength is exceeded.

The material yield strength is the stress level beyond which the material will encounter plastic deformation (or yield). Once yielding has occurred, the component will not resume its original shape after the load has been removed. That is, the component has permanent (or plastic) deformation. The yield point can be used a conservative threshold for failure. Designing components to keep stresses below the yield point will help ensure a robust structural design. However, in many cases, local yielding may be allowable. And it may be difficult or costly to design the part to maintain all stresses below yield, especially in the areas of local stress concentrations around holes or corner radii.

In these cases, the designer may consider the approach of allowing local yielding in the component. Local yielding is often acceptable for components that do not have stringent shape constraints, and are therefore allowed to undergo small amounts of plastic deformation. The difficulty in this approach lies in the non-linear material response once the material has exceeded yield (see Figure 150). This non-linearity requires use on non-linear analysis to maintain accurate stress predictions. Non-linear analysis capabilities are readily available in most of the commercially available FEA packages. However, model set-up time, debug time, and run-time will increase when performing non-linear versus linear analysis.

If local yielding is allowed to occur, then the applicable failure threshold typically shifts to material ultimate strain. In the theoretical limit, the material strain level would be allowed to increase until it reached its point of ultimate strain (also referred to as *elongation limit*). In practice, however, substantial design margin (100% or greater) is used when working to the elongation limit due to the high degree of variability around this material property. Even with such a large design margin, use of the ultimate strain threshold affords substantial amount of additional design space beyond what would be available when using the yield point as the design limit.

Figure 150: Typical Stress-Strain Curve for Ductile Material
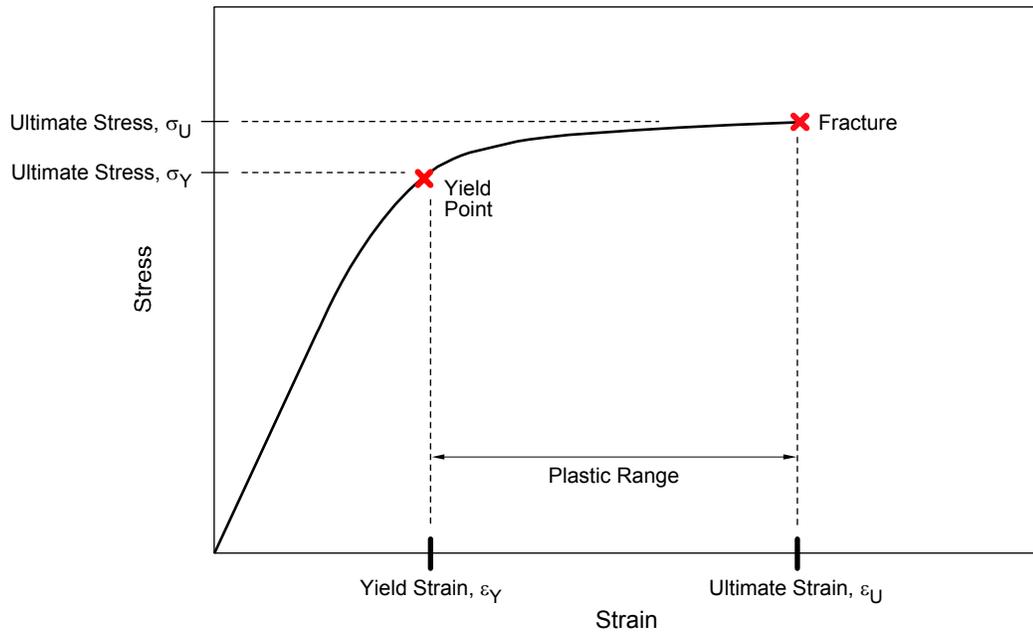
4.4.3 Solder Ball Failure Modes

Solder ball fracture is a second failure mode associated with over-stress. Surface-mounted components (such as the CPU socket or MCH) can experience this type of failure. In conditions where the motherboard is subjected to severe bending deflection, these surface-mount components develop tension and shear stresses within each solder ball in the BGA. Especially susceptible are the corner balls of the array, and more generally, all the balls in the outer ring. These critical balls take substantial loading in a board flexure condition.

In a shock condition, the inertial loads act to flex the motherboard (see Figure 151). This generates curvature throughout the board, including the area under the surface-mounted component. The component will be forced to flex with the board, however, the greater rigidity of the component will resist this flexure, thereby generating loads at the interface between the board and component. Those interface loads are carried through the solder balls in the BGA.

The solder balls at edges of the array carry substantially higher loads than interior balls. The edge balls are responsible for imparting the greatest amount of bending moment into the surface-mounted component. The interior balls are effectively protected by the presence of the edge balls. As a result, however, the edge balls are the most susceptible to fracture. Corner balls are subjected to a compounded effect since they lie on two adjacent edges. As such, the ball loads in the corners are higher than those anywhere else on the edge.

Figure 151: Board Flexure in Shock Causing Solder Ball Stress

OM16779
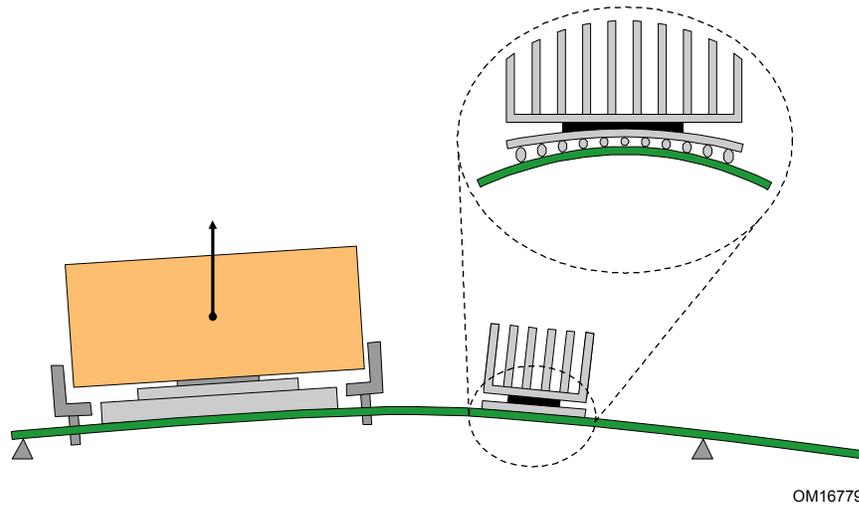
The amount of load carried in the solder balls is primarily dependent on three factors: the amount of board flexure (in the local region of the component), the bending rigidity of the component, and the ball pattern of the BGA.

There is a direct relationship between ball load and motherboard flexure. As Figure 151 suggests, the greater the board flexure, the higher the BGA load. Controlling board flexure is the key to protecting surface-mounted components. Solder ball fracture can be avoided with good system structural design. Strategies for managing board flexure and BGA stress levels are given in Section 4.6 and 4.6.3.

The component bending rigidity and the BGA pattern are fixed for a given component, and therefore, cannot be modified by the system designer to improve performance. The system designer must focus instead on managing motherboard flexure to ensure protection for the surface-mounted components.

4.4.4 Case Study – Applied Engineering for Type I Heatsink Clip Reference Design

The Type I Reference design Thermal Module uses a steel clip to attach the CPU heatsink to the structural duct. The deflection and stress analyses of the clip are provided here as a case study for several of the analysis techniques described above.

The clip used in the Type I Reference Design TMA is made of embossed sheet metal. An image of the clip is shown in Figure 152. The clip has a hole in its center to accommodate the circular copper core of the heatsink. A snap ring is affixed to the bottom of the core to hold the clip and core together. The clip interfaces with the duct at its ends as shown in the figure.

Figure 152: 12.9L Reference Design Heatsink Clip

Heatsink attaches
to center hole

Duct interfaces
at ends of clip
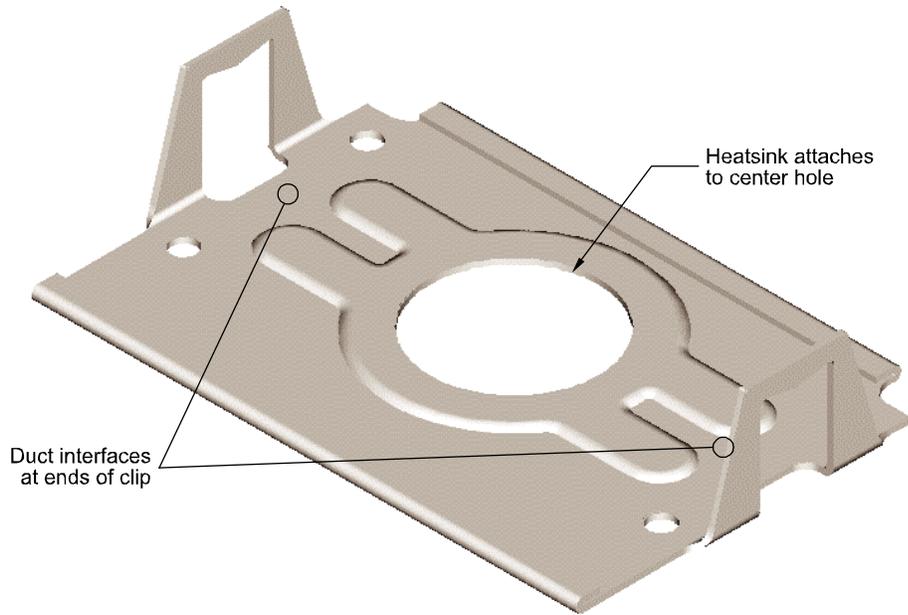
OM17098

4.4.4.1 Clip Bending Response Using Hand-Calculations

The clip serves to transfer load from the heatsink to the duct in both the preload and shock load conditions. The z-direction shock load conditions are considered in this example. In this load case, the heatsink inertial load is transferred into the clip around the periphery of the center hole. The clip is restrained against vertical movement by the points of contact at the duct. These load and displacement constraints are depicted in Figure 152.

This clip bending problem can be approximated with a simple beam bending representation. Although the sheet-metal clip has the geometric form of a thin plate, its load and constraint conditions will generate a one-dimensional bending response along its length. As a result, it takes on the approximate bending characteristics of a beam, and can be represented as such.

Figure 153 shows the simplified beam representation for the problem. The ends of the beam are constrained with pinned supports to represent the duct interface points. The heatsink inertial load is applied as a uniformly distributed load across the center of the beam. The load is distributed over a length equal to the diameter of the copper core. The heatsink inertial load, P is calculated from Equation 71 as follows:

Equation 71: Heatsink Inertial Load

$$P = maf = (0.90 \,\text{kg})\,(30 \,\text{G})\,(9.81 \,\text{m/s}^2)\,(3.0) = 800 \,\text{N}$$

Where:

m = component mass = 0.90 kg

a = peak input acceleration = 30 G * 9.81 m/s^2

f = DAF = 3.0

The beam cross-section is represented with an effective MOI and thickness based on the actual cross-section of the clip. For this clip, the cross-section geometry varies along its length. The bending stress, and to some extent, the effective bending stiffness of the clip, is governed by the cross-sectional geometry near the center of the clip. This is the point where the bending moment is maximum and the MOI is minimum. This combination will generate the highest bending stress at the center of the clip. For this reason, the MOI is calculated for the cross-section at the center of the clip and assigned as the effective MOI for the simple beam representation. Doing so yields an effective MOI of 37 mm^4 and an effective clip thickness of 0.95 mm.

Figure 153: Heatsink Clip Load Condition

OM17099
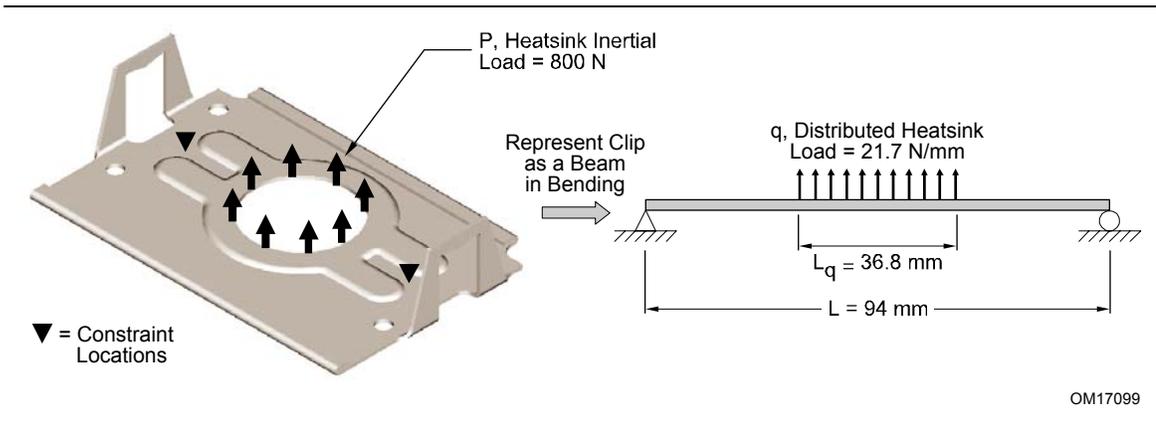
Figure 154 shows the resulting free body diagram for the beam. This diagram is used with Equations Equation 56, Equation 58 and Equation 60 to generate the shear, bending moment, and bending stress diagrams for the clip (see Figure 155). The equation for the maximum bending stress can be derived directly from this process to yield:

Equation 72: Clip Bending Stress

$$\sigma_{max} = \frac{Pc(2L - L_q)}{8I}$$

The beam deflection can also be derived. The closed-form solution for the maximum deflection (occurring at the center of the beam) is given as [reference *Roark's Formulas for Stress and Strain*],

Equation 73: Clip Deflection Equation

$$\Delta_{max} = \frac{P(8L^3 - 4LL_q^{\,2} + L_q^{\,3})}{384EI}$$

And the effective bending stiffness for the clip (defined as the ratio of applied load to maximum deflection) is given by,

Equation 74: Clip Effective Bending

$$k = \frac{8L^3 - 4LL_q^2 + L_q^3}{384EI}$$

Substituting the values, $P = 800$N, $L = 94$ mm , $L_q = 36.8$ mm, $I = 37$ mm^4, $c = 0.95$ mm, and $E = 200$ GPa yields the follow key characteristics for the clip in the shock load condition:

$\sigma_{max} = 388$ MPa

$\Delta_{max} = 1.74$ mm

$k = 460$ N/mm

These equations can be used for a number of evaluations, but most directly, the bending stress can be compared to the material yield point to determine whether the clip is expected to yield in a shock condition, the deflection can be used to evaluate against deflection-based failure modes, and the stiffness can be used to predict preload values.

These hand calculations above clearly require use of simplifying assumptions for the clip geometry and load conditions. As such, the results will only be approximate, but may still be useful in preliminary design to support feasibility studies, part sizing, and trade studies. Results like this are also critical in checking the reasonableness of finite element analysis results.

Figure 154: Free Body Diagram for Heatsink Clip under Shock Load
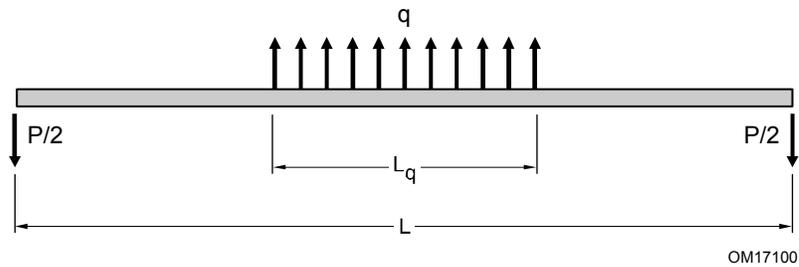
Figure 155: Shear, Bending Moment, and Bending Stress Diagrams for Heatsink Clip under Shock Load

OM17101
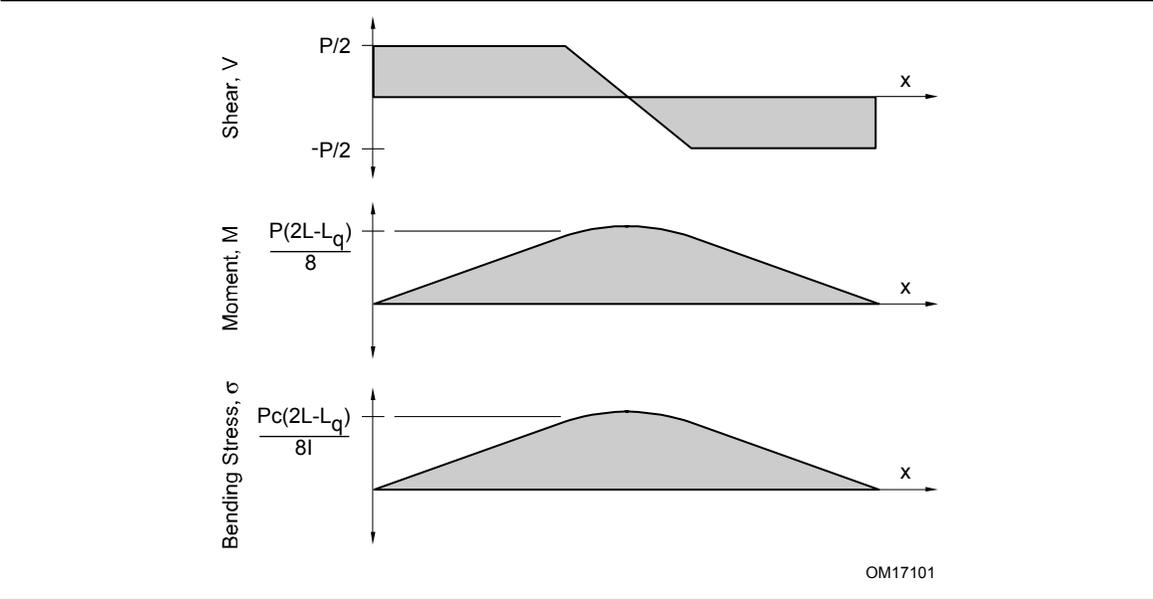
4.4.4.2 Clip Bending Response Using Finite Element Analysis

When the simplifying assumptions associated with hand calculations are not practical, or when a higher level of accuracy is required, finite element analysis can be used to evaluate the component structural response. In this section, the CPU heatsink clip is evaluated in the z-direction shock load condition using FEA. Clip deflection, bending stress and bending stiffness results are presented.

Figure 156 shows the clip load and displacement constraint conditions. In this case, the periphery of the clip hole is constrained in z-direction while the shock load is applied to the duct contact points at the ends. This is slightly different than the physical condition of interest where the ends are constrained and the load is applied at the hole. However, the distinction is immaterial to the FEA result, and the former constraint definition turns out to have more straightforward implementation.

Figure 156: Boundary Conditions for Finite Element Analysis
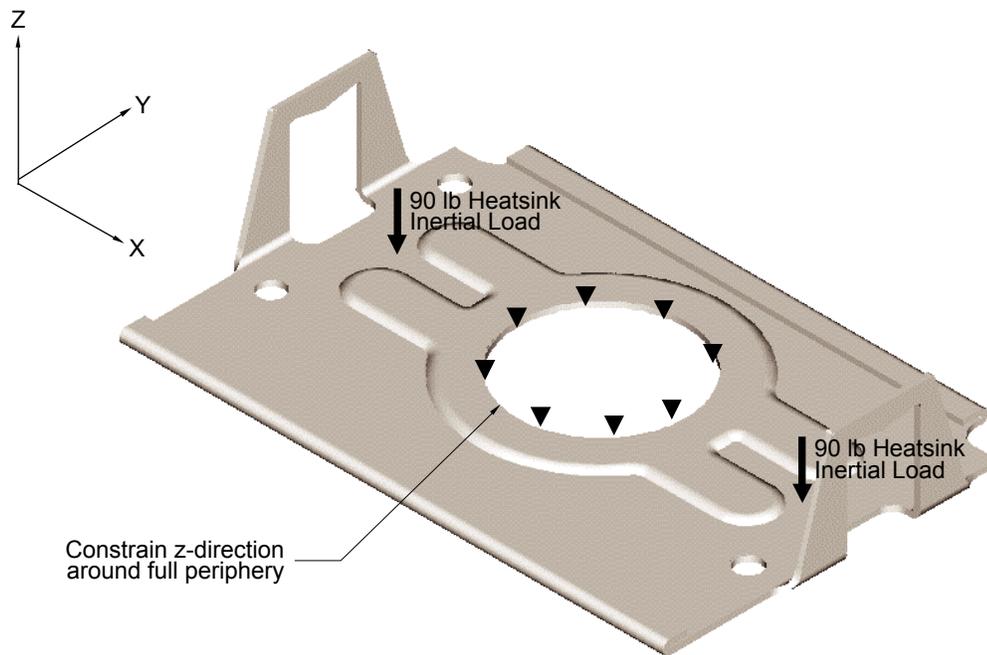
Figure 157 shows the resulting displacement plot. Notice that the displacement contour pattern shows several gradations in displacement as we move from clip center outward toward the edge. This is similar to the expected deflection pattern of a beam, and therefore lends credence to the simplifying assumptions that were made for the hand calculations in the preceding section. Notice the displacement at load point of 2.36 mm. The clip bending stiffness can then be calculated as k = (800 N)/(2.36 mm) = 340 N/mm. Both the displacement and stiffness result from the FEA predict a somewhat more compliant clip that the one predicted in the hand calculation (which generated a stiffness prediction of 460 N/mm).

Figure 158 shows the resulting bending stress plot in terms of Von Mises stress. The results show the highest stresses occurring at the mid-point of the clip (as was predicted in the hand calculations). But the stress is clearly affected by the presence of the core and the fact that the loads and constraints are focused along the centerline of the clip rather than spread out uniformly across its width. The maximum stress is

shown to be 724 MPa. This is higher than the 390 MPa predicted with the hand calculation primarily due to the presence of the core and the position of the load and constraint points.

Figure 157: Displacement Contour Plot for Heatsink Clip under Shock Load Condition
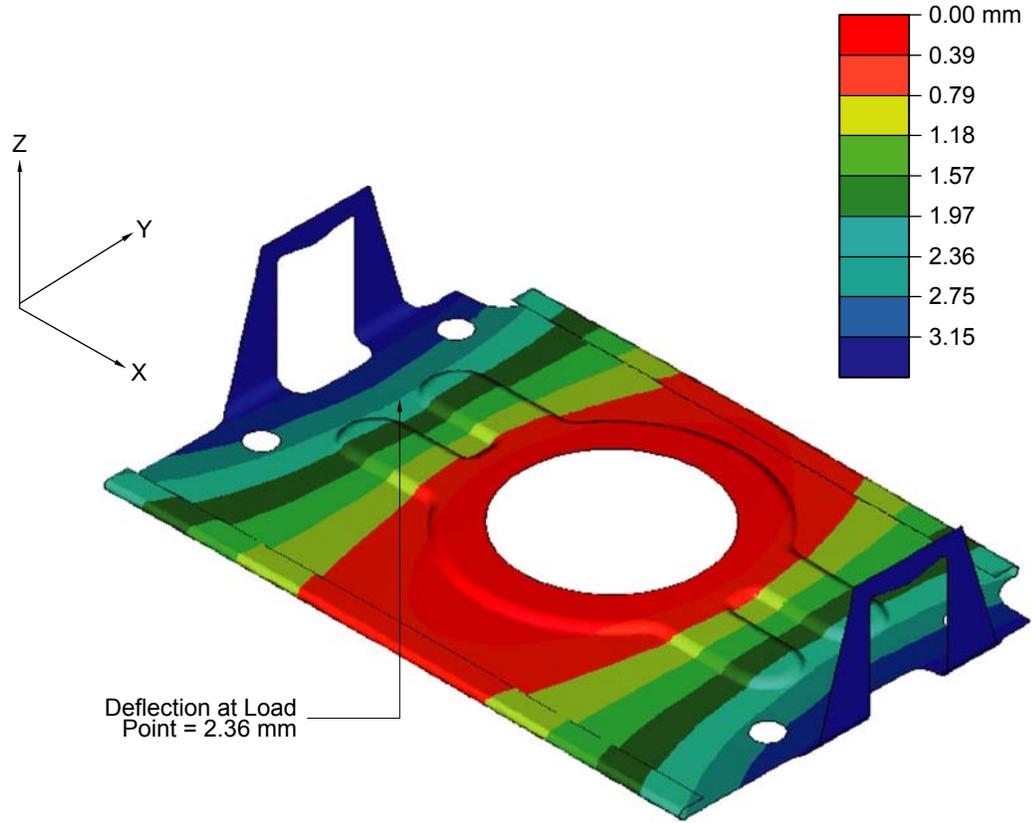
Figure 158: Von Mises Stress Contour Plot for Heatsink Clip under Shock Load Condition

Max stress = 724 MPa

Typical stress for center
section = 500 - 600 MPa
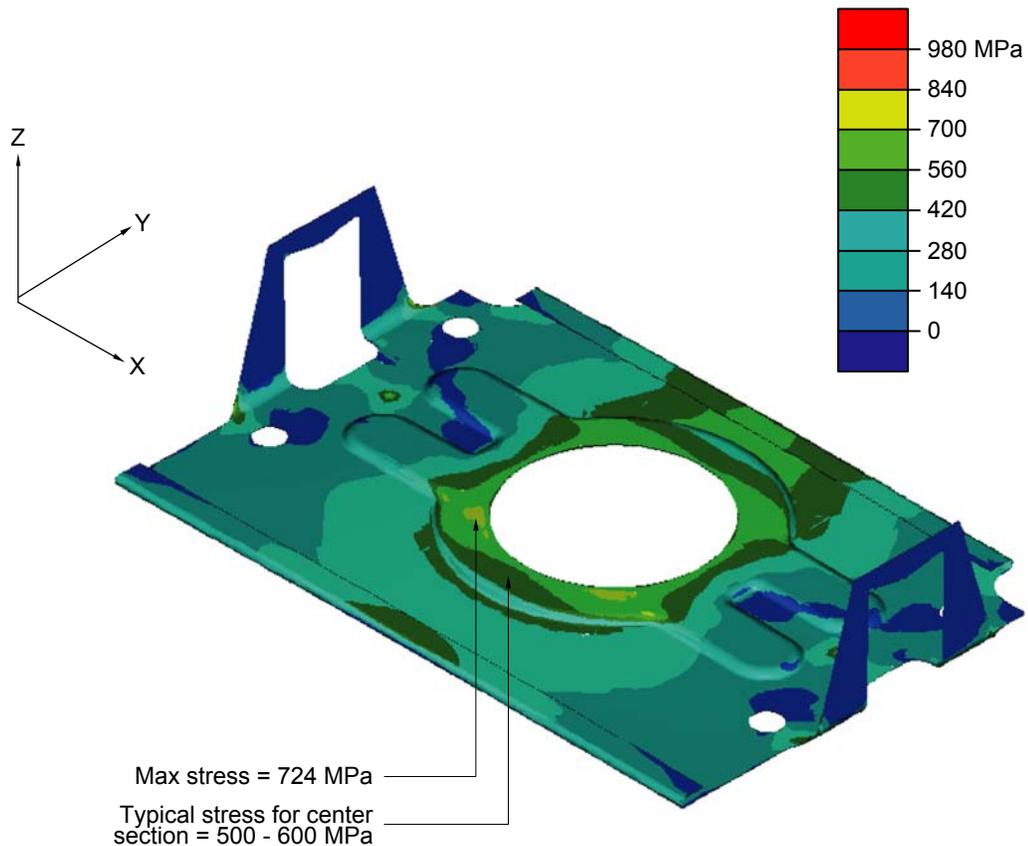
OM17104

This level of accuracy in the hand calculation is typical for all but the most straightforward beam representations. And while the hand calculations are indispensable for preliminary analysis and trade studies, they must typically be followed up with FEA to support detailed design development.

4.5 Load Management Strategies

As stated in earlier sections, platform structural design is centered on management of inertial loads under shock and vibration conditions. This section describes strategies for managing those loads to prevent over-deflection or over-stressing of components.

In general there are three strategies for preventing over-deflection or over-stress of a component in a given load condition:

- Design the component strength to carry the given load.

- Utilize an alternate load path to reduce the amount of load being carried by the component.

- Utilize preload to pre-stress the component in a benign direction of stress.

Any combination of these strategies can be employed to ensure that component load-carrying capability meets the anticipated load condition. Each strategy is described in the following sections.

4.5.1 Designing for Component Strength

Component design characteristics can be tailored to support its specific load carrying requirements. Three primary factors dictate the load-carrying capability of the component: geometry, material, and attach or support conditions. Section 4.4.2.1 showed how both geometry and support conditions can affect the resulting bending stresses for a given load condition. Reference *Roark's Formulas for Stress and Strain* for an extensive discussion on this subject. Similar discussions can be found in a wide range of structural textbooks on the subject.

The impact of material properties on the component structural response can also be seen in Section 4.4.2.1. Plastic and metal materials both have a wide range of stiffness and strength characteristics to choose from. When selecting structural materials, it is important to consider modulus, yield strength, and post-yield behavior, since all can be important in managing component deflections and overall strength. For plastics, creep properties and material degradation at temperature may also be important depending on the targeted usage environment.

4.5.2 Alternate Load Paths

In some situations, there is not enough design space available to allow a component design to meet the load requirements without failure. This can be a result of limited physical space (e.g., form factor volumetric constraints or motherboard keepout zones), manufacturability limitations, cost constraints, etc. In these cases, alternate load paths can be considered to reduce the load carried by the failing component.

An example can be seen in the CPU heatsink support mechanism. The heatsink generates a substantial inertial load during shock. When the heatsink is mounted directly to the motherboard, the board itself carries the inertial load from the heatsink to the motherboard mounting locations. This generates a high degree of flexure in the motherboard, and may lead to board component damage. Since motherboard designers fix the structural design characteristics of the motherboard, system designers have no practical ability to enhance its load carrying capability. In this case, alternate load paths must be pursued to support the inertial loads that should not be carried directly by the board.

In this example, an alternate load path can be established to bridge the heatsink load through additional structural components out to the chassis hard points (i.e., structural ground points) (see Figure 159). This will effectively reduce the amount of load carried by the motherboard. When establishing an alternate load path, it is important to understand the behavior of the existing load path to identify the most effective location for the alternate path. It is also necessary to understand the load path stiffness characteristics in order to manage the amount of load going through each path. This example will be carried through the next two sections to illustrate these concepts.

Figure 159: Alternate Load Path for the Processor Heatsink Load

Inertial Load

All Load Carried Through Board
− Generates Board Flexure

Heatsink

Motherboard

Chassis Pan and Side Panels

Board Deflection

Load Shared by Board and Plate
− Board Flexure Reduced

Heatsink

Structural Plate Providing
Alternate Load Path
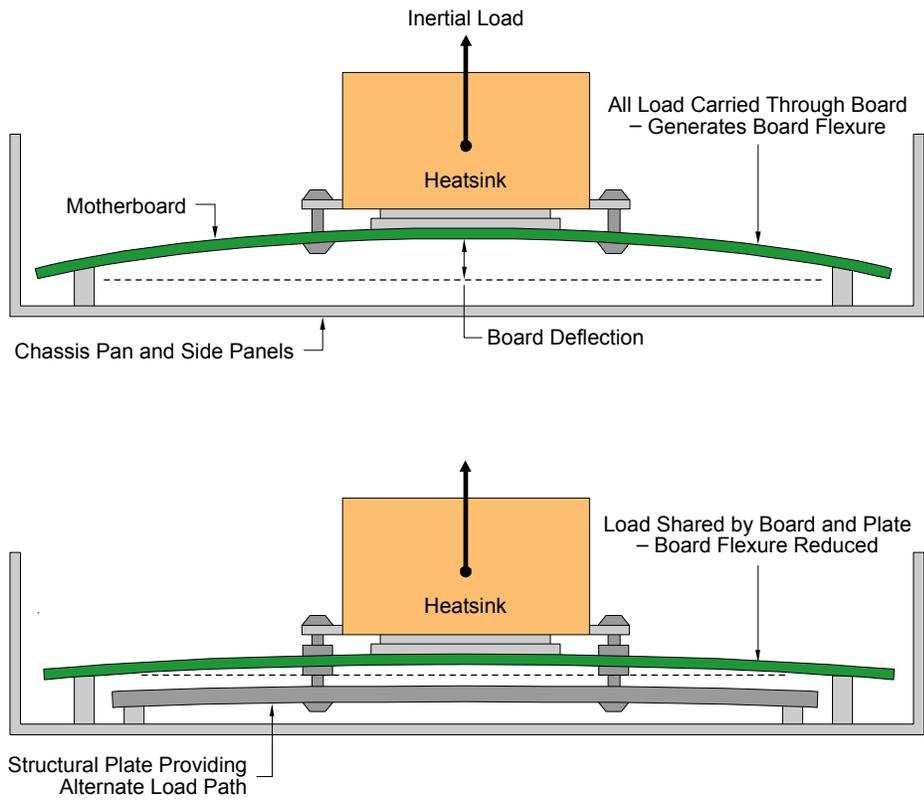
OM16780

4.5.2.1 Unsupported Load Path – Heatsink Example

The load path for the case where the heatsink is mounted directly to the motherboard is shown in Figure 160 through a series of free body diagrams. Notice in the diagram for the motherboard that the inertial forces generate the aforementioned bending condition in the board.

Figure 160: Series of Free Body Diagrams Showing Load Path
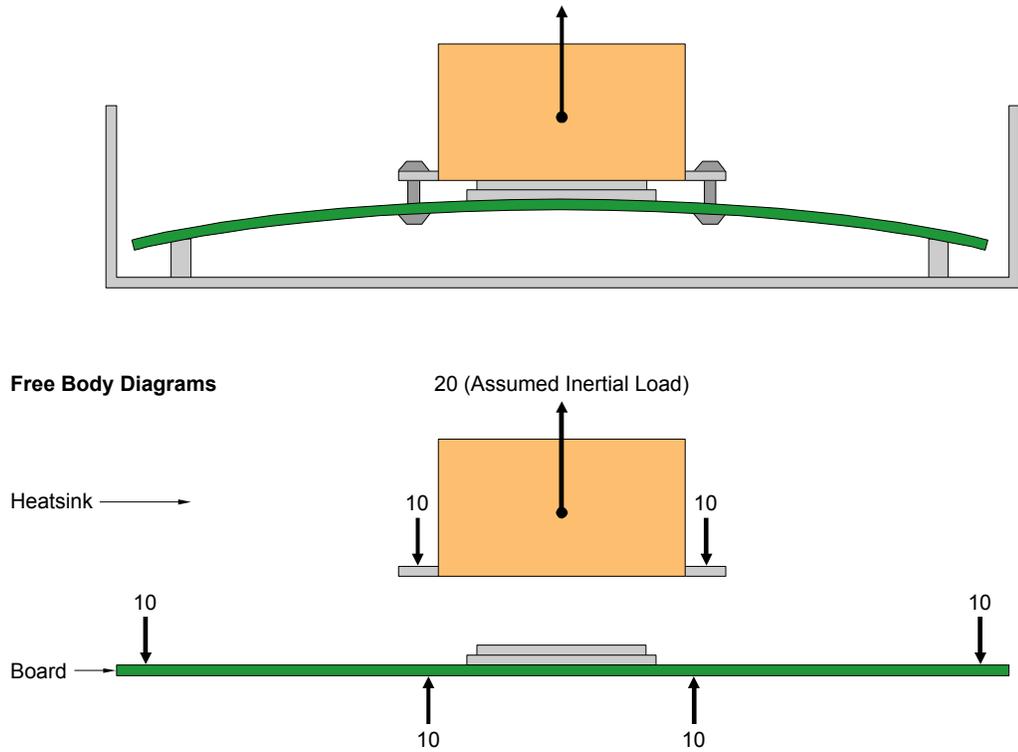
4.5.2.2 Alternate Load Path – Heatsink Example

Introduction of an alternate load path can bridge the heatsink load to the chassis hard points and reduce motherboard loading. In this example, an additional structural element is placed between the heatsink and the chassis (see Figure 160). Thus, a parallel load path is established. The free body diagrams in Figure 161 illustrate this parallel path.

Figure 161: Series of Free Body Diagrams Showing Alternate Load Path
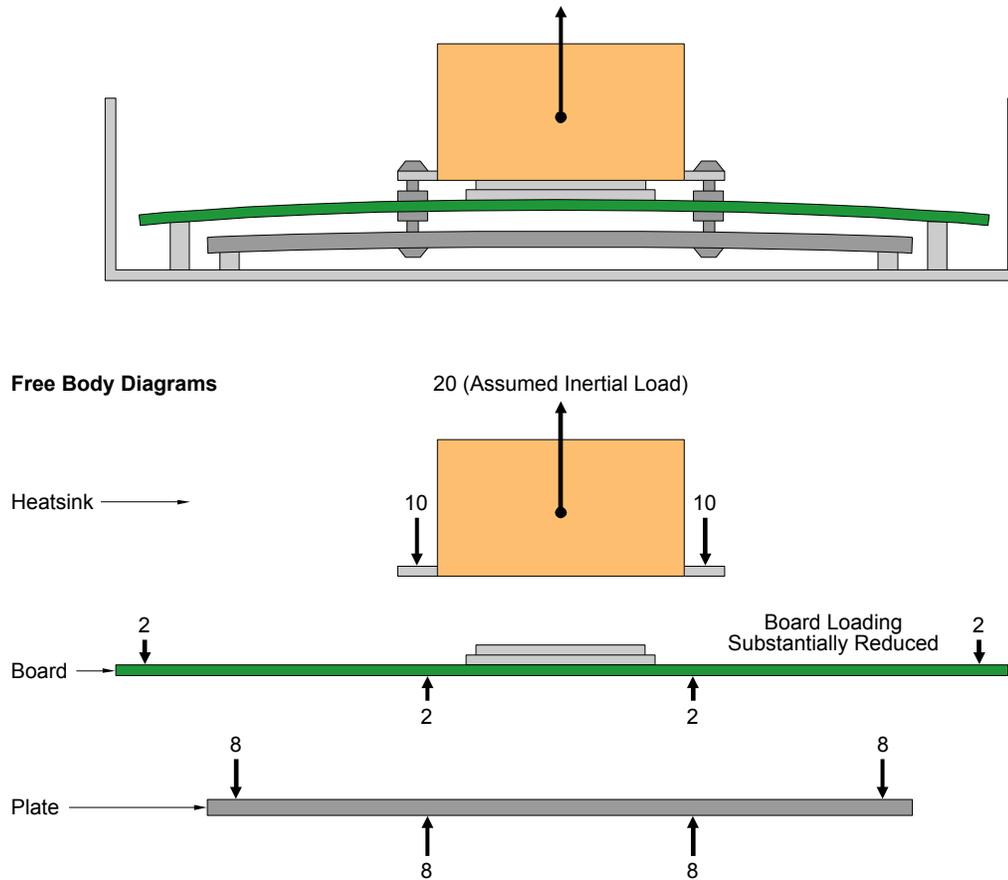
As parallel load paths are established, the relative stiffness characteristics for each path become important in determining the amount of load carried by each element. In general, the stiffer load path will carry a greater portion of the load. So if the intention of the alternate path is to remove load from the original path, then the alternate path must be made stiffer than the original (or the original path must be made more compliant).

The parallel load paths can be represented with a parallel spring model (see Figure 162). In this case, the amount of load carried by each path (or spring) is given by the equations:

Equation 75: Parallel Spring Model (Motherboard)

$$P_{MB} = \frac{Pk_{MB}}{k_{MB} + k_{Plate}}$$

Equation 76: Parallel Spring Model (Plate)

$$P_{Plate} = \frac{Pk_{Plate}}{k_{MB} + k_{Plate}}$$

Notice that as the stiffness of the alternate path (the plate) becomes much greater than that of the original path (the motherboard), the alternate path tends to carry the full load, thereby substantially relieving the load in the original path.

Figure 162: Parallel Spring Model

OM16783
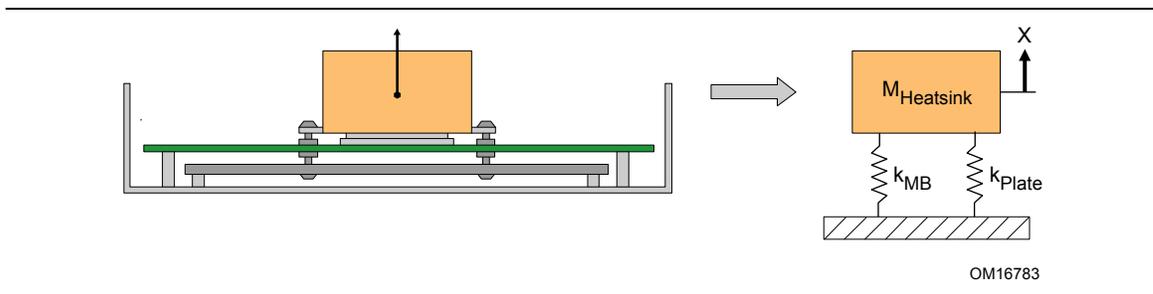
4.5.3 Preload

A third load management strategy is to utilize preload to pre-stress the component in a benign direction of stress. Many system components have a preferred direction of stress due to asymmetric geometry or differences in the tension and compression material strength properties. In these circumstances, the components can be pre-stressed in the preferred direction. When the component then undergoes loading in the opposite direction in a shock condition, the load must first overcome pre-stress generated by the preload to bring the component back to a state of neutral stress, and only then will it begin to load the component in the direction of stress to which it is more susceptible to failure.

This strategy is commonly used to protect surface mounted components such as the CPU socket. An example application of this strategy is illustrated in Section 4.6.3.

4.5.4 Case Study – Applied Structural Engineering for Thermal Module Mounting Locations

As described in previous sections, MCH and LGA775 solder joint protection is provided by the combination of preload, thermal module stiffness, and SRM stiffness. Of particular importance is the flexural shape of the board in the preload and mechanical shock conditions. It is therefore important to consider the span and fixity boundary conditions over which preload acts.

As illustrated in Figure 163, the same load will create different displacement, shear, and moment response in a beam depending on the fixity boundary conditions. With each fixity boundary condition illustrated, the displacement and curvature radius increases as the beam length increases.

The board span between the thermal module mounting locations can be thought of as a beam in bending. The larger the span between mounting locations, the greater the board displacement and curvature radius for any given preload.

Figure 163: Beam Bending Displacement, Shear, and Moment Diagrams Vs Fixity Boundary Conditions
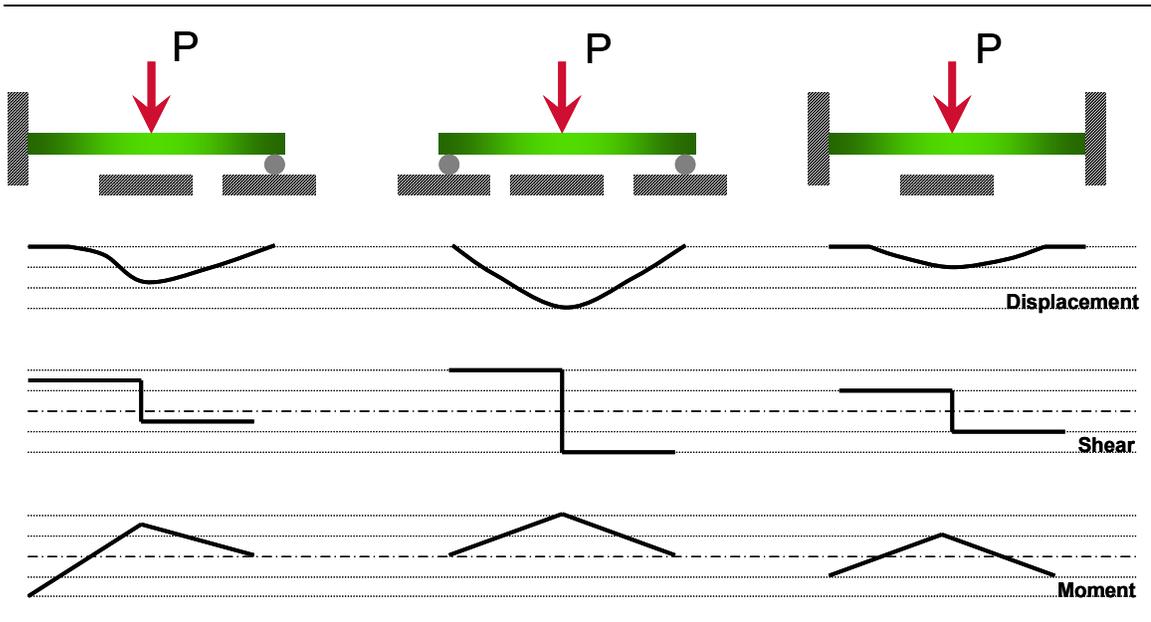
Intel's thermal module reference designs are attached to the SRM at the rear board mounting hole locations (described in detail in the Grantsdale Platform Design Guide) and at the Thermal Module Attach features (6-32 threaded PEM studs) described in the BTX Interface Specification.

The primary concern with the use of the front board mounting holes is the change in the board stiffness over which preload acts. Intel empirical data indicates that the board bending stiffness in the LGA775 area will increase slightly. Tests with the rear two mounting holes fixed (screwed down) and front two mounting holes simply supported indicates that this configuration has 10-15% lower bending stiffness than when all four mounting locations are fixed. This increase in stiffness may require an increase in preload; although it is anticipated that the preload would need to increase by <5%.

In addition to the noted board bending stiffness increase, using the front board mounting holes also changes the preload load path. The bending span of the SRM decreases by approximately 30% when the front board mounting positions are used, which increases the SRM stiffness slightly. However, this increase is insignificant given the relative stiffness of the SRM to the board.

A second concern with the use of the front board mounting hole location for the thermal module attach is the risk of thermal module duct failure. Figure 164 shows the +X shock direction Free Body Diagram for the reference design mounting scheme (on the left), which does not use the front board mounting location, and one that uses only the board mounting holes (on the right). The thermal module duct feature and fastener at the board mounting locations must be sized appropriately if a board mounting hole scheme is used since the load at these locations increases relative to the reference design mounting scheme.

Figure 164: Thermal Module Mounting Scheme Vs Reaction Loads
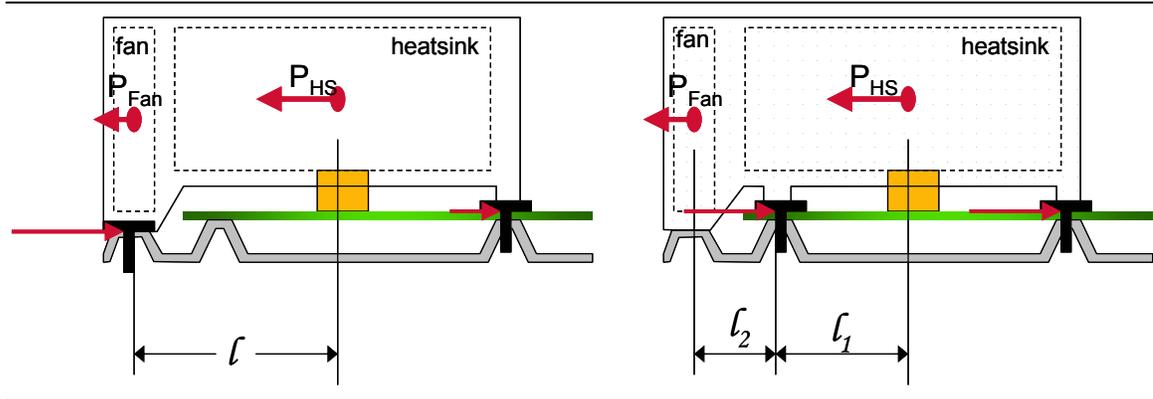
4.6 MCH BGA Protection Strategy

The surface-mounted MCH component is susceptible to solder ball fracture in shock conditions due to its proximity to the CPU heatsink. The motherboard region surrounding the CPU heatsink is subjected to substantial flexure due to the heatsink inertial load. The MCH and other board components in this region can be protected from over-stress with proper system structural design and management of the CPU heatsink inertial load.

The primary load management strategy available in this case is use of an alternate load path. The following sections describe the application of this strategy. A case study for the 12.9L Reference System Design is also presented.

4.6.1 Designing the Alternate Load Path

In an unsupported configuration, the motherboard will suffer substantial deflection under the shock load condition. The processor heatsink inertial load will deflect the board as illustrated in Figure 151. The MCH component is also subjected to severe board flexure in this condition. As such, the MCH is susceptible to solder ball fracture as described in Section 4.4.3.

An alternate load path can be established to manage this flexure. The alternate load path transfers loads from the heatsink to a stiff portion of the system, usually the chassis side walls, or the edge regions of the chassis pan.

Figure 165 shows a pair of design options for alternate load paths. The least effective path is one that simply connects the heatsink to the chassis pan beneath it. Not only is the chassis pan too compliant to provide the necessary load path stiffness, but in some situations, the response of the chassis itself can amplify the amount of bending induced in the motherboard. This is called resonance and occurs when the structural response of the chassis pan constructively couples with the structural response of the motherboard-heatsink system. In these cases, the DAF can increase substantially, causing the shock-induced displacement and flexure in the motherboard to increase rather than decrease.

The more effective alternate load path is created by bridging the load out to the edges of the chassis where it can be reacted directly into the chassis side walls. This generates a high-stiffness load path which will direct the majority of the heatsink load away from the motherboard and will avoid issues with resonance, thereby protecting the motherboard from excessive flexure.

Figure 165: Contrasting Options for Alternate Load Path Design

Ineffective: Attaching to Center of Chassis Pan Relies on This Flat Plate as Alternate Load Path (Low Bending Rigidity)

Most Effective: Attach to Stiffening Plate to Bridge Loads to Edges of Chassis Using High-Stiffness Load Path

OM16784
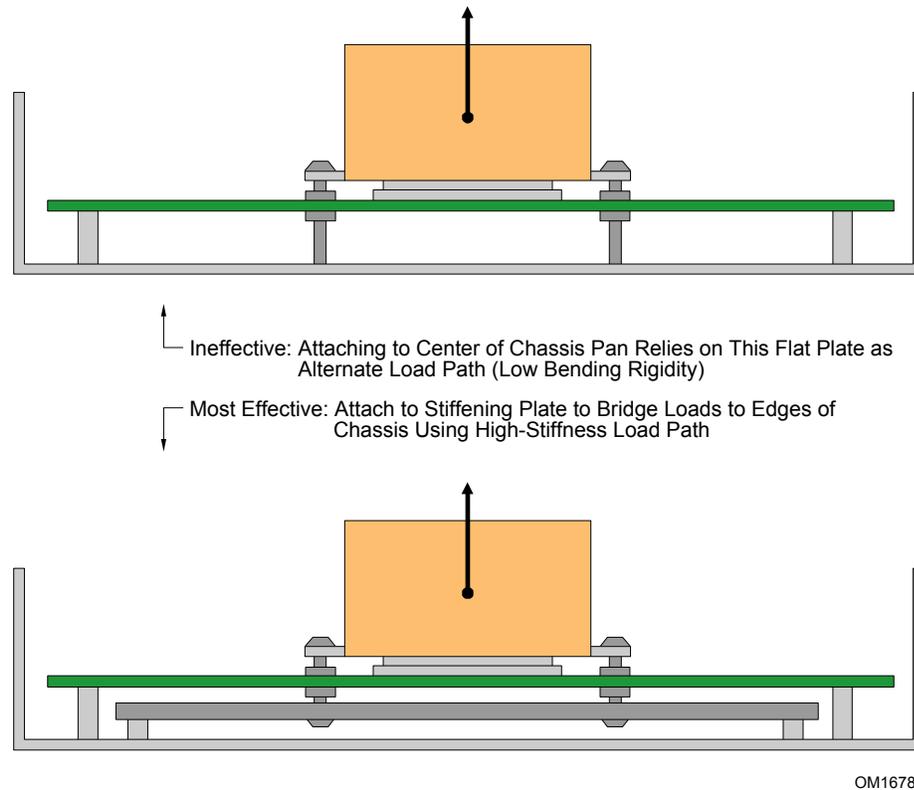
As described in Section 4.5.2.2, in order for the alternate load path to be effective, it must be substantially stiffer than the original load path. Two primary factors determine the load path stiffness: the points of load transfer to the chassis and the stiffness of the structural bridging elements. These factors are described in the following sections.

4.6.1.1 Load Transfer Points

The effectiveness of the alternate load path is dependent the points of load transfer into the chassis. The stiffest parts of the chassis (for the case of a vertical shock condition) are the sidewalls. The sidewalls act as large shear panels transferring load from the chassis pan into the packaging materials outside the system. The heatsink inertial load must eventually be transferred into these sidewalls to be reacted by the packaging.

Figure 166 shows the chassis pan response over a range of load transfer point spacing. The chassis pan, although stiffer than the motherboard, is still relatively compliant. Its flat, thin, non-embossed construction allows it to deflect easily under load. When loads are transferred into its center region, the pan will undergo substantial deflection under the heatsink load. As a result, the effective stiffness of the alternate load path is weak and does not adequately protect the board and MCH.

When a stiff interposer element such as an SRM is used to bridge the loads out to the edges of the pan, the pan deflection is small and the effective stiffness of the alternate load path is adequate to minimize motherboard flexure. The SRM attach points to the chassis that are defined in the *BTX Interface Specification* were selected based on these considerations.

Figure 166: Chassis Pan Stiffness Sensitivity Load Transfer Point Spacing
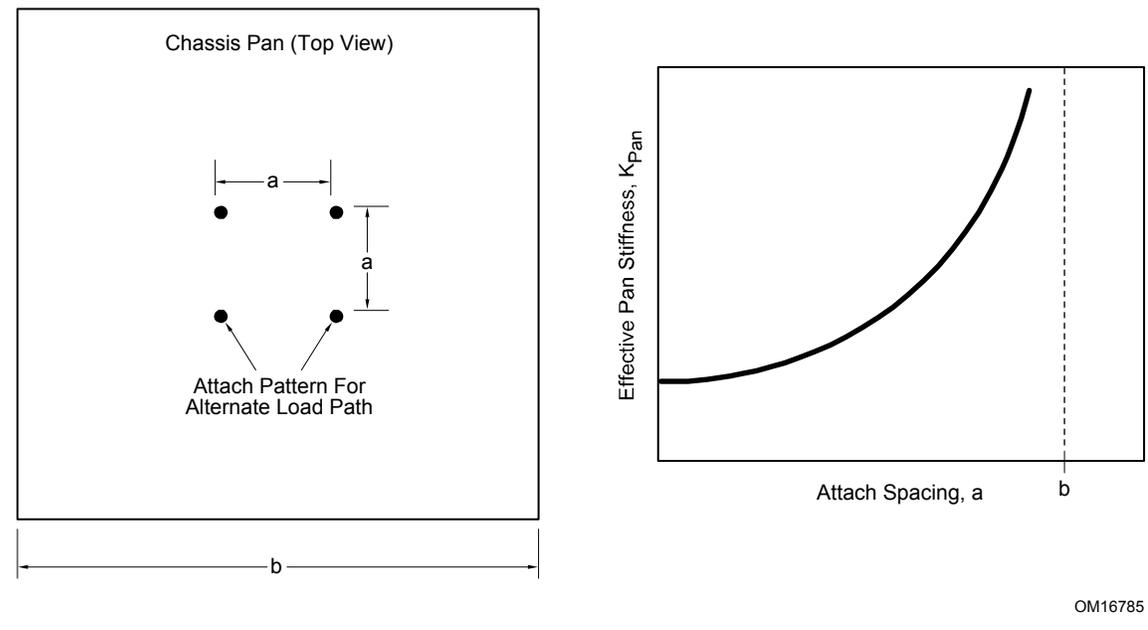
OM16785

4.6.1.2 SRM Stiffness

In addition to load transfer points, the stiffness of the SRM or interposer element is also important in determining the effectiveness of the alternate load path. Sheet-metal interposers with high cross-sectional moments of inertia will provide high-stiffness load paths.

Figure 167 shows the effect of SRM stiffness on load point deflection. In this case, the stiffness sensitivity is derived from the parallel spring model presented in Figure 162. The equation governing the load point deflection is given by:

Equation 77: Deflection in Parallel Spring Model

$$\Delta = \frac{P}{k_{MB} + k_{SRM}}$$

This model assumes that the SRM attach pattern is broad enough to make the chassis pan deflection is small. If the chassis pan deflection is significant, it should be added as a third spring in parallel. It is important to note that as the chassis load transfer points are spread out (to help decrease chassis pan deflection) the increased length and width of the SRM will decrease its bending rigidity (see Section 4.4.1.1.2). Deep emboss features can be used to compensate for this decrease by effectively increasing the moment of inertia and gain the desired level of bending rigidity.

The target value for SRM stiffness is available in the SRM Requirements Document, available separately at www.formfactors.org.

Figure 167: Deflection Sensitivity to SRM Stiffness
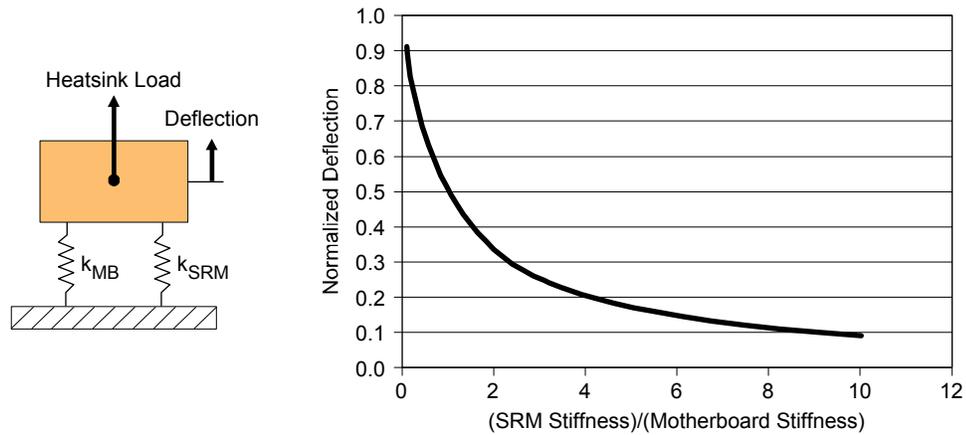
OM16786

4.6.2 Case Study – Applied Engineering for 12.9L Reference SRM Design

The 12.9L system reference design uses an SRM to transfer CPU heatsink inertial loads to the edges of the chassis. This serves to protect the motherboard (and the MCH) by reducing the load carried by the board in a shock condition.

The effectiveness of this solution approach can be illustrated by contrasting it against a simpler attach scheme in which the heatsink mass is attached directly to the motherboard. Figure 168 shows such a configuration. In this case, the heatsink generates 800 N of inertial load in shock. Figure 169 shows a series of free body diagrams for the heatsink and motherboard. That heatsink load is imparted directly to the motherboard. In turn, the motherboard bridges the load out to its chassis mounting points. The board undergoes severe bending as a result of this loading.

Figure 168: Shock Configuration without SRM
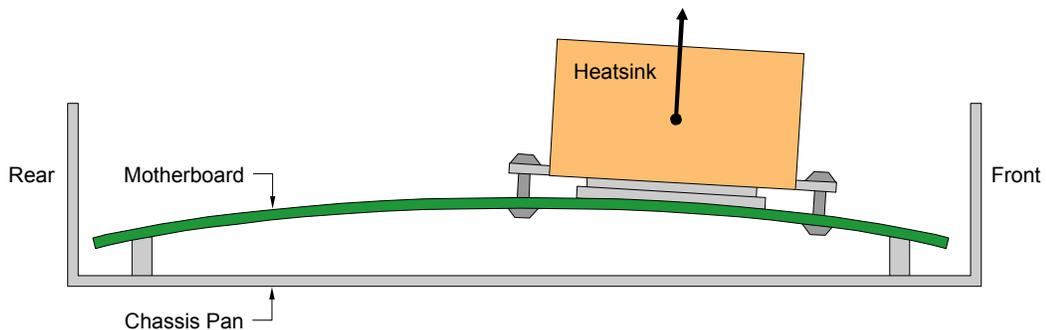
OM17105

Figure 169: Free Body Diagrams for Shock Configuration without SRM
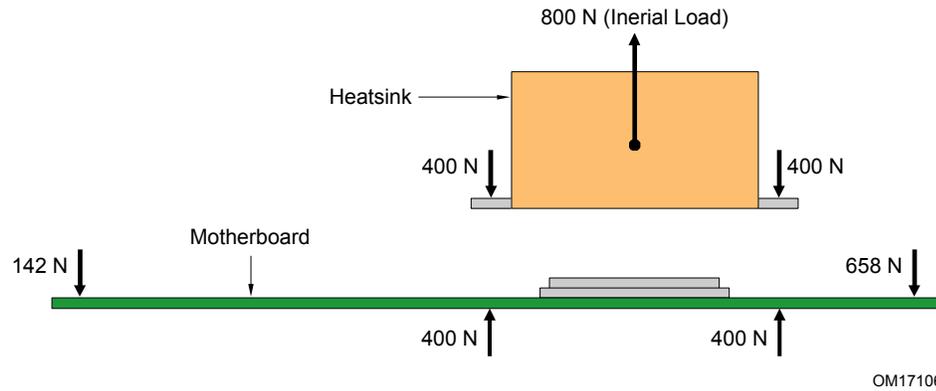
By adding a load transfer plate to the system, the motherboard can be shielded from the heatsink load. The SRM is positioned between the motherboard and chassis as shown in Figure 170. The SRM is constructed of deeply embossed steel, giving it a bending stiffness several times that of the motherboard. Figure 171 shows a schematic of the assembly. The thermal module attaches to both the motherboard and SRM. Figure 172 shows a series of free body diagrams for the heatsink, motherboard and SRM. Finite element analysis of the system was performed to derive the reaction loads for each component. The results show that the heatsink inertial load is transferred primarily into the SRM. The resulting motherboard loads are reduced by a factor of three or greater relative to the configuration shown in Figure 168.

Figure 170: Exploded Assembly Showing SRM Placement in 12.9L Reference System Configuration

Thermal Module Assembly
• Heatsink and Fan
• Clip
• Structural Duct

Motherboard

SRM

Chassis Pan

OM17107
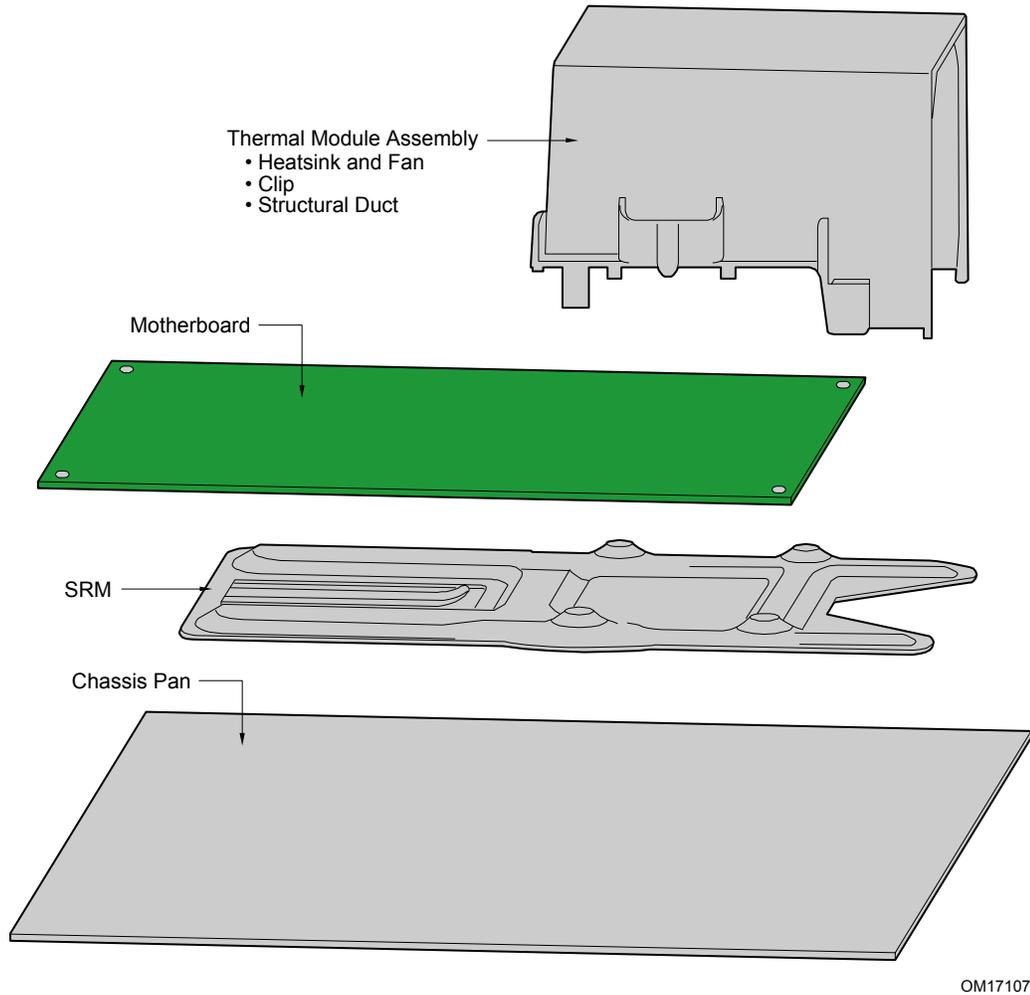
Figure 171: Shock Configuration with SRM Used to Support Heatsink Load

Heatsink

Rear

Motherboard

Front

SRM

OM17108
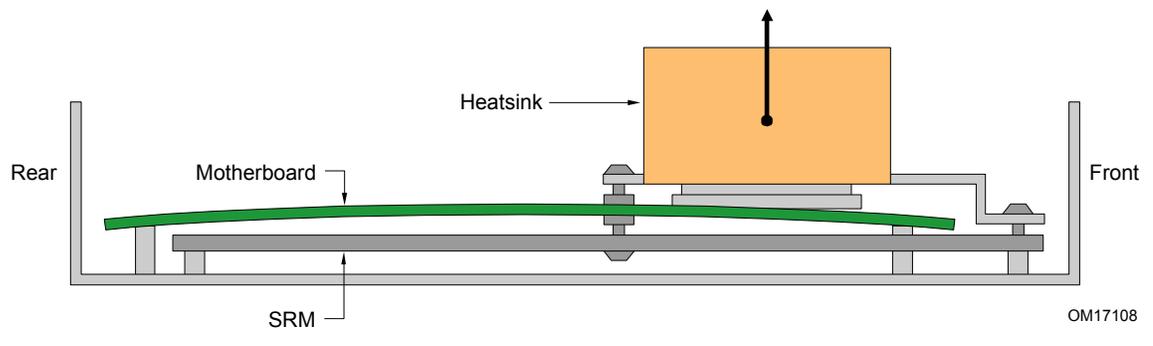
Figure 172: Free Body Diagrams for Shock Configuration with SRM

800 N (Inertial Load)

Heatsink

540 N

260 N

Motherboard

40 N

100 N

140 N

SRM

210 N

190 N

400 N

OM17109
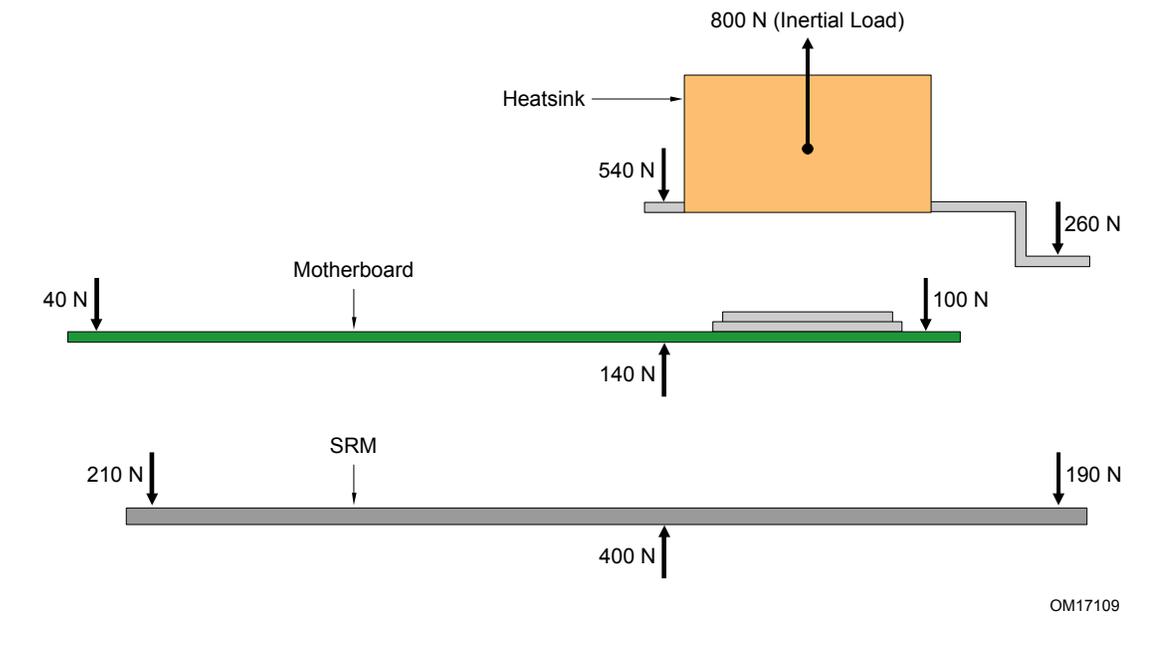
4.6.3 Case Study – Applied Engineering for MCH Protection in With 12.9L Reference SRM Design

The stress levels in the MCH BGA in a shock condition are dependent on many parameters including CPU heatsink mass, input acceleration, dynamic amplification factor, etc. Two parameters of particular interest in this case study are the position and orientation of the MCH on the motherboard. As the MCH position changes for different board layouts, it may shift to locations that have higher or lower board flexural response in shock, and therefore induce more or less severe stress levels in the BGA solder joints. Likewise, as the MCH orientation changes, the BGA pattern may become more or less aligned with primary directions of board curvature, which again, would impact the severity of the solder joints stresses generated in shock.

To investigate this sensitivity, finite element analysis was conducted for a range of MCH positions and orientations. The range is shown in Figure 173. The range was selected to represent practical positions and orientations for the MCH on a typical BTX motherboard.

The structural analysis considered a system-level, +z shock condition (board facing down on shock table). This is generally the worst-case shock orientation for MCH solder joint stress conditions. The finite element model accounted for the primary inertial load contributions and structural load paths generated from the CPU thermal module, motherboard, SRM and chassis pan. The +z shock condition generates bending in the motherboard resulting from the inertial load of the CPU heatsink. This severity of the bending is substantially reduced by use of the SRM, however, some board bending does occur, and it must be assessed for impact to MCH BGA solder joints.

Figure 174 shows the motherboard displacement contour plot. Notice that the MCH is generally located near the center of the board, which coincides with the region of peak board deflection. This region is subjected to substantial board flexure which serve to generate stress in the solder joints of any components mounted in the region.

The 12.9L reference system has the MCH located close to position A and 30-deg orientation (see Figure 173). This represents the baseline configuration for the case study. The solder ball forces were extracted from the finite element analysis to determine the worst cast solder joint force for the +z shock condition. Results showed maximum joint forces to occur on the corner balls. Those forces were compared to a failure threshold to determine strength margin against solder joint fracture. Analysis results showed a margin of 32% in this load condition.

As the position and orientation of the MCH were adjusted, the impact to solder joint protection was evaluated by comparing the solder joint strength margin for the various configurations. Figure 175 summarizes the results. Notice that position B has the least margin at 7%. There is moderate sensitivity to position within this region of movement with the predicted strength margin ranging from 7% to 32%. The sensitivity to orientation is weaker, with the predicted strength margin ranging from 20% to 33%. The 0-deg orientation shows the lowest margin, while the 30-deg and 45-deg offer slight improvements.

It should be noted that in all cases, the highest loaded solder ball in the shock condition is a corner ball. This is significant because all MCH corner balls are sacrificial, i.e. non-functional. Cracks in these sacrificial balls resulting from environmental stress conditions have no impact whatsoever to electrical performance of the system. So actual margin to failure is substantially larger than the values reported in Figure 175. The additional margin to failure for a functional ball is typically on the order of 15% to 20% higher.

In general, the MCH is well-protected in shock over a reasonable range of position and orientations. This protection is due in large part to the performance of the SRM, which significantly limits board curvature in shock.

Figure 173: Range of MCH Movement Considered in MCH Stress Analysis
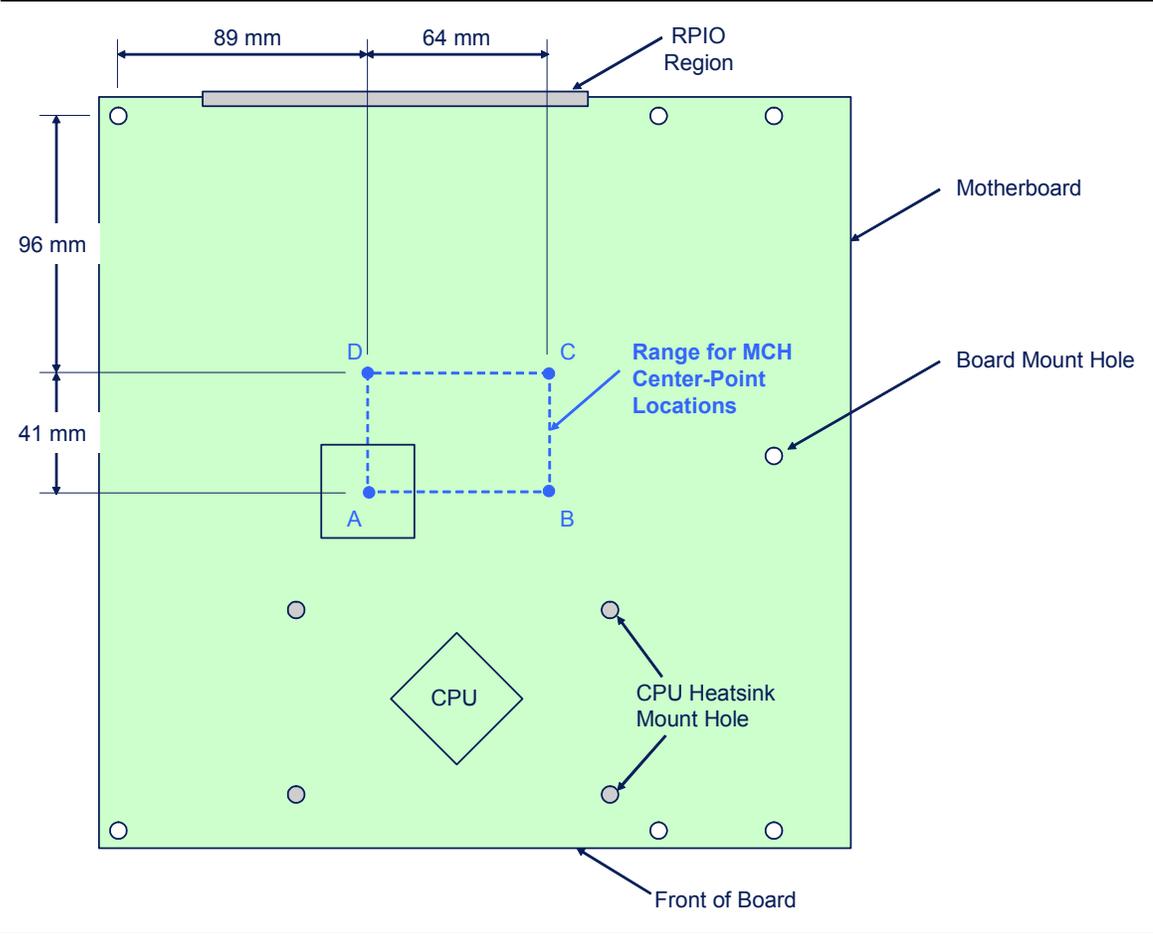
Figure 174: Displacement Contour Plot in +Z Shock Condition
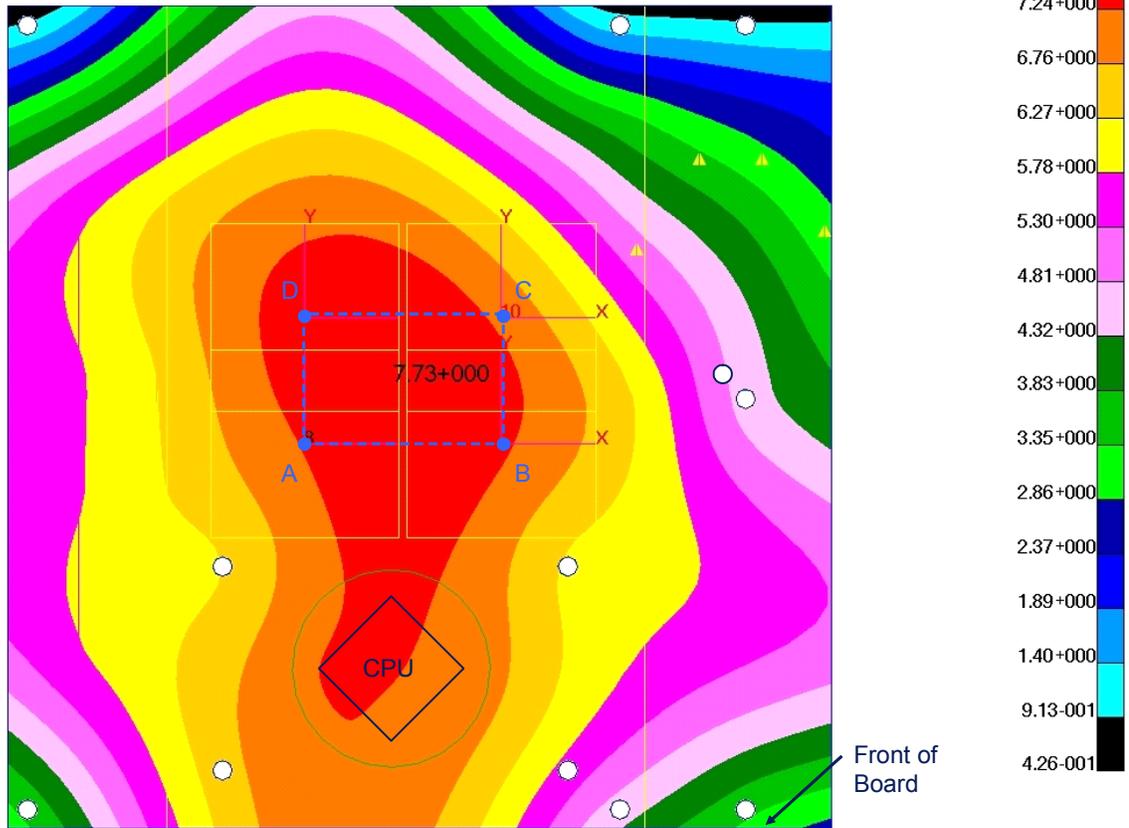
Figure 175: Solder Joint Stress Sensitivity in +Z Shock Condition

Position Sensitivity

Position	Orientation, θ (deg)	Solder Joint Strength Margin
A	0	20%
B	0	7%
C	0	32%
D	0	19%

Front of Board

Orientation Sensitivity

Position	Orientation, θ (deg)	Solder Joint Strength Margin
A	0	20%
A	30	32%
A	45	33%

Front of Board
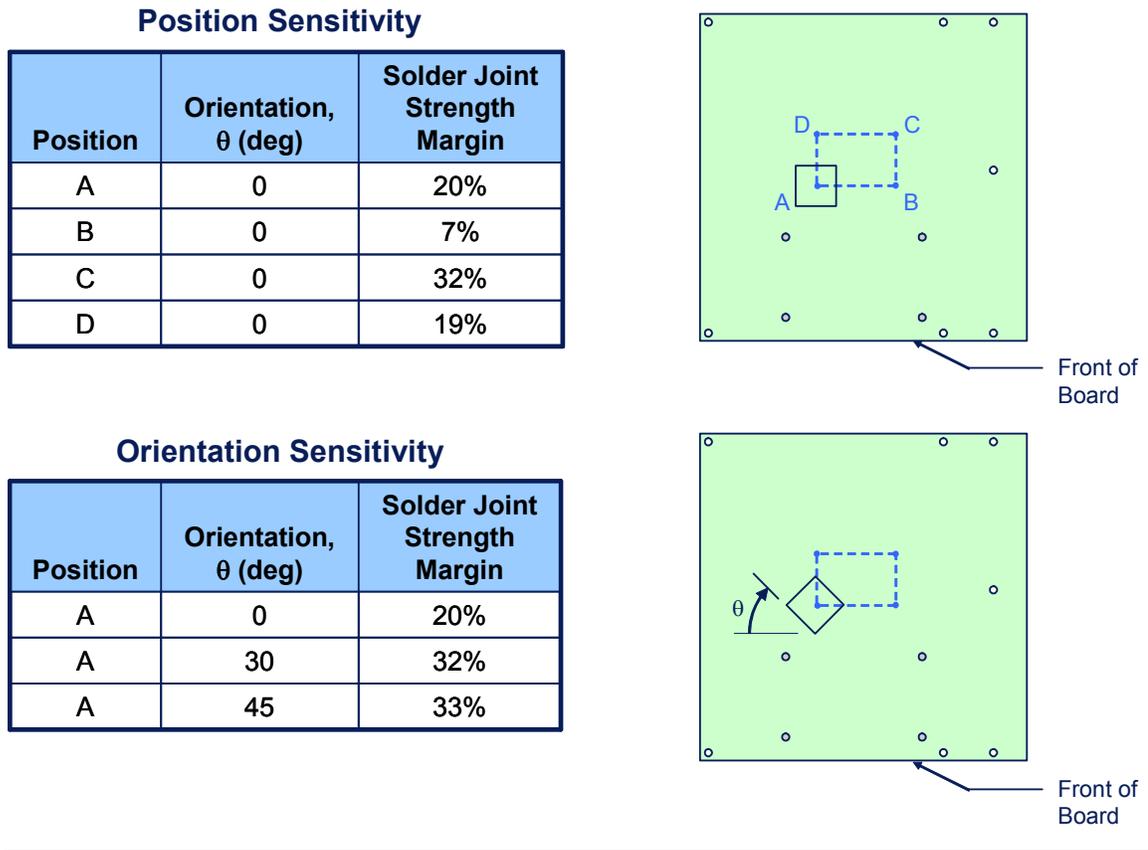
4.7 Socket BGA Protection Strategy

The CPU socket must be mechanically protected to prevent solder ball fracture during shock conditions. Without protection, the processor heatsink inertial loads will generate severe board flexure and may over-stress the socket solder balls.

Several strategies can be employed to protect the socket in this condition. These include the use of alternate load path, top-side stiffening, and preload. Application of the alternate load path strategy was described in Section 4.5.2 as applied to MCH protection. It can be used in the same manner to protect the processor socket. Applications of the top-side stiffening and preload strategies are described in the following sections. A case study for the 12.9L Reference System Design is also presented.

4.7.1 Top-Side Stiffening

The processor socket solder balls can be protected against shock loads by providing stiff overhead structure (typically a heatsink and retention mechanism) in order to constrain the motherboard against flexure within the region of the socket. Section 4.4.3 described the solder ball failure mechanism as it relates to motherboard curvature. The socket and solder ball construction are such that it is most susceptible to solder ball failure in conditions of *upward* board curvature. Preventing this direction of board curvature is key to protecting the socket.

Figure 176 illustrates how the use of topside stiffening elements can reduce upward board curvature in a shock condition. In the unsupported case, the entire board suffers upward curvature. For the case in which the socket is supported by topside stiffening, notice that the local region under the socket remains relatively flat while the remainder of the board suffers curvature. The stiff structure above the socket prevents the socket and package from moving upward, and hence reduces the resulting board curvature.

Figure 176: Top-side Stiffened Board Flexure
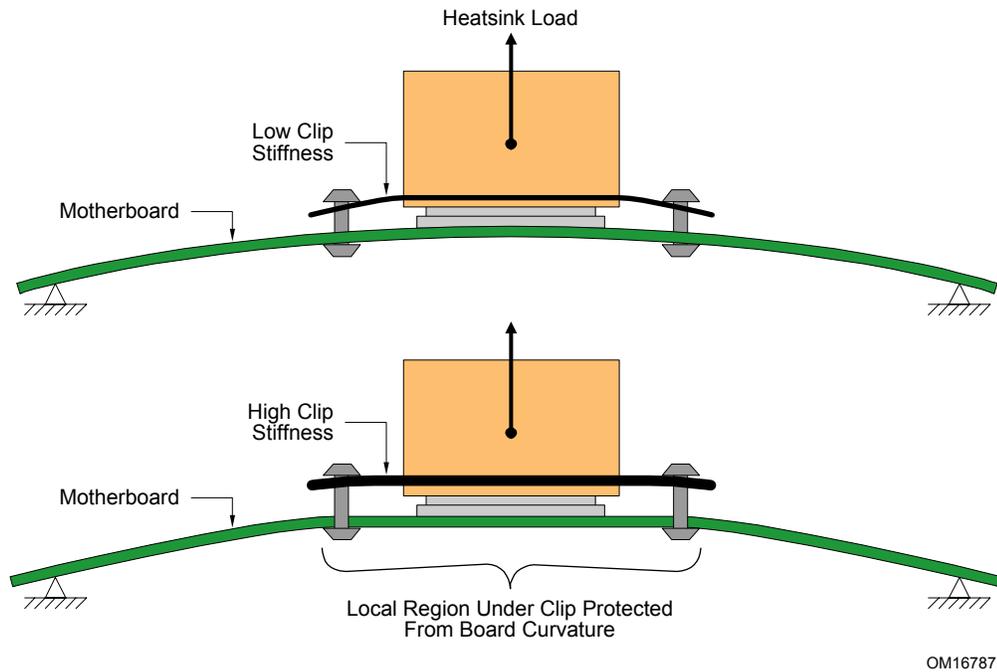
The effectiveness of this strategy is dependent upon the stiffness of the overhead structure. The greater the overhead stiffness, the greater its ability to resist board curvature in shock. Many CPU heatsink retention mechanism designs have migrated from relatively compliance plastic clips to metal clips with embossed features to achieve higher overhead stiffening capability. Figure 177 shows this dependency as a relation between board curvature and clip stiffness. At very low values of clip stiffness, the board flexes as it would in an unsupported condition. But as the clip stiffness increases, the board flexure is reduced. The relationship asymptotes are at zero curvature for a rigid clip.

The target value for Thermal Module Assembly stiffness is available in the Thermal Module Requirements Document, available separately at www.formfactors.org.

Figure 177: Effect of Clip Stiffness on Board Curvature in Shock
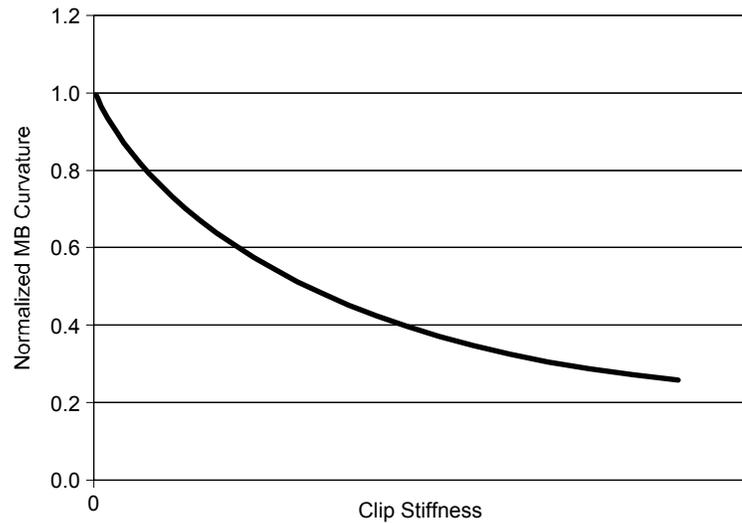
OM16788

4.7.2 Preload

Preload is a second strategy that can be applied to protect socket solder balls. The effect of preload is to place the most susceptible solder balls into compression by applying a static compressive load to the top of the processor package using the heatsink and retention mechanism. Pre-compression of the critical balls effectively delays the onset of solder ball tension load in a shock condition, and thereby protects those balls from fracture.

As described in Section 4.4.3, the inherent flexural response of the motherboard and socket generate the highest BGA loads in the corner balls of the array, and more generally, through every ball in the outer ring. These critical solder balls require protection in a board flexure condition. Figure 178 shows how preload can be used to provide this protection.

Figure 178: Pre-load Board Flexure

Heatsink Shock Load

Heatsink
Motherboard

A. Motherboard Without Preload.

B. Motherboard in Shock Condition, Processor Socket Region Undergoes Upward Curvature.

Region of Preload Effect

C. Apply Preload to Induce Beneficial Downward Board Curvature.

Heatsink Shock Load

Preload Delays Onset of Upward Curvature

D. In Shock Condition, the On-Set of Upward Curvature is Delayed Within the Preloaded Region.

OM16789

When static preload is applied to the processor, the motherboard is forced to undergo slight downward curvature. As a result, the critical corner balls are placed into compression. As the shock load is applied, the board attempts to curve upward, but it must first overcome the preload and the initial downward curvature induced by the preload. In this way, the preload has delayed the onset of upward board curvature and thereby reduced the maximum tension load placed on the critical solder ball during shock.

Note that the preload magnitude may be limited by the maximum allowable static compressive loading for the component under consideration. Refer to the EMTS for static load limits on the processor and the socket Design Guide for limits on the socket.

The target value for Thermal Module preload is available in the Thermal Module Requirements Document, available separately at www.formfactors.org.

4.7.3 Combined Top-Side Stiffening and Preload

The strategies of topside stiffening and preload can be used in tandem. When preload is used to pre-curve the board downward, and thus delay the on-set of upward board curvature, topside stiffening can be used to reduce the maximum upward curvature. Figure 179 shows this effect as it plots board curvature against clip stiffness for a range of preload values.

These two distinct strategies work well together to manage board curvature in the shock condition. One strategy can be used to supplement the other. For example, the maximum allowable static compressive load for the processor or socket limits preload, the clip stiffness can be increased to provide the additional protection needed. Or when volumetric design constraints limit the achievable stiffness for the clip, preload may be added to supplement the socket protection.

Figure 179: Board Curvature versus Clip Stiffness
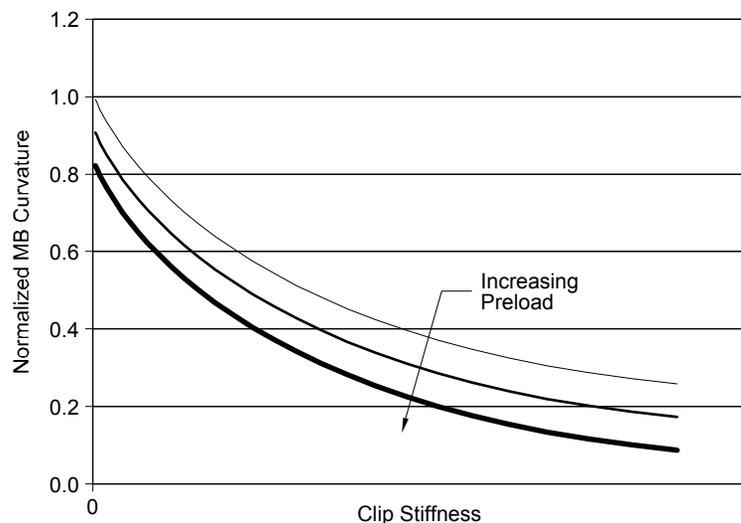
4.7.4 Case Study – Applied Engineering for LGA775 Protection

The stress levels in the LGA775 BGA in a shock condition are dependent on many parameters including CPU heatsink mass, input acceleration, dynamic amplification factor, etc. One parameters of particular interest in this case study is the orientation of the LGA775 socket on the motherboard. Since the position of the LGA775 socket is somewhat tightly constrained (see the Intel® Platform Design Guide for details), the positional sensitivity was not analyzed. As the LGA775 orientation changes, the BGA pattern may become more or less aligned with primary directions of board curvature, which would impact the severity of the solder joints forces generated in shock.

To investigate this sensitivity, finite element analysis was conducted for LGA775 orientations perpendicular to the board edge (0 degree) and at an angle of 45 degrees to the board edge. The BTX Grantsdale RVP uses the 45 degree orientation.

The structural analysis considered a system-level, +z shock condition (board facing down on shock table). This is generally the worst-case shock orientation for LGA775 solder joint force. The finite element model accounted for the primary inertial load contributions and structural load paths generated from the CPU thermal module, motherboard, SRM and chassis pan. The +z shock condition generates bending in the

motherboard resulting from the inertial load of the CPU heatsink. This severity of bending is substantially reduced by use of the SRM. However, some board bending does occur, and it must be assessed for impact to LGA775 solder joints.

Figure 174 shows the motherboard displacement contour plot. Notice that the LGA775 is generally located away from the center of the board and is, therefore, outside the region of peak board deflection. However, the LGA775 region is subjected to an asymmetrical board flexure that may generate different force states depending on the socket orientation.

The relative solder joint margin was evaluated by comparing the solder joint strength margin for the two orientations investigated. In the Intel 12.9L reference design system, the LGA775 solder joint force state had significantly greater margin in the 0 degree orientation.

Since LGA775 solder joint protection is offered through a combination of preload, thermal module stiffness, and SRM stiffness, the preload stiffness requirements have been provided for the worst case LGA775 orientation of 45 degrees. Thermal module and SRM components designed to meet the requirements outlined in their respective Requirements Documents will provide adequate LGA775 protection in all socket orientations.

5. EMC Design Considerations

Electromagnetic compatibility (EMC) considerations have been made to the BTX specification and are illustrated in the Intel reference designs. The purpose of this section is to highlight the considerations that are either part of the *BTX Interface Specification version 1.0a* or related to containment techniques that have been used in the BTX reference designs.

5.1 Rear Panel I/O Design

The *BTX Interface Specification version 1.0a* defines the size of the chassis cutout, surrounding keepout areas, and I/O connector protrusion area to insure board to chassis compatibility and metal to metal contact. The specification does not place requirements for the type or orientation of I/O connectors. The rear panel I/O design shown in Figure 197**Error! Reference source not found.** is derived from the *BTX Interface Specification* and includes edge geometries used in previous form factors to be known to provide sufficient contact for EMC requirements.

5.2 Chassis EMC Design Considerations

The BTX Specification reserves grounding features on the chassis to be used if needed. The position and location of these features was determined using the assumption that the potential risk of standing waves for frequency ranges below 1GHz is increased when the wavelength is greater than ½ λ. The recommended grounding span is determined from equation Equation 78.

Equation 78: MI Grounding Span for Frequencies below 1GHz

½ λ = C/2F

Where

λ = Wavelength

$C = 3 \times 10^8 \text{ m/s}$

$F = \text{frequency } (10^9 \text{ Hz})$

Solving for Equation 78 yields 5.905 inches. From this calculation it is recommended that the grounding span should be ≤ 6 inches to reduce the risks of standing waves. The grounding areas were determined by sweeping a 6 in radius from the motherboard standoffs shown below. The overlapping areas represent the locations for potential chassis to motherboard grounding.

Figure 180: EMC Diagram for BTX Motherboards
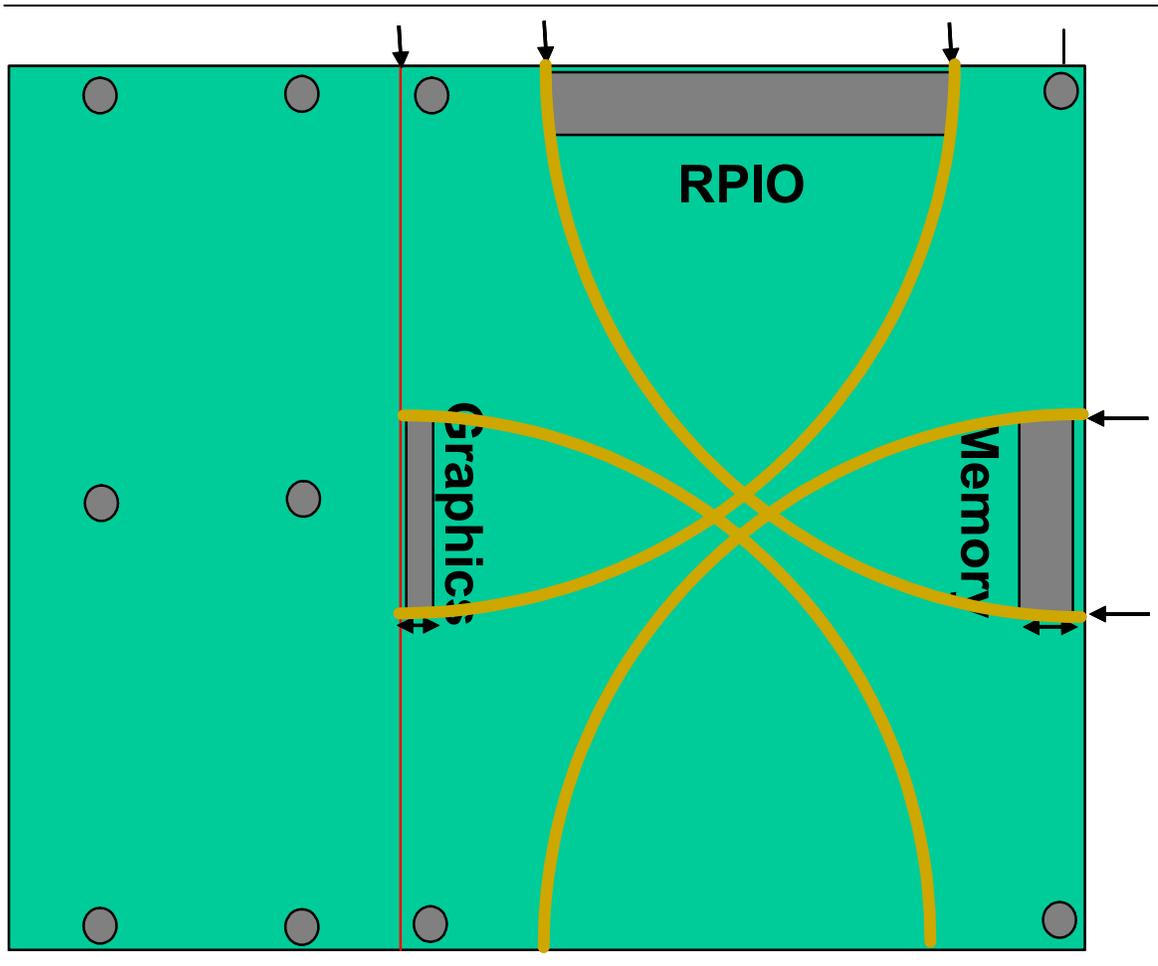
6. System Validation Guidance

6.1 Purpose

The System Validation Guideline provides a description of tests used to assess the Thermal, Mechanical, and Acoustic performance of a BTX system. Each test category includes:

1. The test purpose
2. Recommended sample size
3. Recommended success criteria
4. Recommended test profile, including set-up and EUT configuration

6.2 Definitions and Acronyms

Table 28: Definitions of Terms and Acronyms

Term or Acronym	Definition
AIC	Add-in card (e.g. PCI modem, PCI-Express graphics card, etc.)
Product Configuration	Production level board and system configuration
	Recommendation: Intel Reference Validation Platform (motherboard) populated with memory, AICs, and processor; and installed in final system configuration.
TTB	BTX Thermal Test Board
EUT	Equipment Under Test. The system and supporting peripherals being tested.
FDD	Floppy disk drive
HDD	Hard disk drive
ODD	Optical Disk Drive (i.e. CD-ROM, DVD-ROM, etc.)
TMA	BTX Thermal Module Assembly
TMI	Thermal Module Interface (reference: BTX Interface Specification)
SRM	BTX Support and Retention Module
BTX	Balanced Technology Extended form factor
FSC	Fan Speed Control
Design Condition	The system validation condition at maximum component power loads and maximum external ambient.
Idle Ready	As defined in the Intel Environmental Standards Handbook
Visual Error	An error which does not interfere with the normal operation of the system or cause any loss of data (i.e. read-retries, display flickers, or display bounce); no operator intervention is required to recover.
Intermittent Error	An error that results in permanent corruption of data or requires operator intervention to recover, such as a system hang or unplanned reset of the system.
Hard Failure	An error that results in permanent changes in electrical characteristics or permanent component damage. Any functional failure that requires repair.
Functional Failure	An error that results during the execution of the system functional test. As defined in this document, a functional failure would be an error reported during execution of WinMTA.

6.3   Reference Documents

Table 29: Reference Documents

Reference Material	Location	Description
WinMTA (Functional Test Software)		Comprehensive system level functional test software
Intel Pentium 4 Processor in the 775-Land LGA Package Thermal Design Guidelines	Intel Yellow Cover Document 15058.	Power and temperature requirements, plus instrumentation guidance.
Grantsdale Chipset GMCH Thermal and Mechanical Design Guidelines	Intel Yellow Cover Document # 15743	Memory Controller Hub Chipset power and temperature requirements, plus instrumentation guidance.
Intel® I/O Controller Hub 6 (ICH6) Thermal Design Guidelines	Intel Yellow Cover Document # 15451	I/O Controller Hub Chipset power and temperature requirements, plus instrumentation guidance.
Thermal Readiness for Platform Compatibility Guide '04 B Platforms	Intel Yellow Cover Document # 15913	Processor Voltage Regulation requirements, validation, and instrumentation guidance

6.4   Validation Overview

6.4.1 Test Article Baseline Configurations

Validation activities will use two system test article configurations:

1. The **"PRODUCT CONFIGURATION"** uses a functional Intel® *Reference Validation Platform (RVP)* motherboard or its equivalent in the EUT. The Intel RVP has been designed to be compliant with Intel® Platform Compatibility Guide '04 B guidelines. In any case, system validation requires that the PRODUCT CONFIGURATION be equivalent to a production board in terms of its configuration and performance. Table 30 lists the default system bill of materials for the Intel 2004 BTX RVP configuration.

    Intel suggests that the production or pre-production board of interest be substituted for PRODUCT CONFIGURATION.

2. The **"TTB"** configuration uses a *BTX Thermal Test Board*. TTB is used only for system and Thermal Module Assembly (TMA) thermal performance assessments. Table 31 lists the default system bill of materials for the Intel 2004 TTB configuration.

    The Intel TTB will include surface-mounted MCH and ICH TTVs, LGA775 socket, and memory and AIC connectors for load cards.

Table 30: Intel 2004 RVP System Description

Description	Comments
Base Unit	
Chassis	12.9 liter Slim Tower sheet metal chassis assembly, including cosmetic covers
Bezel	Plastic front bezel
Motherboard	PRODUCT CONFIGURATION
Memory slot #1-4	Recommendation: SDRAM DIMM, 512 MB, DDR2-533
MCH heatsink	Recommendation: Grantsdale BTX MCH heatsink assembly, with TIM
Processor	Intel-provided LGA775 characterized processor
Processor Thermal Module	Recommendation: Type I or Type II BTX LGA775 TMA
Processor heatsink	Included in TMA
System fan	Recommendation: 90mm for Type I TMA, 70mm for Type II TMA Recommendation: TMA fan housing with stator
Duct	Included in TMA
Power supply	Recommendation: ATX12V, CFX12V, or LFX12V
Power supply fan	Included in PSU
Peripheral loads:	
ODD bay	Recommendation: DVD+R/+RW (S-ATA)
FDD bay	
HDD bay	Recommendation: 7200RPM, S-ATA, fluid bearing
Add-in card loads:	
Riser Card	A riser card is recommended only for narrow systems where a required Full Height AIC will not fit.
Card slot #1	Recommendation: PCI-Express x16 (Use this card slot for a riser card, if necessary)
Card slot #2	Recommendation: PCI-Express x1 (Use this card in the lower riser slot, if necessary)
Card slot #3	Recommendation: PCI
Card slot #4	Recommendation: PCI-Express x1

Table 31: Intel 2004 TTB System Description

Description	Comments
Base Unit	
Chassis	12.9 liter Slim Tower sheet metal chassis assembly, including cosmetic covers
Bezel	Plastic front bezel
Motherboard	BTX Thermal Test Board (TTB)
Memory slot #1-4	Resistive load DIMM simulator
MCH heatsink	Recommendation: Grantsdale BTX MCH heatsink assembly, with TIM
Processor	CPU Thermal Test Vehicle (TTV)
Processor Thermal Module	Recommendation: Type I or Type II BTX LGA775 TMA
Power supply	Recommendation: ATX12V, CFX12V, or LFX12V
Peripheral loads:	
ODD bay	Load simulator
FDD bay	Load simulator
HDD bay	Load simulator
Add-in card loads:	
Card slot #1	AIC, resistive load simulator, FH PCI-Express x16
Card slot #2	AIC, resistive load simulator, FH PCI-Express x1
Card slot #3	AIC, resistive load simulator, PCI or LP-PCI
Card slot #4	AIC, resistive load simulator, PCI or LP-PCI

6.4.2 TTB Details

6.4.2.1 Powering Load Simulators and Test Vehicles

It is recommended that external bench power supplies be used to provide and control power to the following TTB components: CPU, MCH, ICH, CPU Voltage Regulation, DIMMs, TMA fan, and PSU fan.

6.4.2.2 Fan Speed Control (FSC)

FSC must be simulated during TTB testing since the FSC circuit will not have input from the processor diode. The processor TTV has neither a diode to report the processor temperature nor a T_Control value that can be used as a reference by the FSC circuit.

Therefore, it may be necessary to estimate the appropriate fan speed settings for each validation test condition based on either CFD analysis or PRODUCT CONFIGURATION testing. It may be also necessary to manually control the TMA and PSU fan speeds. The thermal validation procedures outlined in this document require fan speed operation at the minimum and maximum fan speed settings; however, it may be important to validate thermal performance for intermediate fan speed settings.

6.4.3 Functional Test Software

Unless otherwise noted, the latest version of WinMTA (running on Windows* XP SP1) is used to functionally exercise and test all supported system product functions: memory, CPU, peripherals, parallel port, serial port(s), audio, video, etc.

6.5 Thermal Validation

6.5.1 System Temperature Compliance – Idle Ready Operating Condition

6.5.1.1 Purpose

To demonstrate that the system meets all subsystem temperature requirements in the Idle Ready Operating mode, which is typically the mode where all system fans are operating at minimum speed.

6.5.1.2 Sample Size

Validation: Three systems

6.5.1.3 Success Criteria

Table 32: 2004 Performance System - Recommended Idle Power Loads and Success Criteria

Component	Simulation Target	Applied Power (Watts)	Success Criteria
Processor	Performance Target[1]	0.5 * TDP	T_case[1]
CPU Voltage Regulation	Performance Target[1]	0.5 * TDP	91C [2]
MCH	Grantsdale [3]	6.9	Tj = 110C [3]
ICH	ICH6 [4]	1.2	Tj = 120C [4]
Memory	DIMM, DDR2-533	1.6/ea 6.4 total	Consult memory supplier
PCI Express x16 Graphics Card		37.5	Consult graphics card supplier
Subsystem air temperature requirements			
HDD	7200 RPM	8	55C
ODD	DVD/CD-RW	5	45C
FDD	Mobile	0.3	55C
PCI Express x1 Card		3	55C
PCI Cards		3	55C
Board [5]		6	120C
PSU	ATX12V CFX12V LFX12V	0.5 * rated power	50C [6]

[1] Performance Target refers to the Platform Compatibility Guide '04 B guidance for processor power and case temperature are provided in the Processor Datasheet on intel.com.

[2] Maximum temperature of motherboard under the LGA775 socket. Voltage Regulation component temperature requirements are available from component manufacturers.

[3] Guidance for MCH power and temperature requirements are provided in Intel Yellow Cover Document # 15977.

[4] Guidance for ICH power and temperature requirements are provided in Intel Yellow Cover Document # 15451.

[5] Applied power represents efficiency losses in high amperage transmission lines in the motherboard.

[6] Inlet temperature requirement.

6.5.1.4 Test Profile

EUT Preparation

- Instrument the TTB system with thermocouples per Table 33 and Figure 181through Figure 185.
- Prepare TMA.

Table 33: Recommended System Thermocouple Locations

TC #	Description	Location [Axis]
1	External ambient (T_{EXT_AMB})	Laboratory ambient (Outside of system recirculation or inlet pattern)
2	System intake ambient (T_{SYS_INLET})	*Not required* [X]: 5 mm in front of TMI [Y/Z]: At the corners of TMI
3	MCH ambient	*Not required* [X]: 3 mm in front of MCH Heatsink (inlet side) [Y/Z] Centered on MCH HS extrusion width and height
4	AIC ambient	[X]: 3 mm from AIC PCB leading edge [Y]: 10 mm from slot centerline (powered component side) [Z]: @ ½ card height above motherboard
5	ICH ambient	*Not required* [X]: ICH substrate leading edge [Y]: ICH substrate centerline [Z]: 5 mm above motherboard
6	FSC System Monitor	At system monitor thermistor location
7	Memory ambient	[X/Y]: 10 mm behind and halfway between closed levers of J7J1/J7H2 [Z]: 20 mm above motherboard PCB
8	Memory ambient	[X/Y]: Halfway between and 10 mm in front of closed levers of J7J1/J7H2 [Z]: 20 mm above motherboard PCB
9	HDD ambient	[X/Y]: Front left corner of motherboard PCB [Z]: 0 mm below HDD bottom surface
10	FDD amb	[X/Y]: Front left corner of motherboard PCB [Z]: 0 mm above FDD top surface
11	ODD amb	[X]: -95 from inside surface of chassis front subpanel [Y]: +3 mm to the right of ODD bracket [Z]: 0 mm above ODD top surface
12	PSU intake	[X]: 5 mm in front of PSU vent [YZ]: Geometric center of inlet vent pattern

TC #	Description	Location [Axis]
14	PSU exhaust	*Not required* [X]: 5 mm beyond PSU exhaust vent [Y/Z]: Median fan radius or geometric center of vent pattern
15	Rear Vent	*Not required* [X]: 5 mm beyond chassis vent [Y/Z]: Geometric center of exhaust vent pattern
16	Front Vent	*Not required* [X]: 5 mm in front of chassis vent [Y/Z]: Geometric center of vent pattern
17	Processor T_{CASE}	Per Intel Yellow Cover Document
18	MCH T_J	Per Intel Yellow Cover Document
19	ICH T_J	Per Intel Yellow Cover Document

Figure 181: Thermocouple Locations
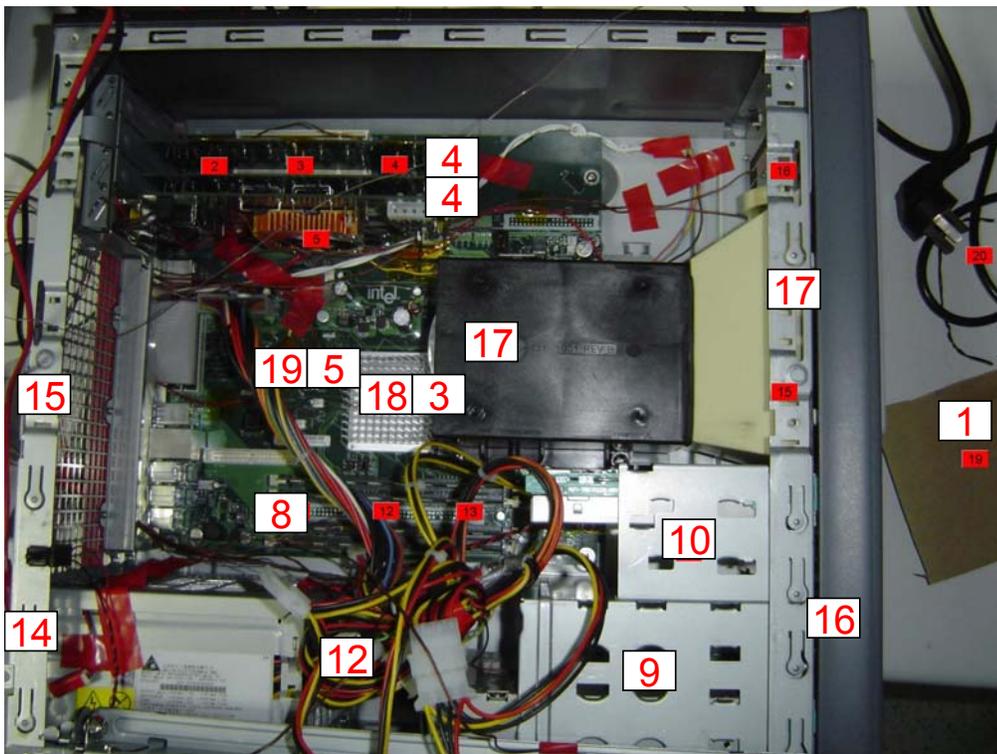
Figure 182: Thermocouple Locations – HDD [Position 11]
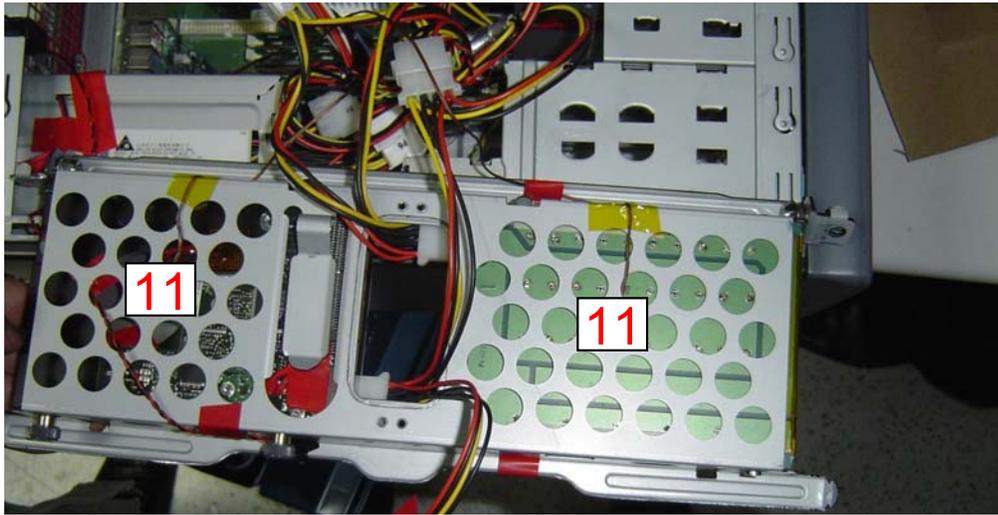
Figure 183: Thermocouple Locations - ICH Ambient [Position 5]
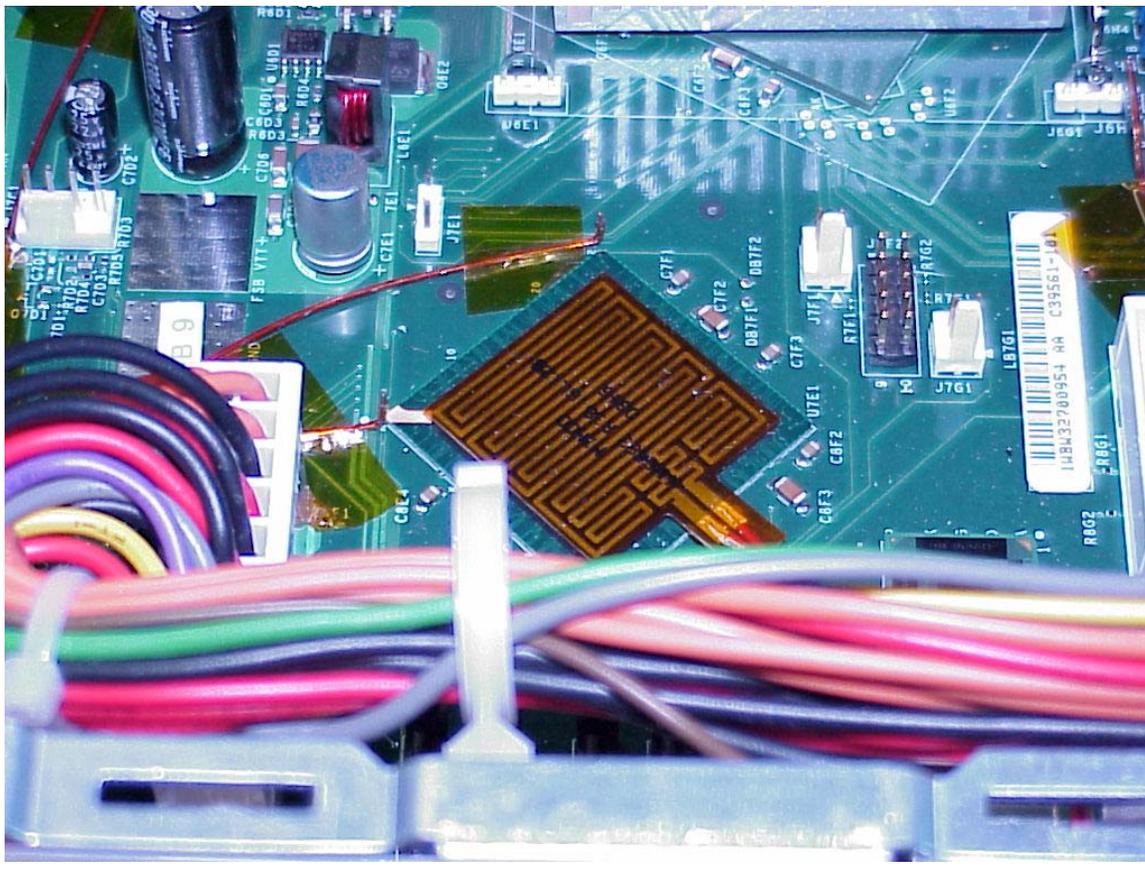
Figure 184: Thermocouple Locations - Memory Ambient [Position 8]
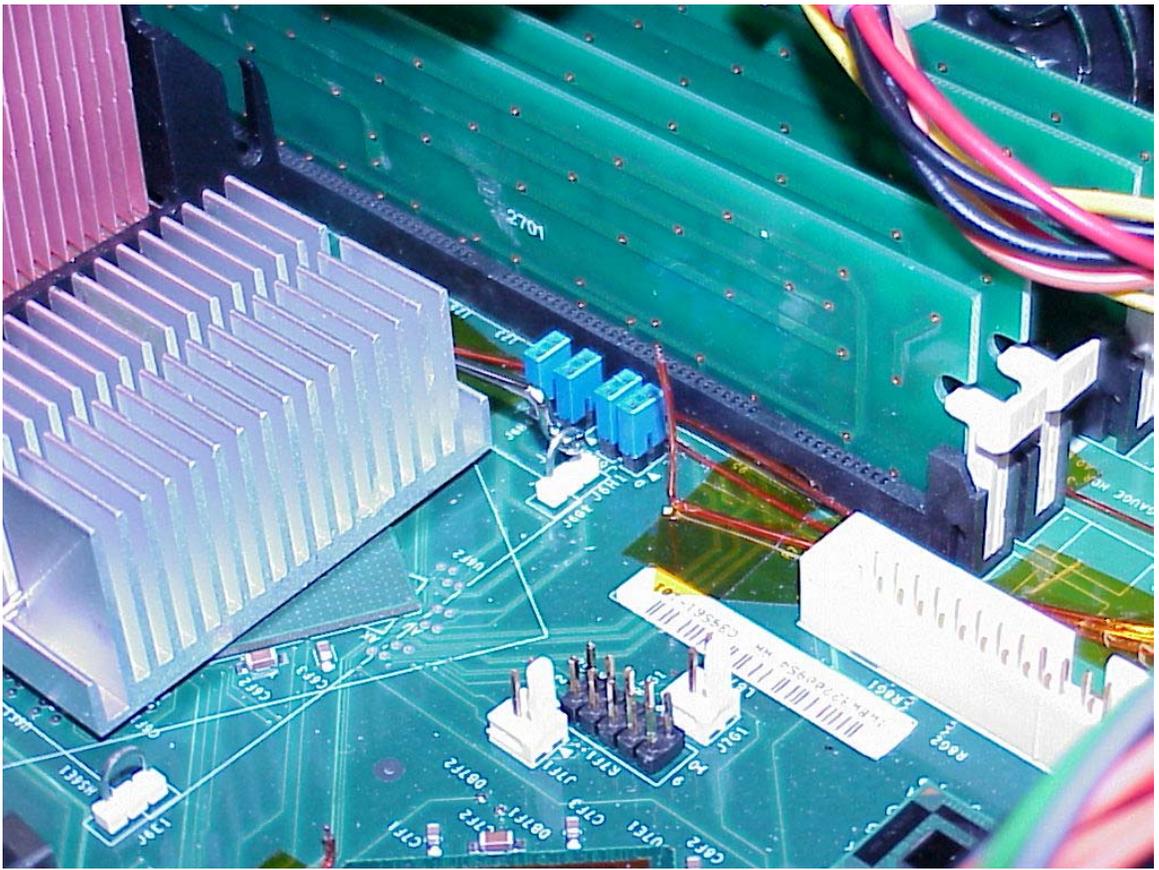
Figure 185: Thermocouple Locations - PSU Exhaust Ambient [Position 14]
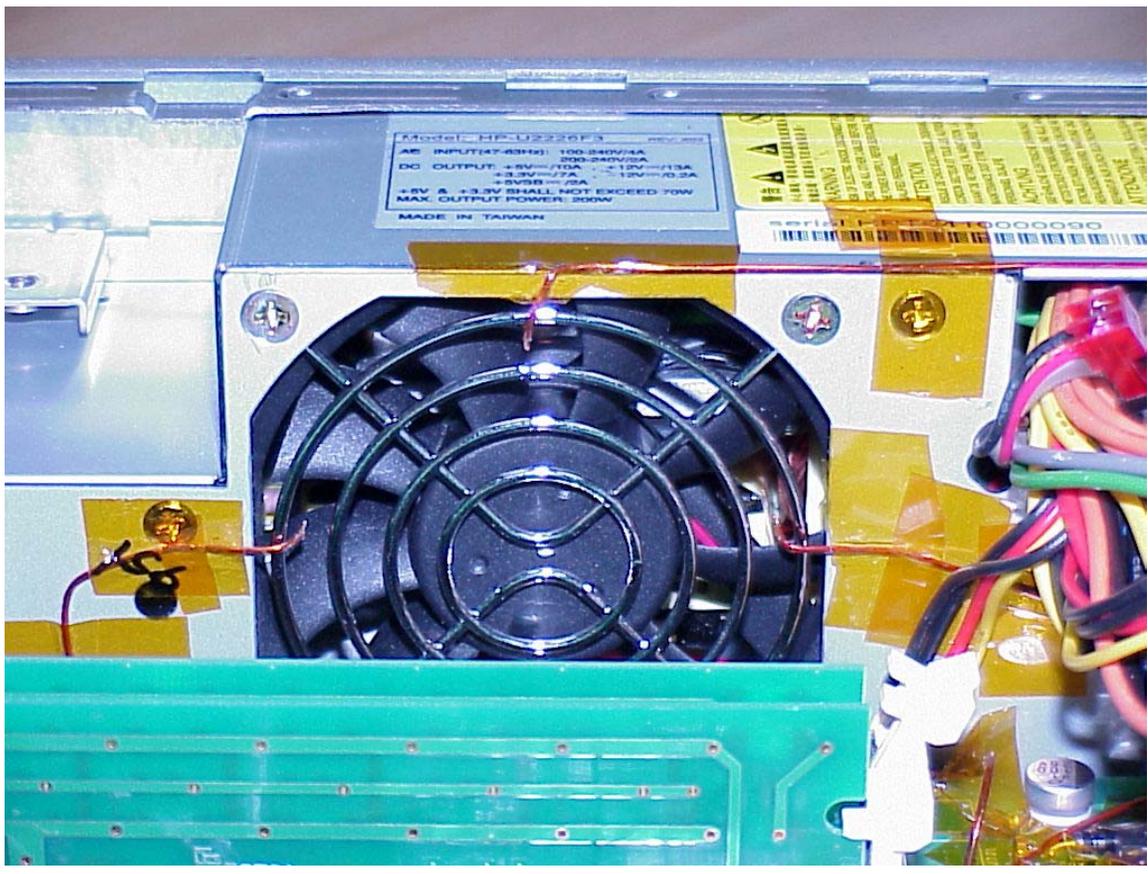
Test Sequence

- Perform all testing at laboratory ambient temperature.
- Apply Idle Power to all components per Table 32.
- Use bench power supply to set TMA and PSU fan speeds to their minimum speed.
- Allow system to reach steady state operating temperature (approximately 30-45 minutes).
- Record the temperature at all measurement locations.
- For each of the three systems, determine compliance with success criteria (see illustration in Table 34):
  - Calculate the difference between the recorded ambient temperature and the defined idle ambient temperature.

    Equation 79: Ambient Offset

    $$\Delta T = T_{ambient,\ Idle} - T_{ambient,\ recorded}$$

  - Calculate the Apply the difference to each recorded system temperature:

    Equation 80: Calculated System Temperature

    $$T_{calculated,i} = T_{recorded,\ i} + \Delta T$$

  - Compare each to the temperature requirement at that location. $T_{calculated,i}$ must be less than or equal to $T_{required,\ i}$

    Equation 81: Comparing Calculated Temperature to Success Criteria

    $$T_{calculated,i} \leq T_{required,\ i}$$

Table 34: Illustration of Compliance Calculation

If measured ambient is 21C but Idle Ready Operating compliance should be assessed at 23C, then Equation 79 yields Δ= +2C

Component		Value at Required Ambient			Success Criteria
	Measured Value	**(+)**	**Compliance Value**		
Processor	50	+2	52		76C
CPU Voltage Regulation	68	+2	70		91C
MCH	88	+2	90		Tj = 110C
ICH	96	+2	98		Tj = 120C
HDD	46	+2	48		55C
ODD	36	+2	38		45C
FDD	47	+2	49		55C
PCI Express x1 Card	47	+2	49		55C
PCI Cards	43	+2	45		55C
PSU	38	+2	40		50C

6.5.2 System Temperature Compliance – Design Operating Condition

6.5.2.1 Purpose

To demonstrate that the system meets all subsystem temperature requirements in the Design Operating mode, which is typically the mode in which all system fans are operating at their maximum operating speeds.

6.5.2.2 Sample Size

Validation: Three systems

6.5.2.3 Success Criteria

Table 35: 2004 Performance System - Recommended Design Power Loads and Success Criteria

Component	Simulation Target	Applied Power (Watts)	Success Criteria
Processor	Performance Target[1]	TDP	T_case[1]
CPU Voltage Regulation	Performance Target[1]	TDP	91C [2]
MCH	Grantsdale [3]	17.7	$Tj = 110C$ [3]
ICH	ICH6 [4]	14.0	$Tj = 120C$ [4]
Memory	DIMM, DDR2-533	3.5/ea 13.9 total	Consult supplier
PCI Express x16 Graphics Card		75	Consult supplier
Subsystem air temperature requirements			
HDD	7200 RPM	10	55C
ODD	DVD/CD-RW	10	45C
FDD	Mobile	0.3	55C
PCI Express x1 Card		5	55C
PCI Cards		5	55C
Board [5]		12	120C
PSU	ATX12V CFX12V LFX12V	Rated Power	50C [6]

[1] Performance Target refers to the Platform Compatibility Guide '04 B guidance for processor power and case temperature are provided in the Processor Datasheet on intel.com.

[2] Maximum temperature of motherboard under the LGA775 socket. Voltage Regulation component temperature requirements are available from component manufacturers.

[3] Guidance for MCH power and temperature requirements are provided in Intel Yellow Cover Document # 15977.

[4] Guidance for ICH power and temperature requirements are provided in Intel Yellow Cover Document # 15451.

[5] Represents efficiency losses in high amperage transmission lines in the motherboard.

[6] Inlet temperature requirement.

6.5.2.4 Test Profile

EUT Preparation

- Instrument the TTB system with thermocouples per Table 33 and Figure 181 through Figure 185.
- Prepare TMA.

Test Sequence

- Perform all testing at laboratory ambient temperature.
    - a. Alternately, the system may be placed in a controlled temperature chamber, where the chamber temperature setpoint matches the Design Ambient.
- Apply Design Power to all components per Table 35.
- Use bench power supply to set TMA and PSU fan speeds to their minimum speed.
- Allow system to reach steady state operating temperature (approximately 30-45 minutes).
- Record the temperature at all measurement locations.
- For each of the three systems, determine compliance with success criteria (see illustration in Table 34):
    - Calculate the difference between the recorded ambient temperature and the defined Design ambient temperature.

Equation 82: Ambient Offset

$$\Delta T = T_{ambient, Idle} - T_{ambient, recorded}$$

- Calculate the Apply the difference to each recorded system temperature:

Equation 83: Calculated System Temperature

$$T_{calculated,i} = T_{recorded, i} + \Delta T$$

- Compare each to the temperature requirement at that location. $T_{calculated,i}$ must be less than or equal to $T_{required, i}$

Equation 84: Comparing Calculated Temperature to Success Criteria

$$T_{calculated,i} \leq T_{required, i}$$

6.5.3 Processor Voltage Regulation Thermal Validation

Please refer to Intel Yellow Cover Document # 15913 for Voltage Regulation thermal validation procedures.

6.6 Acoustics

6.6.1 Purpose

To demonstrate that product with fans, drives (floppy, hard and CD-ROM), or other noise producing components do not produce excessive broadband noise.

6.6.2 Sample Size

Validation: Three systems

6.6.3 Success Criteria

- Published Idle Ready Sound Power ≤ 4.0BA.
    - o Success criteria recommendation is A-weighted in Bels, measured between 100 Hz and 16 kHz. Note: A-weighted means that the measured noise value is corrected to reflect what a healthy human ear hears.
    - o The published sound power level of a system is derived from the measured sound power level of the system and accounts for sound power level variations in measurement repeatability, measurement reproducibility and system production. For this reason, the published sound power level, not the measured sound power level, should be used for the system.
    - o It is recommended that sample population consider sampling into potential variation in acoustic behavior. For example, with two different fan suppliers and two different HDD suppliers, validation testing may require data collection on all four possible combinations of the fan-HDD suppliers.

6.6.4 Test Profile

Acoustics Room and Equipment

- Test room shall provide a free field over a reflecting plane.
- The background ambient noise in the acoustics room (with no systems operating) should be at least 0.6 B quieter than the system under test measured noise values, in each 1/3-octave band between 100 Hz and 16 kHz.
    - o This requirement is based on the measured noise emitted from the system and the room background noise, not the pass/fail criteria.
    - o As systems do not generate significant noise in all 1/3 octave bands, it is often not possible for the test room to meet this criterion across all 1/3 octave bands.
        - o If this occurs, ISO 7779 requires noting in the test report that: "The background noise requirements of ISO 7779 and ISO 3744 have not been satisfied." Typically, this only applies to some of the 1/3 octave bands, not all of them.
    - • Temperature: 23° C, $\pm 2^{\circ}$ C.
    - • Relative Humidity: Between 40% to 70%.
    - • Barometric Pressure: 86 kPa to 106 kPa.
    - • Measurement duration should be at least 30 seconds and shall be the same for all microphones.
    - • Place the microphones specific locations in a hemispheric arrangement around the system at a radius of 1.0 m. (Figure 186, Table 36)
    - • Place the system in the center of the test room
    - • Keyboard and monitor should be located outside the test facility but may be placed beside the system if remote operation is not possible.

Figure 186: EUT and Microphone Placement
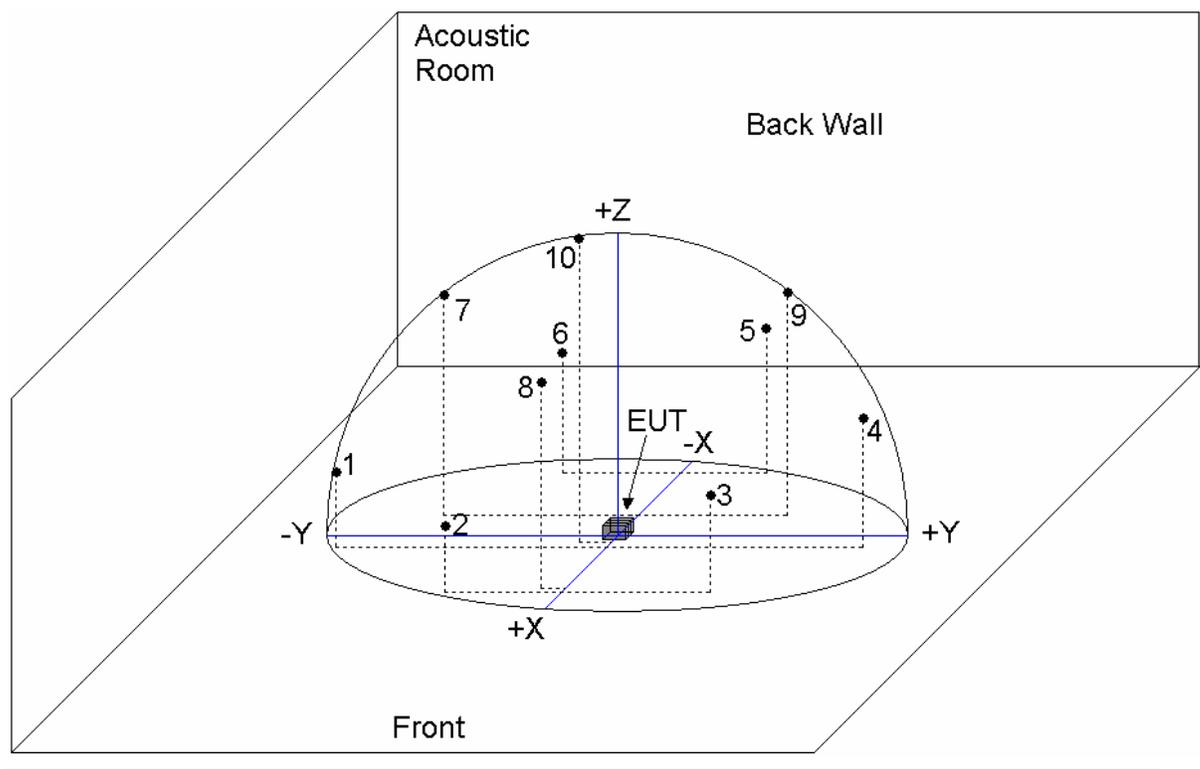
Table 36: Microphone XYZ Location

Position	X	Y	Z	Position	X	Y	Z
1	0.16	-0.96	0.22	6	-0.83	-0.40	0.38
2	0.78	-0.60	0.20	7	-0.26	-0.65	0.71
3	0.78	0.55	0.31	8	0.74	-0.07	0.67
4	0.16	0.90	0.41	9	-0.26	0.50	0.83
5	-0.83	0.32	0.45	10	0.10	-0.10	0.99

All dimensions are in meters.

The origin (0, 0, 0) is the XY center of the system surface coplanar with the test room floor.

System Operational Requirements and Measurements

- Use the Product Configuration (RVP equivalent) system for this test.
- Allow 30 minutes for the system to warm up before collecting acoustic measurements.
- Measurements are made in the Idle Ready mode for each system:
    - Idle Ready: Disks shall be loaded, power on, unit ready to receive and respond to control link commands, with spindle up to speed and read/write heads in track follow mode.

Published Sound Power Noise Values

- The published sound power level of a system is derived from the measured sound power level of the system and accounts for sound power level variations in measurement repeatability, measurement reproducibility and system production. For this reason, the published sound power level, not the measured sound power level, should be used for the system.

- Determine the published noise value from the equations defined below.

Equation 85: Conversion from bels to decibels

$$L_{WA}(dBA) = 10 \times L_{WA}(BA)$$

Where:

$L_{WA}(dBA)$ = Measured Sound Power Level (in dBA)

$L_{WA}(BA)$ = Measured Sound Power Level (in BA)

Equation 86: Average Sound Power (decibels)

$$L_{WAm} = \frac{1}{n}\sum_{i=1}^{n} L_{WAi}$$

Where:

n = 3 (sample size)

Equation 87: Product Standard Deviation

$$S_P = \sqrt{\frac{1}{n-1}\sum_{i=1}^{n}\left(L_{WAi} - L_{WAm}\right)^2}$$

Where

S_P = product standard deviation

Equation 88: Total Standard Deviation

$$S_t = \sqrt{1.5^2 + S_P^2}$$

Where:

S_T = total standard deviation

Equation 89: Published Sound Power (bels)

$$L_{WAd} = \frac{1}{10}\left[L_{WAm} + 1.5S_t + 0.564\left(2.0 - S_t\right)\right]$$

Where:

L_{WAd} = published sound power (in BA), rounded to the nearest 0.1BA

6.7 Temperature Operating

6.7.1 Purpose

To demonstrate that the product will operate successfully at either high or low temperature.

6.7.2 Sample Size

Validation: Three systems

6.7.3 Success Criteria

- No functional test failures during any active functional test period.
- No visual inspection failures.

6.7.4 Test Profile

Chamber Setpoints

- Temperature Operating: +10° C to +35° C
- Temperature Ramp Rate: ≤ 10° C per hour
- Required Performance: +5° C to +40° C
- Number of power cycles: 5

EUT Preparation

- Use the Product Configuration (RVP equivalent) system for this test.
- Perform a visual inspection.
- Perform a functional test.

Test Sequence

- Insert systems into chamber.
- Power up systems.
- Start functional test.
- Set chamber temperature to 25° C.
- Ramp temperature down to 5° C in 2 hours.
- At 5° C:

1. Turn off system power and leave off for 2 hours.
2. Every 90 minutes, turn on system power, restart functional test, then turn off system power.
3. Repeat step 2 for a total of four (4).
- Ramp temperature up to 40° C in 3.5 hours
- At 40° C:
    1. Turn off system power and leave off for 2 hours.
    2. Every 90 minutes, turn on system power, restart functional test, then turn off system power.
    3. Repeat step 2 for a total of four (4).
- Ramp temperature down to 25° C in 1.5 hours.
- At 25° C:
    1. Turn on system power and restart functional test for 3 hours.
- Remove the system from the chamber.
- Perform visual inspection.

Figure 187: Illustration of Chamber Temperature Profile
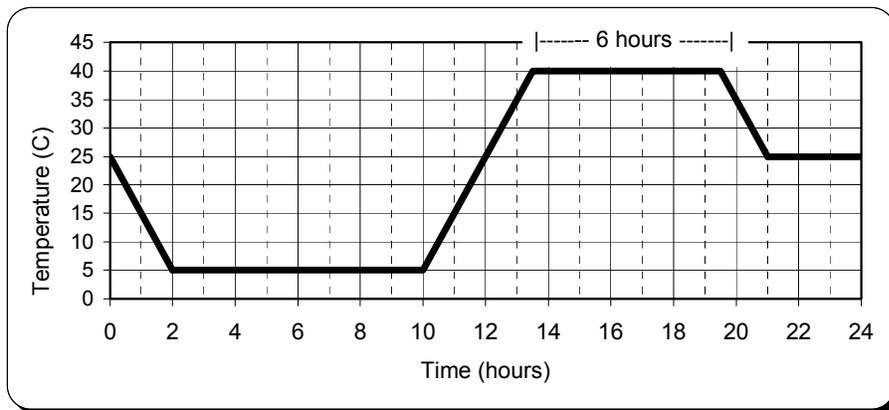
6.8 Mechanical Shock, Unpackaged

6.8.1 Purpose

To demonstrate that the systems can withstand drops from reasonable height that may be encountered in shipping or handling.

6.8.2 Sample Size

Validation: Three systems

6.8.3 Success Criteria

- No functional test failures during any active functional test period.
- No visual inspection failures.

6.8.4 Test Profile

Shock Table Set-up

- Product must be mounted to the shock table so that no cushions or feet attenuate the applied stress.

- Testing will be conducted on a programmable shock table.

- Table drop height and programmer pressure is adjusted to produce the desired delta V and acceleration.

- Table response used to set the desired input shock is to be filtered using a frequency 20 times the fundamental frequency of the input waveform (Equation 90).

Equation 90: Shock Table Frequency

Example: f = (1 / (t x 2)) x 20

- Filter frequency for any waveform data collection is typically 225 Hz.

- The input shock pulse is defined by velocity change and the faired acceleration, (the average of the oscillations between the first time that the response goes negative and the last time the wave goes positive.)

- Rise time is defined as the time required for the waveform to rise from 10% of the faired magnitude to 90% of the faired magnitude and shall be less than 1 ms.

Figure 188: Typical Shock Table Input Waveform

EUT Preparation

- Use the Product Configuration (RVP equivalent) system for this test.
- Perform a visual inspection.
- Perform a functional test.

- Mount the system(s) to the shock table.

    o Support must be under the structural members of the chassis, typically under each of the four corners of the bottom surface. (Figure 189)

Figure 189: Typical Shock Table Fixturing
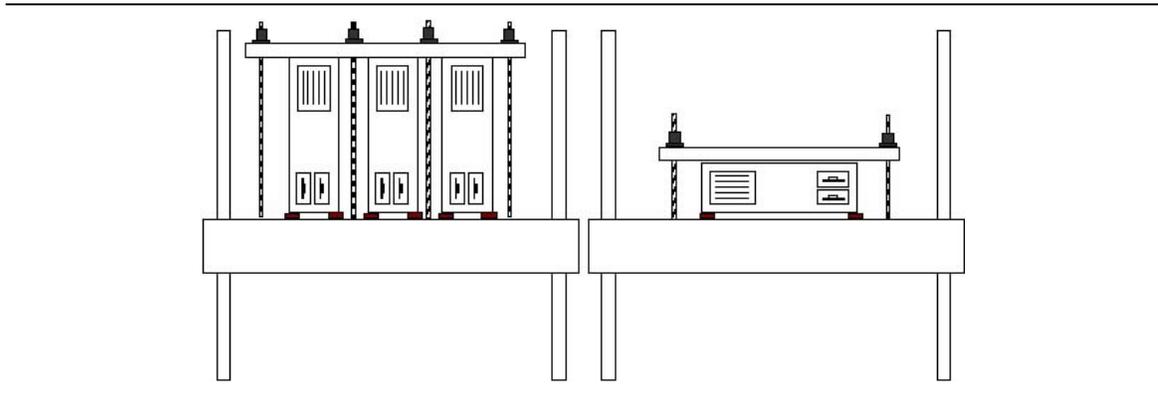
- Conduct two drops on each of the six sides of the system.
- Remove the system(s) from the shock table.
- Perform a functional test.
- Perform a visual inspection.

6.9 Vibration, Unpackaged

6.9.1 Purpose

To ensure the system is sufficiently robust to withstand vibration typically encountered in shipping.

6.9.2 Sample Size

Validation: Three systems

6.9.3 Success Criteria

- No functional test failures during any active functional test period.
- No visual inspection failures.

6.9.4 Test Profile

Vibration Table Setup

- Product must be mounted to the shock table so that no cushions or feet attenuate the applied stress.
- Testing will be conducted on a programmable vibration table.
- Random Profile: 5 Hz @ .001 g^2/Hz to 20 Hz @ 0.01 g^2/Hz (slope up)

    20 Hz to 500 Hz @ 0.01 g^2/Hz (flat)

Input acceleration is 2.20 g RMS

10 minutes per axis in all 3 axes on all samples

Random control limit tolerance is ± 3 dB

EUT Preparation

- Use the Product Configuration (RVP equivalent) system for this test.
- Perform a visual inspection.
- Perform a functional test.

Test Sequence

- Mount the system(s) to the vibration table.

    o Support must be under the structural members of the chassis, typically under each of the four corners of the bottom surface. (Figure 190)

Figure 190: Typical Vibration Table Fixturing
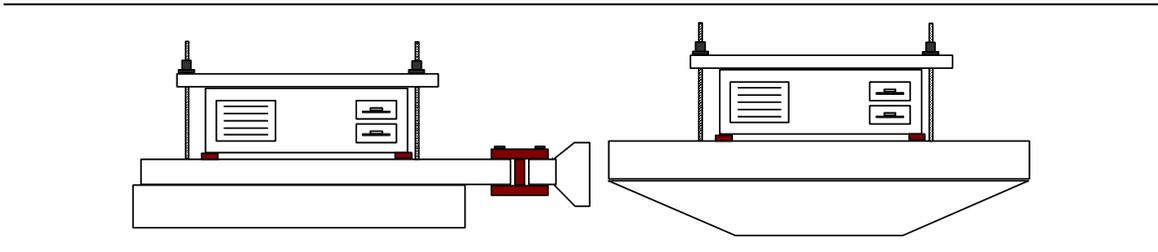
    o Power on system and start functional test.

- Apply Random Profile while system is performing functional test to determine intermittent errors induced by the applied stress.

- Turn off system power.

- Remove the system(s) from the table.

- Perform a functional test.

- Perform a visual inspection.

6.10 Radiated Emissions

6.10.1 Purpose

To investigate whether the level of emissions radiating from a product will interfere with other electronic products. This demonstrates compliance to U.S., European and International legal requirements.

6.10.2 Sample Size

Validation: One system

6.10.3 Success Criteria

Table 37: US and International Radiated Emissions Limits

U.S. (FCC) Radiated Limits for
Radiating Class A and B Devices

Frequency range MHz	Class A 10 meters* limits dB (µV/m)	Class B 3 meters* limits dB (µV/m)
30 to 88	39	40
88 to 216	43.5	43.5
216 to 960	46	46
Above 960	**49	**54

International (CISPR 22) Radiated Limits for
Radiating Class A and B Devices

Frequency range MHz	Class A 10 meters* limits dB (µV/m)	Class B 10 meters* limits dB (µV/m)
30 to 230	40	30
230 to 1000	47	37

*   Quasi Peak detector used $\leq$ 1 GHz
** Average detector used > 1 GHz
Note: The tighter limit shall apply at the band edges

6.10.4 Test Profile

- Use the Product Configuration (RVP equivalent) system for this test.
- The product test setup shall meet the ANSI C63.4: 2001 requirements for U.S. testing, EN 55022: 1994 requirements for European testing, and CISPR 22: 1997 requirements for International testing.
- If the product is tested to the International limits (excluding the U.S.), then the maximum frequency limit for measurement is 1 GHz.
- If the product is tested to the U.S. limits, then the maximum frequency limit for measurement is based on Table 38.

Table 38: Radiated Emissions Maximum Frequency Limit

Highest frequency generated or used in the device or on which the device operates or tunes (MHz)	Upper frequency of measurement range (MHz)
1.705 to 108	1000
≥ 108 to 500	2000
≥ 500 to 1000	5000
≥ 1000	5th Harmonic limited to 40 GHz

Note: The tighter limit shall apply at the band edges

6.11 Radiated Immunity

6.11.1 Purpose

To demonstrate that system products can withstand electromagnetic fields such as those generated by portable radio transceivers (walkie-talkies) or any other devices that will generate continuous wave radiated electromagnetic energy. This demonstrates compliance to European legal requirements.

6.11.2 Sample Size

Validation: One system

6.11.3 Success Criteria

Performance Criteria A, as specified by EN 55024:1998, is used for this test as follows:

"The apparatus shall continue to operate as intended without operator intervention. No degradation of performance or loss of function is allowed below a level specified by the manufacturer, when the apparatus is used as intended. The performance level may be replaced by a permissible loss of performance. If the minimum performance level or the permissible performance loss is not specified by the manufacturer, then either of these may be derived from the product description and documentation and what the user may reasonably expect from the apparatus if used as intended."

6.11.4 Test Profile

Equipment and EUT Setup:

- Use the Product Configuration (RVP equivalent) system for this test.
- The product test setup and immunity test equipment shall meet the IEC 61000-4-3: 1995 specification.
- All testing of equipment shall be performed in conditions as close as possible to installed conditions.
- The power cord of the equipment under test shall be plugged into 230 V, 50 Hz power.
- Wiring shall be consistent with the manufacturer's recommended procedures.
- Equipment shall be in its housing with all covers and access panels in place, unless otherwise stated.
- The system is placed on the front edge of the test table; this should be flush with the vertical uniform electromagnetic field.

Perform Test:

- Use Table 39 to perform the test:

Table 39: Radiated Immunity Test Profile

Frequency Band (MHz)	RF Field Strength (Volts/meter)	Frequency Step Size
80 to 1000	6	4%
80 to 1000	3	1%

Note: Either method may be used for compliance as per the regulation. The 6 V/m method will reduce test time.

- Dwell on each frequency for a minimum of one complete software diagnostic test cycle.

- Test all four sides of the system (face the side under test directly toward the antenna).

- Test in both the horizontal and vertical antenna polarities.

6.12 Power Line Conducted Emissions

6.12.1 Purpose

To investigate whether the level of emissions conducting from a product will interfere with other electronic products. This demonstrates compliance to U.S. (FCC), European (EN) and International (CISPR) legal requirements.

6.12.2 Sample Size

Validation: One system.

6.12.3 Success Criteria

Table 40: Conducted Limits for Class A and B Devices

Frequency range MHz	Class A * QP limits dB (µV)	Class A ** Ave. limits dB (µV)	Class B * QP limits dB (µV)	Class B ** Ave. limits dB (µV)
0.15 to 0.5	79.0	66.0	[2] 66.0 to 56.0	[2] 56.0 to 46.0
0.5 to 5.0	73.0	60.0	56.0	46.0
5.0 to 30.0	73.0	60.0	60.0	50.0

* QP = Quasi Peak   ** Ave. = Average

Note 1: The tighter limit shall apply at the band edges.

Note 2: The limit decreases linearly with the logarithm of the frequency in the range 0.15 MHz to 0.50 MHz

6.12.4 Test Profile

Equipment and EUT Setup:

- The product test setup shall meet the ANSI C63.4: 2001 requirement for U.S. testing and CISPR 22: 1997 requirement for International testing, respectively.

- System diagnostics are to be run on the product during all emission testing.

- The maximum frequency limit for measurement is 30 MHz.

- The EUT input power terminals will be tested at both 120 V, 60 Hz and 230 V, 50 Hz voltage levels.

- Use the Product Configuration (RVP equivalent) system for this test.

6.13 Electrical Fast Transient

6.13.1 Purpose

To demonstrate that bursts or spikes generated by environmental noise (motors, relays, etc.) will not impair the operation of the product. EFT is applied to the input AC and DC power and all I/O cables. This demonstrates compliance to international legal requirements.

6.13.2 Sample Size

Validation: One system

6.13.3 Success Criteria

- No component damage under any condition

- No unsafe operation is under any condition

- Performance Criteria B, as specified by EN 55024: 1998 is used for this test

- A summary of the pass criteria follows:

    o After the test, the equipment shall continue to operate as intended without operator intervention

    o During the test, degradation of performance is allowed. However, no change of operating state or stored data is allowed to persist after the test

6.13.4 Test Profile

Equipment and EUT Setup:

- The setup shall meet the EN 61000-4-4: 1995 specification.

- The test climatic conditions shall comply with the following requirements:

    - Ambient temperature: 15° C - 35° C

    - Relative humidity: 45% - 75%

    - Atmospheric pressure: 86 kPa (860 mbar) – 106 kPa (1060 mbar)

- AC/DC power lines are to be directly connected to the EFT generator coupler at nominal 230 V and 50 Hz or DC power.

- All I/O and communication cables > 3 meters are to be tested in a capacitive clamp with data communication active (USB, SCSI, LAN, etc.).

- Monitor, keyboard and mouse cable testing is not required.

- Use the Product Configuration (RVP equivalent) system for this test.

Test Sequence:

- Start the functional test program.

- Expose the system to a 15 ms burst of high frequency spikes (pulse repetition rate of 5 kHz).

- Repeat every 300 ms for 1 minute. (Table 41 and Table 42).

Table 41: I/O and DC Power Cable Tests

	I/O and DC Power Cables		Comments
a.	Voltage levels	2	250 V and 500 V EFT waveform
b.	Polarities	2	Positive and negative pulses
c.	Pulses per burst	75	
d.	Bursts per minute	200	
e.	Lead combination terminals	1	Capacitive clamp around I/O lines One minute for each I/O line
Total pulses (a x b x c x d x e)		60 k	Per I/O line

AC/DC Power Input Test		Comments
a. Voltage levels	2	AC 500 V and 1.0 kV EFT waveform* DC 250 V and 500 V EFT waveform*
b. Polarities	2	Positive and negative pulses
c. Pulses per burst	75	
d. Bursts per minute	200	
e. Lead combination terminals	6	Applied to six combinations of input AC/DC power (L1, L1-L2, L1-PE, L1-L2-PE, L2, L2-PE) One minute for each combination
Total pulses (a x b x c x d x e)	360 k	Pulses for each power supply

Notes: Connect a second PC on the data lines for data communication during test.
*All voltages are measured open circuit

Table 42: AC/DC Power Input Test

- Ensure the EFT generates the waveform in Figure 191.

Figure 191: EFT Waveform

The figure shows a voltage waveform with V(t) / Vp on the vertical axis (0 to 1) and Time (ns) on the horizontal axis (0 to 100). It is labeled with "Rise Time = 5ns" and "Duration = 50ns".
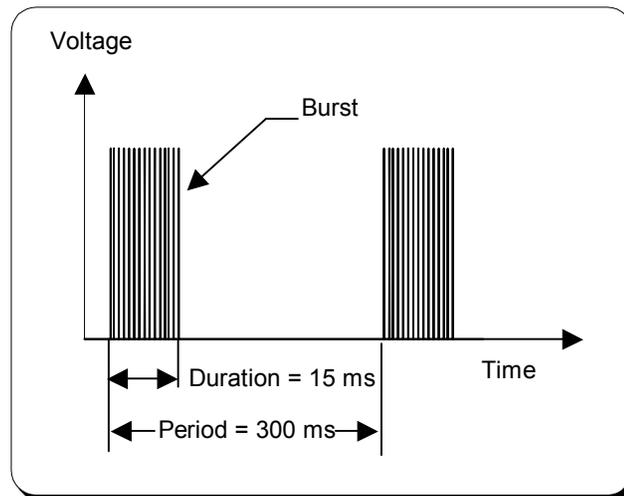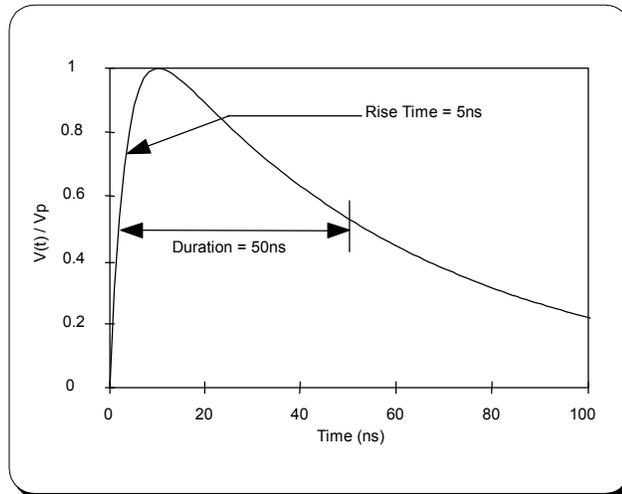
o   Move the peripheral device cables (e.g. mouse, keyboard and joystick) three times when each cable is in the capacitive clamp: before, during and after EFT testing of each cable, respectively.

o   Record the temperature, humidity and atmospheric pressure levels of the environment.

6.14 ESD

6.14.1 Purpose

To demonstrate that system products can withstand static discharges encountered in normal handling or operation of the equipment. This profile is a hybrid of the Intel ESD test requirements and the IEC test profile.

6.14.2 Sample Size

Validation: One system

6.14.3 Success Criteria

Table 43: ESD Voltage Levels

Test Location (Example)	Surface Type	Voltage Levels					
		2 kV (±)	4 kV (±)	6 kV (±)	8 kV (±)	12 kV (±)	15 kV (±)
Chassis Front							
Switches	Nonconductive	Air	Air	Air	Air	Air	Air
LED's	Nonconductive	Air	Air	Air	Air	Air	Air
Drive Surfaces[1]	Both	Air	Air	Air	Air	Air	Air
Power/Reset Button	Nonconductive	Air	Air	Air	Air	Air	Air
Keyholes	Conductive[2]	Contact	Contact	Contact	Contact	Air	Air
Chassis Back							
Voltage Switch	Conductive[2]	Contact	Contact	Contact	Contact	Air	Air
Power Cord Area	Conductive[2]	Contact	Contact	Contact	Contact	Air	Air
Power Cord Area	Nonconductive	Air	Air		Air		Air
I/O Ports	Conductive[2]	Contact	Contact	Contact	Contact		
Add-In Cards/Panels	Conductive[2]	Contact	Contact	Contact	Contact	Air	Air
Other							
Chassis Surfaces	Both	A and C	A and C	A and C	A and C	Air	Air
Vertical Coupling Plane	Conductive[2]	Contact	Contact	Contact	Contact		
Horizontal Coupling Plane	Conductive[2]	Contact	Contact	Contact	Contact		
Failure Classification (Per Polarity)							
Air Visual and Audio Errors [3]		0	1/20	2/20	4/20	10/20	20/20
Air Re-settable Errors [4]		0	0	0	0	2/20	5/20
Air Hard Failure [5]		0	0	0	0	0	0
Contact Visual and Audio Errors [3]		0	1/25	2/25	4/25	10/25	20/25
Contact Re-settable Errors [4]		0	0	0	0	2/25	5/25
Contact Hard Failure [5]		0	0	0	0	0	0

1. Air discharge only. Limited to maximum drive specification, but must be at least 8 kV.
2. For conductive surfaces a minimum of 50 contact discharges, 25 in each polarity are required.
3. Visual and Audio Errors: Errors that do not interfere with the normal operation of the system or cause any loss of data (e.g., read-retries or display flicker or bounce, speaker noise); no operator intervention is required to recover.
4. Re-settable Errors: Errors or faults that require operator intervention to recover from: system hang, unplanned reset, change of operating state or stored data.
5. Hard Failure: Any functional failures that require repair

6.14.4 Test Profile

Equipment and EUT Setup:

- Test the product in a climate-controlled room, within the following ranges:

    - Ambient temperature: 15° C - 35° C

    - Relative humidity: 30% - 60%

    - Atmospheric pressure: 86 kPa (860 mbar) – 106 kPa (1060 mbar)

- The product test setup shall meet the IEC 61000-4-2: 1995 specification.

- The ESD gun(s) shall meet the IEC 61000-4-2: 1995 specification.

- Use the Product Configuration (RVP equivalent) system for this test.

Identify Test Points:

- Only points and surfaces of the product that are normally accessible to personnel during normal usage are to be considered as test discharge point candidates.

- Determine the exact chassis test points by conducting a complete investigative air discharge scan at 15 kV except on the I/O panel, where an 8 kV contact discharge investigatory scan should be performed. Chassis test points should be between 4 to 6 inches apart over the entire front, back, side, and top surfaces.

- In addition to the chassis test points, include all LEDs, disk drive surfaces, power or reset buttons, keyholes, and around external connector sockets with cables plugged into them.

- One of the test points must include the coupling plane for all systems configurations. Refer to the IEC 61000-4-2: 1995 specification to determine the proper coupling plane setup for each configuration type.

- Mark and record all points at which an actual discharge occurred.

Test Sequence:

- Start functional test.

- Record the temperature, humidity and atmospheric pressure levels of the environment.

- All external connectors must have cables attached and should be functionally operated during the ESD testing if possible.

- Use Table 1 to determine whether air or contact discharge should be applied to a test point.

- Typically 50 contact discharges (25 in each polarity) shall be applied to conductive surfaces up to 8 kV.

- Air discharge shall be applied to the non-conductive surfaces from 2 kV to 15 kV. Application of electrostatic discharges to the contacts of connectors is not required.

- The mouse is to be moved and the keyboard Num Lock, Scroll Lock and Caps Lock keys pushed 3 times at each voltage level per test point: Once before ESD discharging to the test point, once after 1/2 the ESD discharges have been discharged to the test point and once after all discharges have been discharged to the test point, respectively.

- Set the ESD generator to 2 kV and apply 20 air discharges, with a minimum of 1 second between each discharge, to each test point. Ensure that the tester is fully charged before each discharge.

- Comparing the results to Table 43 after each series of 20 discharges, record the number of visual and audio errors, re-settable errors and hard failures. Repeat for all marked test points.

- Record the results of all test points at each voltage level before ascending to the next voltage level (e.g., 2 kV, 4 kV, 6 kV, 8 kV, 12 kV, and 15 kV) and repeating the 20 air discharges at each test point.

- Air discharge at slots and apertures and insulating surfaces as per regulatory requirements referencing EN 55024: 1998 include the following guidance: "On those parts of the system where it is not possible to perform contact discharge testing, the equipment should be investigated to identify user accessible points where breakdown may occur; examples are openings at edges of keys or in the covers of keyboards and telephone handsets. Such points are tested using the air discharge method. This investigation should be restricted to those areas normally handled by the user."

- Repeat the complete air discharge test in the negative polarity.

- 50 contact discharges (25 in each polarity) shall be applied to conductive surfaces and to the horizontal and vertical coupling planes at 2 kV, 4 kV, 6 kV, and 8 kV with a contact discharge gun. The application of electrostatic discharges to the contacts of connectors is not required.

- For compliance to regulatory requirements reference EN 55024: 1998: "The system shall be exposed to at least 200 discharges, 100 each at negative and positive polarity, at a minimum of four test points (a minimum of 50 discharges at each point). One of the test points shall be subjected to at least 50 indirect discharges (contact) to the center of the front edged of the horizontal coupling plane. The remaining three test points shall each receive at least 50 direct contact discharges."

- Repeat the complete contact discharge test in the negative polarity.

6.15 AC, DC, and I/O Surge

6.15.1 Purpose

To demonstrate power supply and system tolerance for transients on the AC and DC power lines caused by switching or lightning.

6.15.2 Sample Size

Validation: Three systems

6.15.3 Success Criteria

- No unsafe operation is allowed under any condition.

- Diagnostic failures or system hangs are an overall test failure.

- Intermittent system errors that are not captured by diagnostics or do not hang the system are not failures.

- No component damage under any condition.

- Performance Criteria B, as specified by EN 55024: 1998, is used for this test.

- A summary of the pass criteria follows:

  o After the test, the equipment shall continue to operate as intended without operator intervention

  o During the test, degradation of performance is allowed. However, no change of operating state or stored data is allowed to persist after the test

6.15.4 Test Profile

Equipment and EUT Setup:

- The test and equipment setup shall meet IEC 61000-4-5: 1995.

- The input waveforms are to be verified and shall meet the IEC 61000-4-5: 1995. (Figure 192)

- All tests shall be performed with the product power switch in the "on" position.

- Note: System voltages are automatically set for auto-switching power supplies.

- The product test setup and immunity test equipment shall meet the IEC 61000-4-3: 1995 specification.

- The power cord of the equipment under test shall be plugged into 230 V, 50 Hz power.

- Wiring shall be consistent with the manufacturer's recommended procedures.

- Equipment shall be in its housing with all covers and access panels in place, unless otherwise stated.

- The system is placed on the front edge of the test table; this should be flush with the vertical uniform electromagnetic field.

Use the Product Configuration (RVP equivalent) system for this test.

Test Sequence:

- Start functional test.

- All testing of equipment shall be performed in conditions as close as possible to installed conditions.
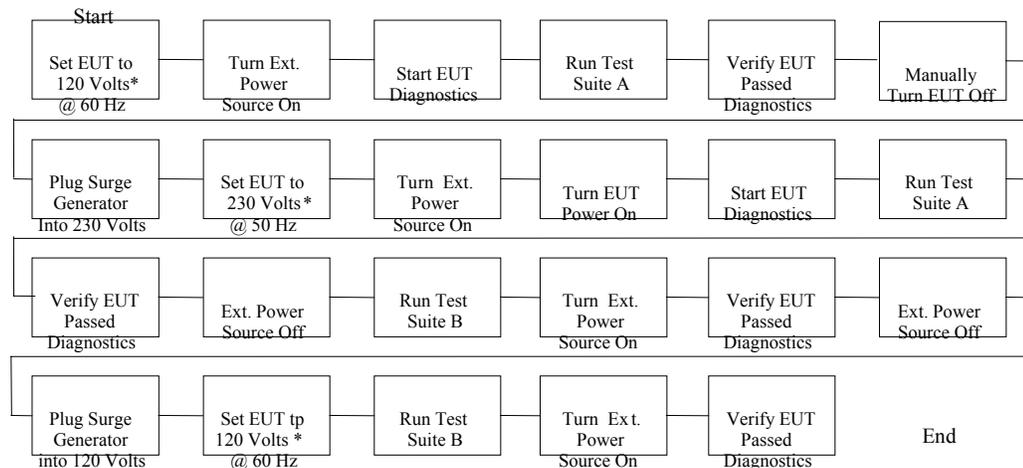
Table 44: Test Suite A (Power Source On)

Qty	Pulse Count During Test	Reference
4	Lead combinations	Table 3
5	Pulses per test voltage level at phase angles 0°, 90°, 180°, 270°, and 360°	Table 4
6	Applied voltage levels (3 for unidirectional, 3 for ringwave)	Table 4
2	Positive and negative polarity	Table 4
240	Total pulse count (4 x 5 x 6 x 2)	

Table 45: Test Suite B (Power Source Off)

Qty	Pulse Count During Test	Reference
4	Lead combinations	Table 3
5	Pulses per test voltage level (random)	Table 4
6	Applied voltage levels (3 for unidirectional, 3 for ringwave)	Table 4
2	Positive and negative polarity	Table 4
240	Total pulse count (4 x 5 x 6 x 2)	

Table 46: Test Suite Lead Combinations

Test Suite Lead Combinations
Lead 1 to Lead 2
Lead 1 to Grd
Lead 2 to Grd
Lead 1 and 2 to Grd

Table 47: AC Inputs

AC Inputs	Pulses Applied To Each Lead			
Applied Voltage	Unidirectional		Ringwave	
Level	+	-	+	-
0.5 kV	5	5	5	5
1.0 kV	5	5	5	5
2.0 kV	5	5	5	5
Optional 3.0 kV	N/A	N/A	5	5

3 kV test level is recommended, 2 kV is required

Table 48: DC Inputs

		Pulses Applied Each Lead to ground			
	Applied Voltage Level	Unidirectional		Ringwave	
		+	-	+	-
DC Inputs	0.5 kV	5	5	5	5
I/O	1.0 kV	5	5	N/A	N/A

- The waveforms in Figure 192 are generated by the surge generator and are applied to the AC/DC and I/O power leads. These input waveforms are to be verified with an oscilloscope before the start of the test process using the criteria listed in EN 61000-4-5: 1995.

- I/O surge testing on all ports connected to outdoor cables is required.

- Applicable only to ports which according to the manufacturer may connect directly to outdoor cables.

Figure 192: Surge Generator Waveforms

Ring Wave
(0.5 uSec - 100 KHz)
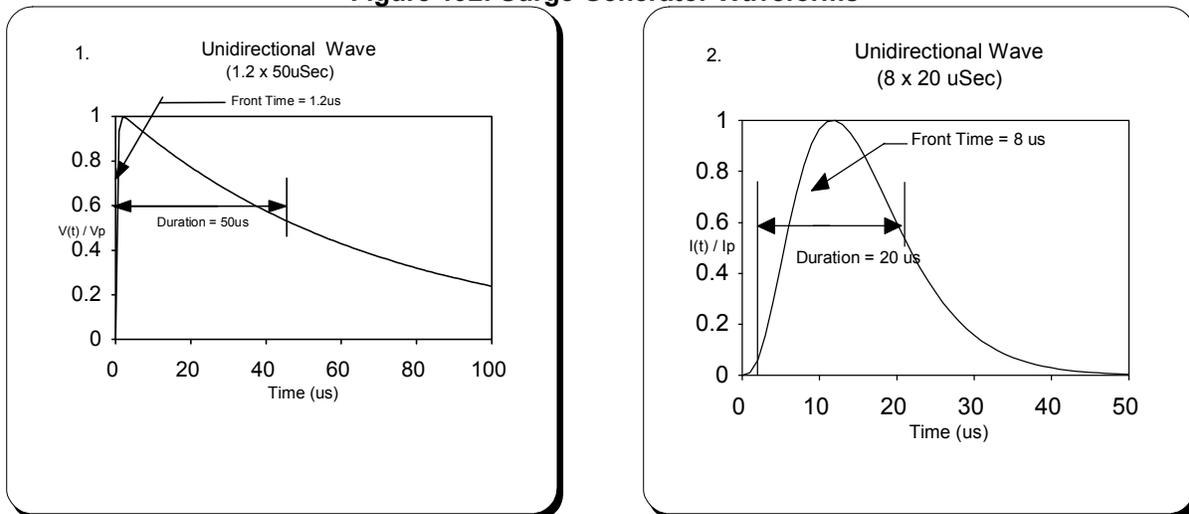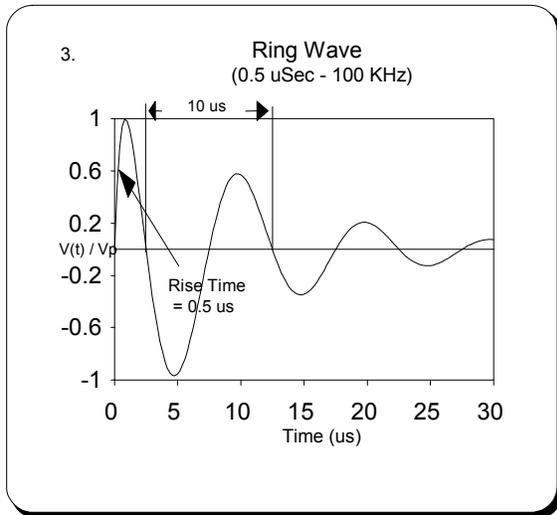
Graph notes
- Open Circuit Voltage
  Unidirectional generator impedance
    L1 to L2 is 2 ohms
    L1 or L2 to PE is 12 ohms
- Short Circuit Current
- Open Circuit Voltage

6.16 AC Voltage, Frequency, and Source Interrupt

6.16.1 Purpose

To demonstrate power supply performance with maximum and typical test loads over a range of AC line voltages and frequencies and to simulate power line interruptions lasting for 20 ms.

6.16.2 Sample Size

Validation: Three systems

6.16.3 Success Criteria

- No component failures.
- No functional failures during or after the stress.
- No system interrupts requiring a reboot.
    - o During the test, degradation of performance is allowed. However, no change of operating state or stored data is allowed to persist after the test.

6.16.4 Test Profile

Test Sequence:
- Use the Product Configuration (RVP equivalent) system for this test.
- Start functional test.
- Voltage/frequency test:
    - o Test once with maximum system load to power supply for each range of voltage/frequency.
    - o Test once with typical system load to power supply for each range of voltage/frequency.

Table 49: Input Setting Vs Frequency

AC POWER SUPPLY INPUT SWITCH SETTING	VOLTAGE FREQUENCY	TIME
100 V - 120 V	90 V at 47 Hz	15 min
	132 V at 47 Hz	15 min
	90 V at 63 Hz	15 min
	132 V at 63 Hz	15 min
200 V - 240 V	180 V at 47 Hz	15 min
	264 V at 47 Hz	15 min
	180 V at 63 Hz	15 min
	264 V at 63 Hz	15 min

- Source Interrupt test

    o Apply maximum motherboard load.

    o Power dropouts for a minimum of one cycle (20 ms at both 60 Hz and 50 Hz), at nominal line voltages (120 V and 230 V).

    o Repeat every 10 seconds starting at the zero-cross-over on the voltage waveform for 30 minutes.

This page is intentionally left blank.

7. 12.9 Liter Reference Design System Validation Results

This section contains the validation results for the 12.9L Intel reference solution.

7.1 Processor Thermal Results

Table 50: Processor Thermal Results

Test Description	Success Criterion	Validation Result	Disposition
End Of Line Thermals	≤ 72 °C	69.7 °C	Pass
Post-Shock & Vibration Thermal	≤ 72 °C	68.4 °C	Pass
Beginning of Line Temperature Cycle Thermals	≤ 72 °C	69.5 °C	Pass
Post-Bake Thermals	≤ 72 °C	69.7 °C	Pass
Post-Power Cycle Thermals	≤ 72 °C	68.2 °C	Pass

7.2 MCH Thermal Results

Table 51: MCH Thermal Results

Test Description	Success Criterion	Validation Result	Disposition
End Of Line Thermals	≤ 110 °C	86.7 °C	Pass
Post-Shock & Vibration Thermal	≤ 110 °C	80.6 °C	Pass
Beginning of Line Temperature Cycle Thermals	≤ 110 °C	86.1 °C	Pass
Post-Bake Thermals	≤ 110 °C	85.0 °C	Pass
Post-Power Cycle Thermals	≤ 110 °C	89.3 °C	Pass

7.3 Subsystem Thermals By Load Condition

7.4 High Load/High Temperature

Table 52: High Load/High Temperature

Component	Success Criterion	Validation Result	Disposition
Processor	≤ 72 °C	68.6 °C	Pass
MCH	≤ 110 °C	75.2 °C	Pass
ICH	≤ 120 °C	94 °C	Pass
Memory (DDR400)	≤ 93 °C	80 °C	Pass
Memory (DDR 2-533)	≤ 95 °C	76 °C	Pass
HDD	≤ 55 °C	42 °C	Pass
ODD	≤ 45 °C	43 °C	Pass
FDD	≤ 55 °C	45 °C	Pass
PCI Express* Graphics Card	≤ 55 °C	39 °C	Pass
PCI Cards (Low Profile)	≤ 55 °C	40 °C	Pass
PCI Cards (Full Hieght / Half Length)	≤ 55 °C	39 °C	Pass
PSU	≤ 50 °C	41 °C	Pass

7.5 High Load/Low Temperature

Table 53: High Load/Low Temperature

Component	Success Criterion	Validation Result	Disposition
Processor	≤ 72 °C	63.7 °C	Pass
MCH	≤ 110 °C	91.8 °C	Pass
ICH	≤ 120 °C	97 °C	Pass
Memory (DDR400)	≤ 93 °C	87 °C	Pass
Memory (DDR 2-533)	≤ 95 °C	83 °C	Pass
HDD	≤ 55 °C	36 °C	Pass
ODD	≤ 45 °C	39 °C	Pass
FDD	≤ 55 °C	43 °C	Pass
PCI Express* Graphics Card	≤ 55 °C	39 °C	Pass
PCI Cards (Low Profile)	≤ 55 °C	34 °C	Pass
PCI Cards (Full Hieght / Half Length)	≤ 55 °C	36 °C	Pass
PSU	≤ 50 °C	38 °C	Pass

7.6 Idle Load/Low Temperature

Table 54: Idle Load/Low Temperature

Component	Success Criterion	Validation Result	Disposition
Processor	≤ 72 °C	57.6 °C	Pass
MCH	≤ 110 °C	80.8 °C	Pass
ICH	≤ 120 °C	60 °C	Pass
Memory (DDR400)	≤ 93 °C	73 °C	Pass
Memory (DDR 2-533)	≤ 95 °C	70 °C	Pass
HDD	≤ 55 °C	36 °C	Pass
ODD	≤ 45 °C	38 °C	Pass
FDD	≤ 55 °C	40 °C	Pass
PCI Express* Graphics Card	≤ 55 °C	48 °C	Pass
PCI Cards (Low Profile)	≤ 55 °C	47 °C	Pass
PCI Cards (Full Hieght / Half Length)	≤ 55 °C	43 °C	Pass
PSU	≤ 50 °C	39 °C	Pass

7.7 Acoustic Results

Table 55: Acoustic Results

Test Description	Performance Metric	Success Criterion	Validation Result	Disposition
System-Level Acoustics	Sound Power, LwAm	≤ 4.0 BA	3.73 BA	Pass

7.8 Mechanical Shock and Vibration/Assembly/Fit

Table 56: Mechanical Shock and Vibration/Assembly/Fit

Test Description	Success Criterion	Validation Result	Disposition
Mechanical Shock	0/10 boot failures	0/10 failures	Pass
	0/10 visual inspection failures	3/10 failures	Fail[1]
	0/10 Socket SJ cracks >50% in length	0/10 failures	Pass
	0/10 MCH SJ cracks >50% in length	0/10 failures	Pass
	0/10 ICH SJ cracks >50% in length	0/10 failures	Pass
Vibration	0/10 boot failures	0/10 failures	Pass
	0/10 visual inspection failures (e.g. cracking, disengagement, etc.)	1/10 failures	Pass
	0/10 Socket SJ cracks >50% in length	0/10 failures	Pass
	0/10 MCH SJ cracks >50% in length	0/10 failures	Pass
	0/10 ICH SJ cracks >50% in length	0/10 failures	Pass
MCH Clip Actuation Force	≤ 63.6 N (m+2.5s)	26.8 N	Pass
Deflection Control Bkt Actuation Force	≤ 63.6 N (m+2.5s)	Not tested	Pass
Volumetric Keep-In Compliance	≤ 10.16 mm (m+2.5s) - Clip height above board	10.34 mm	Pass

Note: 1 Deflection Control Bracket failed in lateral shock. Design has been modified and modled to fix this issue. No plans to retest.

7.9 Processor VR and Motherboard Thermals

Table 57: Processor VR and Motherboard Thermals

Component	Success Criterion	Validation Result	Disposition
Motherboard outside socket perimeter	≤ 120 °C	87 °C	Pass
Motherboard under socket	≤ 91 °C	61 °C	Pass
FET's	≤ 150 °C	91 °C	Pass
Input Bulk Capacitors	≤ 105 °C	67 °C	Pass
I/O Inductors	≤ 85 °C	50 °C	Pass
Gate Drivers	≤ 125 °C	95 °C	Pass
Output Capacitors	≤ 105 °C	63 °C	Pass

7.10 Operating Temperature Range

Table 58: Operating Temperature Range

Test Description	Success Criterion	Validation Result	Disposition
Operating Temperature	0/3 Boot Failures at 5 °C	0/3 failures	Pass
	0/3 Boot failures at 40 °C	3/3 failures	Fail[1]

Note: 1 Failures were due to a known compatibility issue with the units under test. There are no plans to re-run this test.   This needs a better description of the compatibility issue and a justification for why the test was not re-done.

7.11 PSU Electricals

Table 59: PSU Electricals

Test Description	Success Criterion	Validation Result	Disposition
PSU AC Surge Test	0/3 Failures at 2kV across AC Low to Gnd	0/3 Failures	Pass
	0/3 Failures at 2kV across AC High to Gnd	0/3 failures	Pass
	0/3 Failures at 2kV across AC Low to High	0/3 failures	Pass
PSU Line Voltage / Frequency, Source Interrupt Test	0/3 Failures @ 90, 132V @ 47Hz & 63Hz for 15 min	0/3 failures	Pass
	0/3 Failures @ 180, 264V @ 47Hz & 63Hz for 15 min	0/3 failures	Pass
	0/3 Failures, Power Dropout (20mSec @ 50 & 60 Hz for 120 & 230 V	0/3 failures	Pass

7.12 EMI and ESD

Table 60: EMI and ESD

Test Description	Success Criterion	Validation Result	Disposition
System Radiated Emission	FCC B baseline scan w/ 100MB/s LAN	0.45dB Margin	Pass
	FCC B baseline scan w/ 10MB/s LAN	1.98dB Margin	Pass
ESD	Multiple Air discharges 2 - 15KV	System functions after 15KV	Pass
	Multiple Contact discharges 2 - 8KV	System functions after 12KV	Pass

This page is intentionally left blank.

8. 6.9L Validation Results

8.1 Processor Thermal Results

Table 61: Processor Thermal Results

Test Description	Success Criterion	Validation Result	Disposition
End Of Line Thermals	≤ 72 °C	71.2 °C	Pass
Post-Shock & Vibration Thermal	≤ 72 °C	70.0 °C	Pass
Beginning of Line Temperature Cycle Thermals	≤ 72 °C	71.0 °C	Pass
Post-Bake Thermals	≤ 72 °C	71.2 °C	Pass
Post-Power Cycle Thermals	≤ 72 °C	69.7 °C	Pass

8.2 MCH Thermal Results

Table 62: MCH Thermal Results

Test Description	Success Criterion	Validation Result	Disposition
End Of Line Thermals	≤ 110 °C	93.4 °C	Pass
Post-Shock & Vibration Thermal	≤ 110 °C	92.1 °C	Pass
Post-Bake Thermals	≤ 110 °C	91.7 °C	Pass
Post-Power Cycle Thermals	≤ 110 °C	95.9 °C	Pass

8.3 Subsystem Thermals By Load Condition

8.4 High Load/High Temperature

Table 63: High Load/High Temperature

Component	Success Criterion	Validation Result	Disposition
Processor	≤ 72 °C	70.3 °C	Pass
MCH	≤ 110 °C	79.4 °C	Pass
ICH	≤ 120 °C	117.1 °C	Pass
Memory (DDR 2-533)	≤ 93 °C	91.6 °C	Pass
HDD	≤ 55 °C	45 °C	Pass
ODD	≤ 50 °C	48.1 °C	Pass
Express Card* Module	≤ 65 °C	57.9 °C	Pass
PSU	≤ 55 °C	51.1 °C	Pass

8.5 High Load/Low Temperature

Table 64: High Load/Low Temperature

Component	Success Criterion	Validation Result	Disposition
Processor	≤ 72 °C	65.5 °C	Pass
MCH	≤ 110 °C	82.7 °C	Pass
ICH	≤ 120 °C	116.2 °C	Pass
Memory (DDR 2-533)	≤ 93 °C	86.9 °C	Pass
HDD	≤ 55 °C	40.4 °C	Pass
ODD	≤ 50 °C	44.4 °C	Pass
Express Card* Module	≤ 65 °C	53.0 °C	Pass
PSU	≤ 55 °C	49.0 °C	Pass

8.6 Idle Load/Low Temperature

Table 65: Idle Load/Low Temperature

Component	Success Criterion	Validation Result	Disposition
Processor	≤ 72 °C	55.9 °C	Pass
MCH	≤ 110 °C	74.0 °C	Pass
ICH	≤ 120 °C	116.5 °C	Pass
Memory (DDR 2-533)	≤ 93 °C	68.3 °C	Pass
HDD	≤ 55 °C	42.9 °C	Pass
ODD	≤ 50 °C	45.5 °C	Pass
Express Card* Module	≤ 65 °C	48.8 °C	Pass
PSU	≤ 55 °C	51.3 °C	Pass

8.7 Acoustic Results

Table 66: Acoustic Results

Test Description	Performance Metric	Success Criterion	Validation Result	Disposition
System-Level Acoustics	Sound Power, LwAm	≤ 4.0 BA	4.05 BA	Fail[1]

Note: Acoustic performance was impacted due to design changes to the thermal management of the system. The order of magnitude of this failure from the success criteria is perceived to be negligible to the end user and will not be modified at this time.

8.8 Mechanical Shock and Vibration/Assembly/Fit

Table 67: Mechanical Shock and Vibration/Assembly/Fit

Test Description	Success Criterion	Validation Result	Disposition
Mechanical Shock	0/10 boot failures	0/10 failures	Pass
	0/10 visual inspection failures	0/10 failures	Pass
	0/10 Socket SJ cracks >50% in length	0/10 failures	Pass
	0/10 MCH SJ cracks >50% in length	0/10 failures	Pass
	0/10 ICH SJ cracks >50% in length	0/10 failures	Pass
Vibration	0/10 boot failures	0/10 failures	Pass
	0/10 visual inspection failures	0/10 failures	Pass
	0/10 Socket SJ cracks >50% in length	0/10 failures	Pass
	0/10 MCH SJ cracks >50% in length	0/10 failures	Pass
	0/10 ICH SJ cracks >50% in length	0/10 failures	Pass
MCH Clip Actuation Force	$\leq$ 63.6 N (m+2.5s)	26.8 N	Pass
Volumetric Keep-In Compliance	Go, No-go gage compliant	Compliant	Pass

8.9 Processor VR and Motherboard Thermals

The processor VR and motherboard thermal analysis shown in section 7.9 was not repeated on the 6.9L Intel reference system.   Due to similarities in the core layout, results were anticipated to be the same.

8.10 Operating Temperature Range

Table 68: Operating Temperature Range

Test Description	Success Criterion	Validation Result	Disposition
Operating Temperature	0/3 Boot Failures at 5 °C	0/3 failures	Pass
	0/3 Boot failures at 40 °C	0/3 failures	Pass

8.11 PSU Electricals

Table 69: PSU Electricals

Test Description	Success Criterion	Validation Result	Disposition
PSU AC Surge Test	0/3 Failures at 2kV across AC Low to Gnd	2/2 failures	Fail[1]
	0/3 Failures at 2kV across AC High to Gnd	2/2 failures	Fail[1]
	0/3 Failures at 2kV across AC Low to High	2/2 failures	Fail[1]

Note: 1 Two (out of two) systems failed AC surge test at 2.0 kV; the third system and additional testing was not completed due to these failures. Initial investigation indicates that the issue can be resolved within the PSU. A similarly configured system has been shown to function properly in this stress condition with alternate power supplies. These failures are believed to be at the ingredient level and do not impact compliancy to the BTX Interface specification. Investigation into PSU performance is ongoing, but performance of remainder of system is considered low risk for this stress condition. Future updates will be included in later revisions of this document if applicable.

8.12 EMI and ESD

Table 70: EMI and ESD

Test Description	Success Criterion	Validation Result	Disposition
System Radiated Emission	FCC B baseline scan w/ 100MB/s LAN	Over limit at 80 MHz and 90 MHz	Fail[1]
	FCC B baseline scan w/ 10MB/s LAN	Over limit at 80 MHz and 90 MHz	Fail[1]

Note: 1 The system does not meet radiated or conducted emission design validation requirements. Initial investigation indicates that the issue can be resolved within the PSU. The system has been verified to meet radiated emissions requirements when an AC line filter is added to the PSU power cord. This is sufficient to confirm that the chassis enclosure is compliant with the radiated emissions design requirements, and that the emissions are limited to the PSU power cord itself. These failures are believed to be at the ingredient level and do not impact compliancy to the BTX Interface specification. Investigation into PSU EMI performance is ongoing, but performance of remainder of system is considered low risk for this test condition. Future updates will be included in later revisions of this document if applicable.

This page is intentionally left blank.

A Investigative Testing

The tests outlined in this section are ones that are not required. These tests may, however, provide useful information. If conducted prior to the Validation Procedures outlined above, they may provide identify an issue that should be addressed before validation. These investigative tests may also be helpful in identifying root cause in the event of a validation failure.

A.1 Airflow Measurements

A.1.1 Recommendation

Investigative testing of total system airflow and subsystem airflow characteristics may be useful in identifying system design characteristics that may be preventing compliance with thermal or acoustic requirements. Monitoring the following characteristics is recommended:

Characteristic	Method
System Airflow	Wind tunnel measurement of system operating point at minimum and maximum TMA and PSU fan speed setpoints
TMA Effective Fan Curve	Wind tunnel measurement of TMA fan curve at minimum and maximum fan speed setpoints
PSU Effective Fan Curve	Wind tunnel measurement of PSU fan curve at minimum and maximum fan speed setpoints
AIC airflow	Hot-wire anemometer measurement of linear airflow speed at the entrance to the AIC
Memory airflow	Hot-wire anemometer measurement of linear airflow speed at the entrance to the memory modules. Since memory airflow is typically comprised of two inlet streams (from the TMA VR exhaust and flow deflected from the rear panel and PSU), it may be necessary to measure memory airflow at several locations in each of the memory module channels.
MCH heatsink airflow	Hot-wire anemometer measurement of linear airflow speed at the entrance to the MCH heatsink

A.2 Air Temperature Measurements

A.2.1 Recommendation

Investigative testing of subsystem airflow temperature can be useful in identifying system design characteristics that may be preventing compliance with thermal or acoustic requirements. Monitoring the air temperature with a thermocouple at the following locations during thermal testing is recommended:

- Thermal Module Inlet (TMI)

- Memory airflow inlet

- ICH airflow approach

- MCH heatsink inlet

- AIC airflow inlet

B Recommended Visual Inspection

Table 71: Visual Inspection Procedure and Equipment

Overhead task lighting source	100 - 200 candles fluorescent (SECC)
	60-120 candles fluorescent (boards)

Personnel performing visual inspection must be certified.

Personnel must have vision (corrected to) 20/20-20/25. Eye exam results must not be more than one (1) year and records must be retained at the site where inspection is conducted.

All inspection is to be done at 1X (no magnification) at 18" to 24" distance viewed directly along each inspection axis with very minimal rotation.

Recommended viewing period is 7 to 10 seconds per component or surface being inspected.

A 3 diaper (1.75x magnification) may be used to verify defects but is not intended for use to identify defects.

Table 72: Visual Inspection Defect Criteria

Defect Type	Applicable to:	Description	Reject when:
Crack	Components	A narrow break or split in the material	Any surface crack is visible (including complete breakage)
Dent	Motherboard	Any depression on a surface caused by impact.	Any dent is visible.
Stress Discoloration	(Plastic) Components	Evidence of plastic deformity (white stress dislocation lines).	Any evidence of stress discoloration.
Scratch	Components Motherboard	A shallow groove or line that is formed as a result of an abrasive action that removes material.	Any depth with dimensions greater than .020" X .090" (0.5 X 2.3 mm)
Debris	Components Motherboard	Particulate matter that is generated through abrasive action from the components in the system being tested.	Any particle greater in size (any dimension) than .020" (0.5mm)
Wire Damage	System wires	Wire housing damaged in such a way as to expose underlying wire metal.	Any visible bare metal.
Unseated Connector	Connector housings Sockets /Connectors	For connectors with locking tabs: tabs disengaged.	Any tab not fully engaged in the mating hole.
		For connectors with no locating tabs: a discernable gap between the connector and cable housing or between the connector and the mating surface (e.g. motherboard)	Gap or tilt greater than .040" (1mm).
Dislodged Clip	Heat Sink Clip	Missing clip, loose clip, or discernable gap between the heat sink clip and the item to which is intended to be mated.	Clip is missing. Clip is visibly loose or becomes disengaged in the removal of the DUT or during inspection. Any gap between mating surfaces greater than .040" (1mm).
Dislodged Heat Sink	Heat Sink	Missing or displaced (vertical movement relative to mating surface of processor) heat sink	Displacement of greater than .040" (1mm)
Fastener Back out	Fastener Elements	Displacement of a fastener from the mating feature in a direction opposite insertion.	Displacement of greater than .040" (1mm)

Figure 193: Connector Seating Inspection Examples

Acceptable

Unacceptable

Unseated
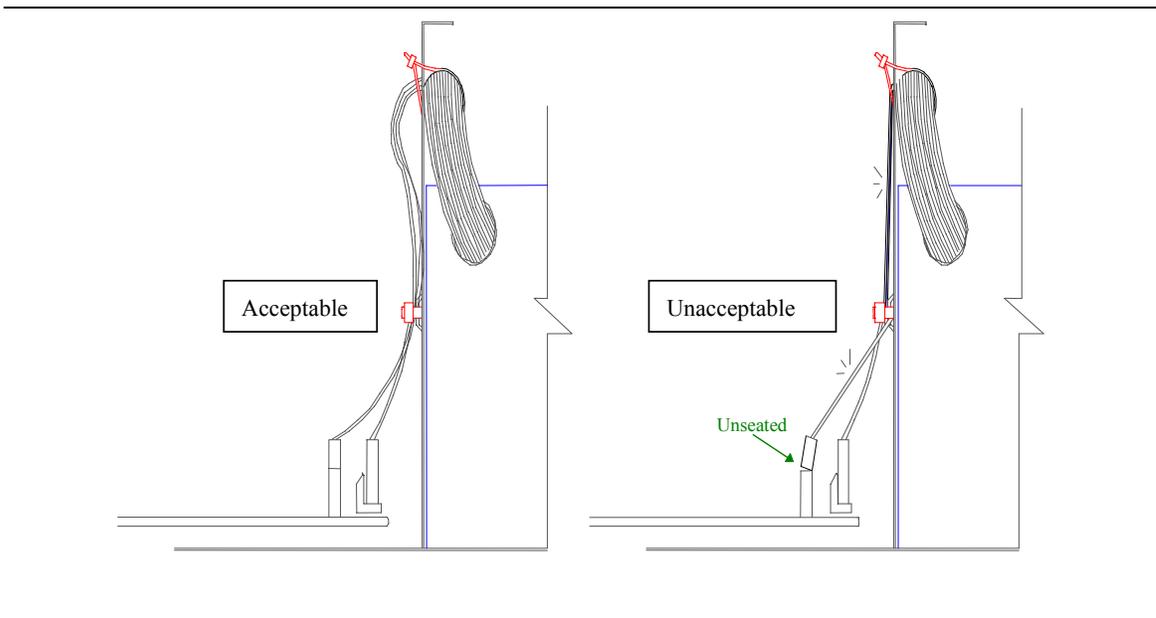
Figure 194: Connector Seating Inspection Examples

<u>Acceptable</u>   **<u>Acceptable</u>**   **<u>Not Acceptable</u>**

Not seated but locked

Not seated and not locked

<u>Acceptable</u>   **<u>Not Acceptable</u>**

Not seated
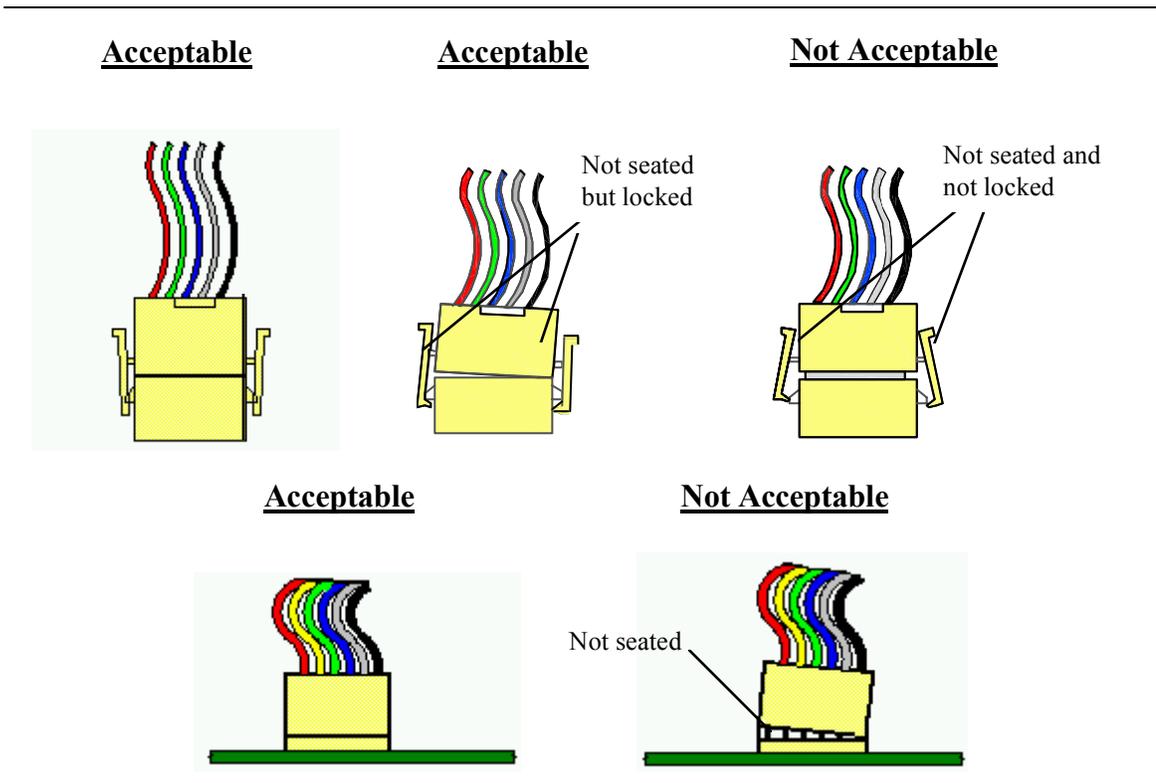
Figure 195: Cable Harness Inspection Examples

Not Acceptable

Wire Exposed

Cable

For connections terminated with a solder ball, exposed wire up to 1/8" from the solder ball is acceptable.
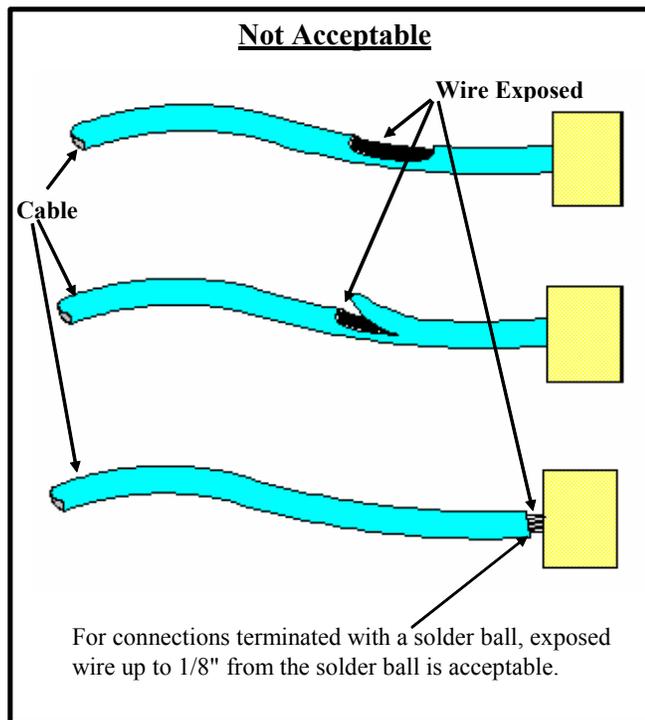
C Bench Top Power Supply Control of System Fans

For fans that use a 4-pin header with voltage and PWM regulated FSC:

- Connect a function generator "output" a second (synthesized) function generator's Input "Trigger".

  - Use 50 Ohm coaxial cabling to connect the terminals.

- Attach a 50 Ohm pass-thru load to the "Function" terminal of the second function generator.

- Attach the 50 Ohm load center signal conductor to the fourth terminal (PWM) of the fan header.

- Connect the outer shield conductor to the fan header ground terminal.

- Configure the first function generator as follows:

  - Square wave

  - Amplitude – 3V peak to peak

  - Frequency – 25 kHz

- Configure the second (synthesized) function generator as a pulse generator as follows:

  - Modulation – Burst

  - Trig Source – External

  - Burst Count – 1

  - Sweep – On

  - Offset – 2.5 Vpp

- The PWM control of the fan is determined by the second function generator's frequency and phase settings.
  - Phase setting of 0° sets a PWM range of 50-100%.

  - Phase setting of 180° selects a range of 0-50%.

  - The frequency setting determines the PWM pulse width.

    - Higher frequencies increase pulse PWM widths per Equation 91:

Equation 91: PWM width Vs Input Frequency

Freq (MHz) = 0.0125 / desired PWM (0-0.5)

- For PWM of 0 – 0.5 (0-50%) use a phase setting of 180%.

- For PWM settings of 0.5 to 1.0 (50-100%) use the inverse of the desired and a phase setting of 0%.

  - **Note:** Due to equipment limitations the actual frequency range may be 25.2 kHz to 1.0 MHz which translates to a PWM range of 1.25% to 49.6% and 50.4% to 98.75%.

- Monitor the output waveform with an oscilloscope to confirm signal frequency, duty cycle, and voltage offset.

- Monitor fan speed with tachometer.

For PSU fans:

- Remove PSU from chassis.

- Cut each of the four wires from the end of the 2x12 PSU power connector (Red, Black, Orange, and Yellow).

- Disconnect fan power lead from PSU circuit board.

- Route the PSU fan power lead through the fan grill opening.

- Connect PSU fan female connector to bench top power supply.

- Install modified PSU into chassis.

- Monitor fan speed with tachometer.

C.1 Combined Top-Side Stiffening and Preload

The strategies of topside stiffening and preload can be used in tandem. When preload is used to pre-curve the board downward, and thus delay the on-set of upward board curvature, topside stiffening can be used to reduce the maximum upward curvature. Figure 196 shows this effect as it plots board curvature against clip stiffness for a range of preload values.

These two distinct strategies work well together to manage board curvature in the shock condition. One strategy can be used to supplement the other. For example, the maximum allowable static compressive load for the processor or socket limits preload, the clip stiffness can be increased to provide the additional protection needed. Or when volumetric design constraints limit the achievable stiffness for the clip, preload may be added to supplement the socket protection.

Figure 196: Board Curvature Versus Clip Stiffness
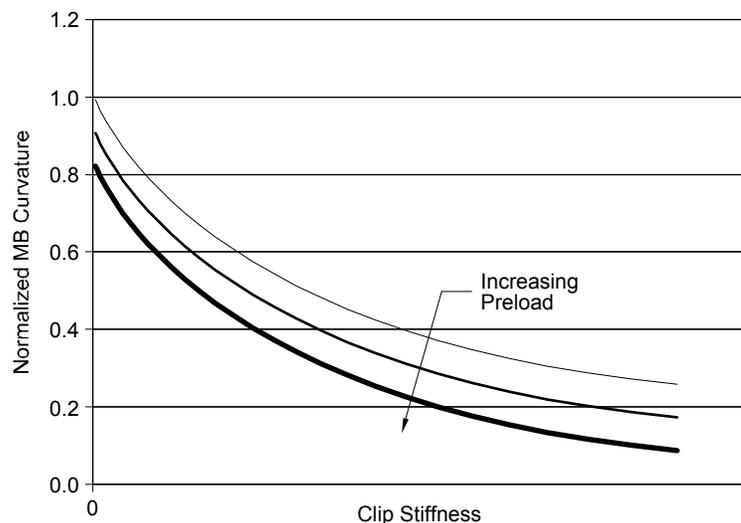
D BTX Rear Panel I/O Shield

The drawing shown in Figure 197 demonstrates the I/O shield geometry used for the Intel reference designs. As stated in Section 5.1 the location and orientation of the I/O connectror features are not driven by the specification and have been removed from Figure 197 for clarity.

Figure 197. BTX Complinat Rear Panel I/O Shield Example

I/O Connector Openings Removed
for clarity