



64-bit Intel[®] Xeon[™] Processor MP with 1 MB L2 Cache

Thermal/Mechanical Design Guidelines

March 2005

Document Number: 306750-001





INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life saving, life sustaining applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The Intel Pentium and Xeon processors may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available upon request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents, which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting [Intel's Web Site](#).

Intel, Pentium and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Copyright © 2005, Intel Corporation. All rights reserved.

* Other brands and names may be claimed as the property of others.



Contents

1	Introduction	7
1.1	Objective	7
1.2	Scope	7
1.3	References	7
1.4	Definition of Terms	8
2	Thermal/Mechanical Reference Design	9
2.1	Mechanical Requirements	9
2.1.1	Performance Target	9
2.1.2	Critical Interface Dimensions (CID)	10
2.2	Thermal Requirements	11
2.2.1	Thermal Profile	11
2.2.2	T _{CONTROL} Definition	13
2.2.3	Performance Targets	14
2.3	Characterizing Cooling Solution Performance Requirements	15
2.3.1	Fan Speed Control	15
2.3.2	Processor Thermal Characterization Parameter Relationships	16
2.3.3	Chassis Thermal Design Considerations	18
2.4	Thermal/Mechanical Reference Design Considerations	19
2.4.1	Heatsink Solutions	19
2.4.2	Thermal Interface Material	19
2.4.3	Summary	20
2.4.4	Assembly Overview of the Intel Reference Thermal Mechanical Design	20
2.4.5	Thermal Solution Performance Characteristics	21
2.4.6	Structural Considerations of Cooling Solution	22
2.4.7	Components Overview	23
A	Mechanical Drawings	27
B	Testing Methods	37
B.1	Case Measurement	37
B.2	Supporting Test Equipment	37
B.3	Thermal Calibration and Controls	38
B.4	IHS Groove	38
B.5	Thermocouple Conditioning and Preparation	39
B.6	Thermocouple Attachment to the IHS	40
B.7	Curing Process	43
B.8	Thermocouple Wire Management	43
B.9	Local Air Thermocouple Placement	45
C	Safety Requirements	47
D	Quality and Reliability Requirements	49
D.1	Intel Verification Criteria for the Reference Designs	49
D.1.1	Reference Heatsink Thermal Verification	49
D.1.2	Environmental Reliability Testing	49
D.1.3	Material and Recycling Requirements	51
E	Enabled Suppliers Information	53

F	Processor Thermal Management Logic and Thermal Monitor Features.....	55
F.1	Thermal Management Logic and Thermal Monitor Feature	55
F.1.1	Processor Power Dissipation.....	55
F.1.2	Thermal Monitor Implementation	55
F.1.3	Operation and Configuration.....	57
F.1.4	Thermal Monitor 2.....	57
F.1.5	System Considerations	58
F.1.6	Operating System and Application Software Considerations.....	58
F.1.7	Legacy Thermal Management Capabilities	59
F.1.8	Cooling System Failure Warning	61

Figures

2-1	Critical Interface Dimensions (Sheet 1 of 2)	10
2-2	Critical Interface Dimensions (Sheet 2 of 2)	11
2-3	Thermal Profile Diagram	12
2-4	T _{CONTROL} and Thermal Profile Interaction	13
2-5	Thermal Profile for the 64-bit Intel® Xeon™ Processor MP with 1 MB L2 Cache	14
2-6	T _{CONTROL} and Fan Speed Control	15
2-7	Processor Thermal Characterization Parameter Relationships.....	17
2-8	Exploded View of Cooling Solution Thermal Solution Components.....	21
2-9	2U+ Cooling Solution Heatsink Thermal Performance	22
2-10	Isometric View of the 2U+ Cooling Solution Heatsink	23
2-11	Hat Spring Isometric View.....	25
2-12	Isometric View of Hat Spring Attachment to the Base Board	25
A-1	2U Cooling Solution Heatsink (Sheet 1 of 4)	28
A-2	2U Cooling Solution Heatsink (Sheet 2 of 4)	28
A-3	2U Cooling Solution Heatsink (Sheet 3 of 4)	29
A-4	2U Cooling Solution Heatsink (Sheet 4 of 4)	29
A-5	Cooling Solution Hat Spring (Sheet 1 of 3).....	30
A-6	Cooling Solution Hat Spring (Sheet 2 of 3).....	30
A-7	Cooling Solution Hat Spring (Sheet 3 of 3).....	31
A-8	Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 1 of 5).....	32
A-9	Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 2 of 5).....	33
A-10	Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 3 of 5).....	34
A-11	Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 4 of 5).....	34
A-12	Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 5 of 5).....	35
B-1	IHS Groove	38
B-2	Groove to Pin Indicator	39
B-3	IHS Groove	39
B-4	Bending Tip of Thermocouple.....	40
B-5	Securing Thermocouple Wires with Kapton* Tape	41
B-6	Thermocouple Bead Placement.....	41
B-7	Thermocouple Placement.....	41
B-8	3D Micromanipulator to Secure Bead Location	42



B-9	Measuring Resistance between Thermocouple and IHS	42
B-10	Applying the Adhesive on the Thermocouple Bead	42
B-11	Thermocouple Wire Management in the Groove	43
B-12	Removing Excess Adhesive from the IHS	44
B-13	Filling the Groove with Adhesive	44
B-14	Thermocouple Wire Management	45
B-15	Local Air Thermocouple Placement for Passive Heatsinks	45
B-16	Local Air Thermocouple Placement for Active Heatsinks (Side View)	46
B-17	Local Air Thermocouple Placement for Active Heatsinks (Plan View)	46
D-1	Random Vibration PSD	49
D-2	Shock Acceleration Curve	50
F-1	Thermal Sensor Circuit	56
F-2	Concept for Clocks under Thermal Monitor Control	56
F-3	Thermal Monitor 2 Frequency and Voltage Ordering	58
F-4	On-Die Thermal Diode Sensor Time Delay	60

Tables

2-1	Performance Target Table (Sheet 1 of 2)	9
2-2	Performance Target Table (Sheet 2 of 2)	14
2-3	Fan Speed Control, T_{CONTROL} and T_{DIODE} Relationship	16
2-4	Recommended Thermal Grease Dispense Weight	24
A-1	Mechanical Drawing List	27
B-1	Test Equipment	37

Revision History

Revision Number	Description	Date
001	<ul style="list-style-type: none"><li data-bbox="500 443 829 468">• Initial release of the document.	March 2005

Note: Not all revisions may be published.



1 Introduction

1.1 Objective

This guide describes the reference thermal/mechanical solution and design parameters required for the 64-bit Intel® Xeon™ processor MP with 1 MB L2 cache. It is also the intent of this document to comprehend and demonstrate the processor cooling solution features and requirements.

Furthermore, this document provides an understanding of the processor thermal characteristics, and discusses guidelines for meeting the thermal/mechanical requirements imposed on the entire life of the processor. The thermal/mechanical solutions described in this document are intended to aid component and system designers to develop and evaluate a processor compatible solution.

1.2 Scope

The thermal/mechanical solutions described in this document pertain only to a solution(s) intended for use with the 64-bit Intel Xeon processor MP with 1 MB L2 cache in 2U+ form factor systems. This document contains the mechanical and thermal requirements of the processor (or processor family) cooling solution. In case of conflict, the data in the *64-bit Intel® Xeon™ Processor MP with 1 MB L2 Cache Datasheet* supersedes any data in this document. Additional information is provided as a reference in the appendix section(s).

1.3 References

Material and concepts available in the following documents may be beneficial when reading this document.

Document	Comment
<i>64-bit Intel® Xeon™ Processor MP with 1MB L2 Cache Datasheet</i>	http://developer.intel.com
<i>mPGA604 Socket Design Guidelines</i>	http://developer.intel.com
<i>64-bit Intel® Xeon™ Processor MP with 1MB L2 Cache Processor Cooling Solution Mechanical Models</i>	http://developer.intel.com
<i>64-bit Intel® Xeon™ Processor MP with 1MB L2 Cache Mechanical Models</i>	http://developer.intel.com
<i>64-bit Intel® Xeon™ Processor MP with 1MB L2 Cache Thermal Test Vehicle and Cooling Solution Thermal Models</i>	http://developer.intel.com
European Blue Angel Recycling Standards	http://www.blauer-engel.de
Thin Electronics Specification (Server System Infrastructure (SSI) Specification) for Rack Optimized Servers	www.ssiforum.com

NOTE: Contact your Intel field sales representative for the latest revision and order number of this document.

1.4 Definition of Terms

Term	Description
Bypass	Bypass is the area between a passive heatsink and any object that can act to form a duct. For this example, it can be expressed as a dimension away from the outside dimension of the fins to the nearest surface.
FMB	Flexible Motherboard Guideline: an estimate of the maximum value of a processor specification over certain time periods. System designers should meet the FMB values to ensure their systems are compatible with future processor releases.
FSC	Fan Speed Control
IHS	Integrated Heat Spreader: a component of the processor package used to enhance the thermal performance of the package. Component thermal solutions interface with the processor at the IHS surface.
mPGA604	The surface mount Zero Insertion Force (ZIF) socket designed to accept the 64-bit Intel® Xeon™ processor MP with 1 MB L2 cache.
P_{MAX}	The maximum power dissipated by a semiconductor component.
Ψ_{CA}	Case-to-ambient thermal characterization parameter (psi). A measure of thermal solution performance using total package power. Defined as $(T_{CASE} - T_{LA}) / \text{Total Package Power}$. Heat source should always be specified for Ψ measurements.
Ψ_{CS}	Case-to-sink thermal characterization parameter. A measure of thermal interface material performance using total package power. Defined as $(T_{CASE} - T_S) / \text{Total Package Power}$.
Ψ_{SA}	Sink-to-ambient thermal characterization parameter. A measure of heatsink thermal performance using total package power. Defined as $(T_S - T_{LA}) / \text{Total Package Power}$.
T_{CASE}	The case temperature of the processor, measured at the geometric center of the topside of the IHS.
T_{CASE_MAX}	The maximum case temperature as specified in a component specification.
TCC	Thermal Control Circuit: Thermal monitor uses the TCC to reduce the die temperature by using clock modulation when the die temperature is very near its operating limits.
$T_{CONTROL}$	A processor unique value, which defines the lower end of the thermal profile and is targeted to be used in fan speed control mechanisms.
Offset	A value programmed into each processor during manufacturing that can be obtained by reading IA_32_TEMPERATURE_TARGET MSR. This is a static and a unique value.
TDP	Thermal Design Power: Thermal solution should be designed to dissipate this target power level. TDP is not the maximum power that the processor/chipset can dissipate.
Thermal Monitor	A feature on the processor that can keep the processor's die temperature within factory specifications under nearly all conditions.
Thermal Profile	Line that defines case temperature specification of a processor at a given power level.
TIM	Thermal Interface Material: The thermally conductive compound between the heatsink and the processor case. This material fills the air gaps and voids, and enhances the transfer of the heat from the processor case to the heatsink.
T_{LA}	The measured ambient temperature locally surrounding the processor. The ambient temperature should be measured just upstream of a passive heatsink or at the fan inlet for an active heatsink.
T_{SA}	The system ambient air temperature external to a system chassis. This temperature is usually measured at the chassis air inlets.
U	A unit of measure used to define server rack spacing height. 1U is equal to 1.75 in, 2U equals 3.50 in, etc.



2 Thermal/Mechanical Reference Design

2.1 Mechanical Requirements

The mechanical performance of the processor cooling solution satisfies the requirements and volumetric keepouts as described in this section.

2.1.1 Performance Target

Table 2-1. Performance Target Table

Parameter	Minimum	Maximum	Unit	Notes
Volumetric Requirements and Keepouts				Refer to drawings in Appendix A
Heatsink Mass		1000 2.2	g lbs	
Static Compressive Load	44 10	222 50	N lbf	1, 2, 3, 4
	44 10	288 65	N lbf	1, 2, 3, 5
Dynamic Compressive Load		222 N + 0.45 kg * 100 G 50 lbf (static) + 1 lbm * 100 G	N lbf	1, 3, 4, 6, 7
		288 N + 0.45 kg * 100 G 65 lbf (static) + 1 lbm * 100 G	N lbf	1, 3, 5, 6, 7
Transient		445 100	N lbf	1, 3, 8

NOTES: In the case of a discrepancy, the most recent processor datasheet supersedes targets listed in the above table.

1. These specifications apply to uniform compressive loading in a direction perpendicular to the IHS top surface.
2. This is the minimum and maximum static force that can be applied by the heatsink and retention solution to maintain the heatsink and processor interface.
3. These parameters are based on limited testing for design characterization. Loading limits are for the package only and do not include the limits of the processor socket.
4. This specification applies for thermal retention solutions that allow baseboard deflection.
5. This specification applies either for thermal retention solutions that prevent baseboard deflection or for the Intel enabled reference solution.
6. Dynamic loading is defined as an 11 ms duration average load superimposed on the static load requirement.
7. Experimentally validated test condition used a heatsink mass of 1 lbm (~0.45 kg) with 100 G acceleration measured at heatsink mass. The dynamic portion of this specification in the product application can have flexibility in specific values, but the ultimate product of mass times acceleration should not exceed this validated dynamic load (1 lbm x 100 G = 100 lb).
8. Transient loading is defined as a 2 second duration peak load superimposed on the static load requirement, representative of loads experienced by the package during heatsink installation.

2.1.2 Critical Interface Dimensions (CID)

The CID drawing illustrates the key interfaces between the package and the thermal/mechanical solution for the processor. Should there be any conflict, this drawing is superseded by the drawing in the processor datasheet.

Figure 2-1. Critical Interface Dimensions (Sheet 1 of 2)

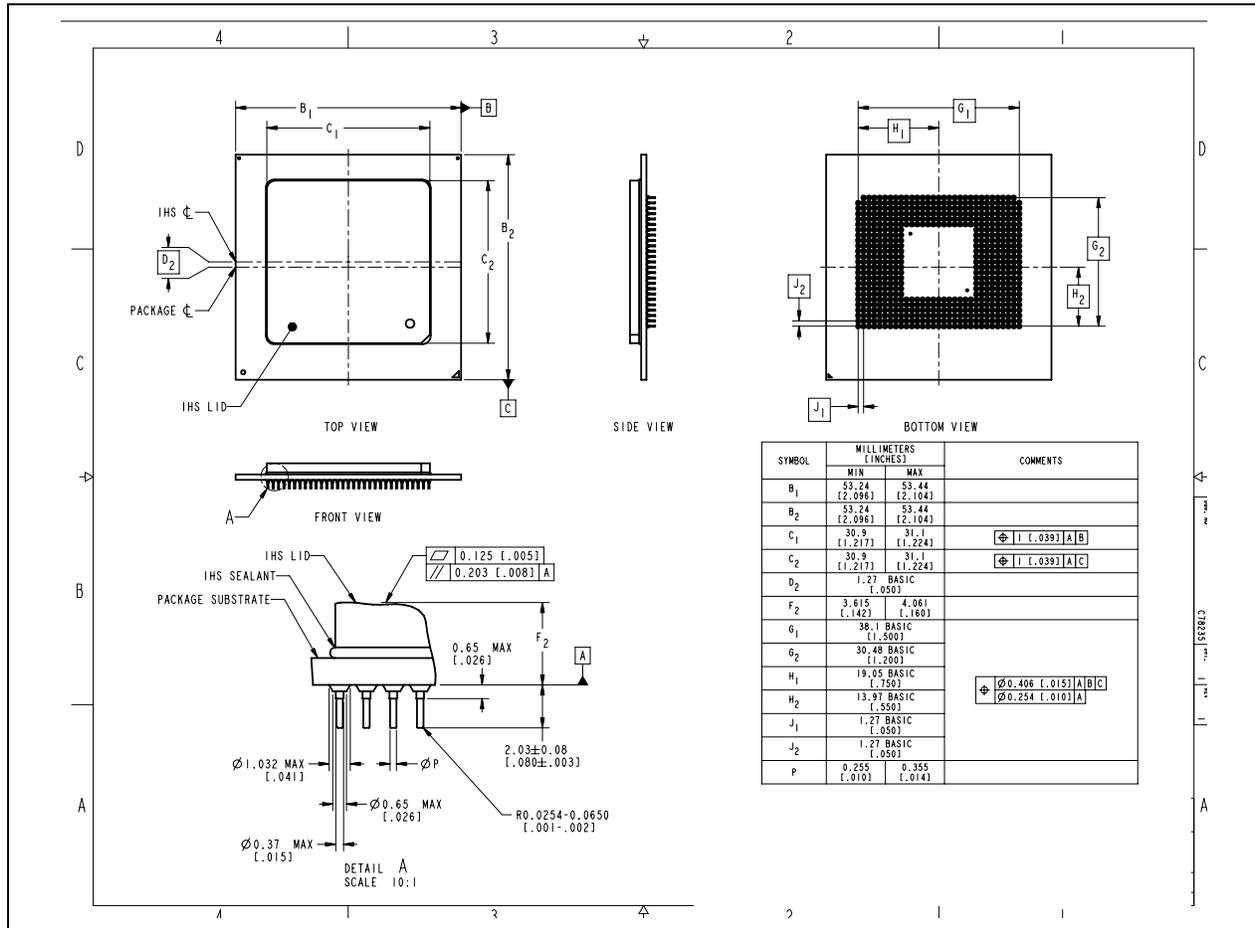
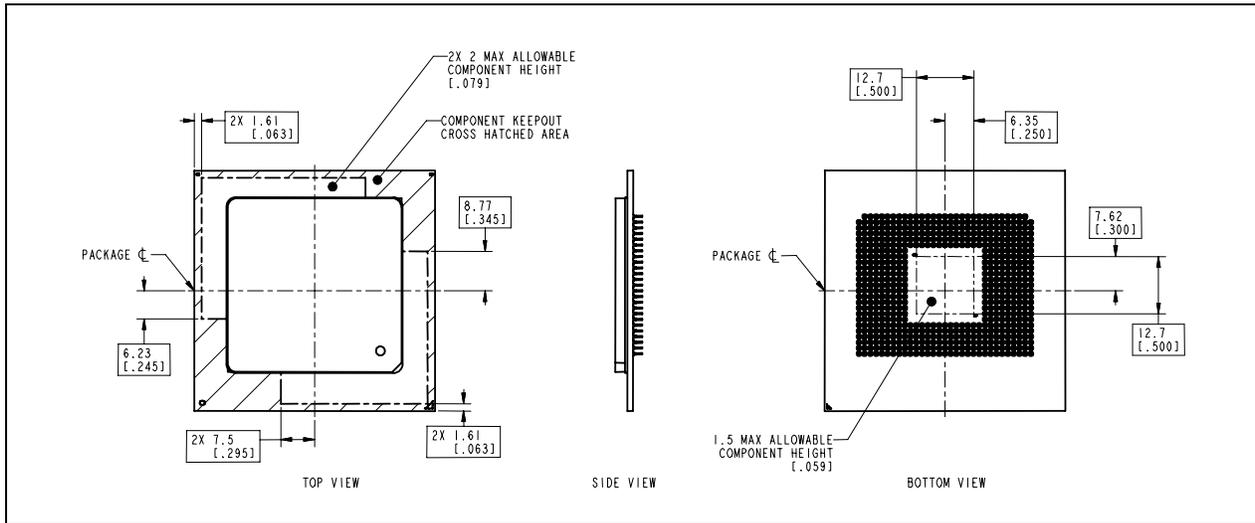


Figure 2-2. Critical Interface Dimensions (Sheet 2 of 2)


2.2 Thermal Requirements

In order to remain within a certain case temperature (T_{CASE}) specification to achieve optimal operation and long-term reliability, the thermal specification methodology, referred to as the thermal profile, is used. The intent of the thermal profile specification is to support acoustic noise reduction through fan speed control and ensure the long-term reliability of the processor. For information on thermal testing, please see Appendix B.

To ease the burden on thermal solutions, the Thermal Monitor feature and associated logic have been integrated into the silicon of the processor. By taking advantage of the Thermal Monitor feature, system designers may reduce thermal solution cost by designing to Thermal Design Power (TDP) instead of maximum power. The TDP is defined as the power level at which the processor thermal solutions be designed to dissipate. TDP is not the maximum power that the processor can dissipate. TDP is based on measurements of processor power consumption while running various high power applications. This data set is used to determine those applications that are interesting from a power perspective. These applications are then evaluated in a controlled thermal environment to determine their sensitivity to activation of the thermal control circuit. This data set is then used to derive the TDP targets published in the processor datasheet. Thermal Monitor can protect the processor in rare workload excursions above TDP. Therefore, thermal solutions should be designed to dissipate the target TDP level.

The relationship between TDP to the thermal profile, and thermal management logic and thermal monitor features, is discussed in the sections to follow. The thermal management logic and thermal monitor features are discussed in extensive detail in Appendix F.

2.2.1 Thermal Profile

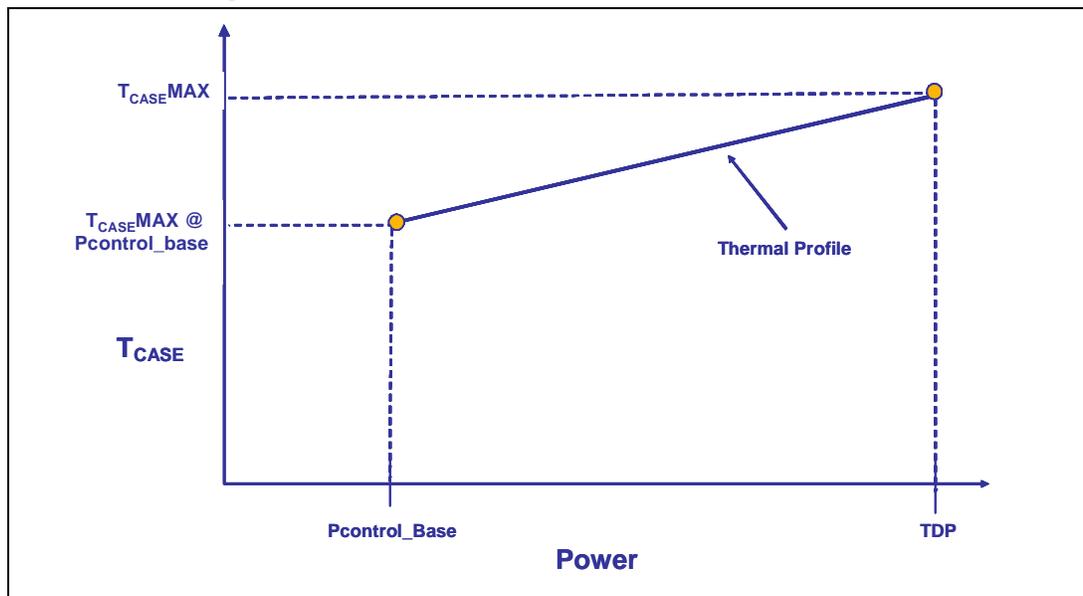
The thermal profile is a linear line that defines the relationship between a processor's case temperature and its power consumption as shown in Figure 2-3. The equation of the thermal profile is defined as:

$$y = ax + b \quad \text{Equation 1}$$

Where:

- y = Processor case temperature, T_{CASE} ($^{\circ}C$)
- x = Processor power consumption (W)
- a = Case-to-ambient thermal resistance, Ψ_{CA} ($^{\circ}C/W$)
- b = Processor local ambient temperature, T_{LA} ($^{\circ}C$)

Figure 2-3. Thermal Profile Diagram



The higher end point of the Thermal Profile represents the processor's TDP and the associated maximum case temperature ($T_{CASE\ MAX}$). The lower end point of the Thermal Profile represents the power value ($P_{control_base}$) and the associated case temperature ($T_{CASE\ MAX\ @\ P_{control_base}}$) for the lowest possible theoretical value of $T_{CONTROL}$ (see Section 2.2.3). This point is also associated with the $T_{CONTROL}$ value defined in Section 2.2.2. The slope of the Thermal Profile line represents the case-to-ambient resistance of the thermal solution with the y-intercept being the local processor ambient temperature. The slope of the Thermal Profile is constant between $P_{CONTROL\ BASE}$ and TDP, which indicate that all frequencies of a processor defined by the Thermal Profile will require the same heatsink case-to-ambient resistance.

To satisfy the Thermal Profile specification, a thermal solution must be at or below the Thermal Profile line for the given processor when its diode temperature is greater than $T_{CONTROL}$ (refer to Section 2.2.2). The Thermal Profile allows the customers to make a trade-off between the thermal solution case-to-ambient resistance and the processor local ambient temperature that best suits their platform implementation (refer to Section 2.3.3). There can be multiple combinations of thermal solution case-to-ambient resistance and processor local ambient temperature that can meet a given Thermal Profile. If the case-to-ambient resistance and the local ambient temperature are known for a specific thermal solution, the Thermal Profile of that solution can easily be plotted against the Thermal Profile specification. As explained above, the case-to-ambient resistance represents the slope of the line and the processor local ambient temperature represents the y-axis intercept. Hence the T_{CASE} values of a specific solution can be calculated at the TDP and $P_{control_base}$ power levels. Once these points are determined, they can be joined by a line, which represents the Thermal Profile of the specific solution. If that line stays at or below the Thermal Profile specification, then that particular solution is deemed as a compliant solution.

2.2.2 T_{CONTROL} Definition

T_{CONTROL} is a temperature specification based on a temperature reading from the processor's thermal diode. T_{CONTROL} defines the lower end of the Thermal Profile line for a given processor, and it can be described as a trigger point for fan speed control implementation. The value for T_{CONTROL} is calibrated in manufacturing and configured for each processor individually. For the 64-bit Intel Xeon processor MP with 1 MB L2 cache, the T_{CONTROL} value is obtained by reading a processor model specific register (MSR) and adding this offset value to a base value. The equation for calculating T_{CONTROL} is:

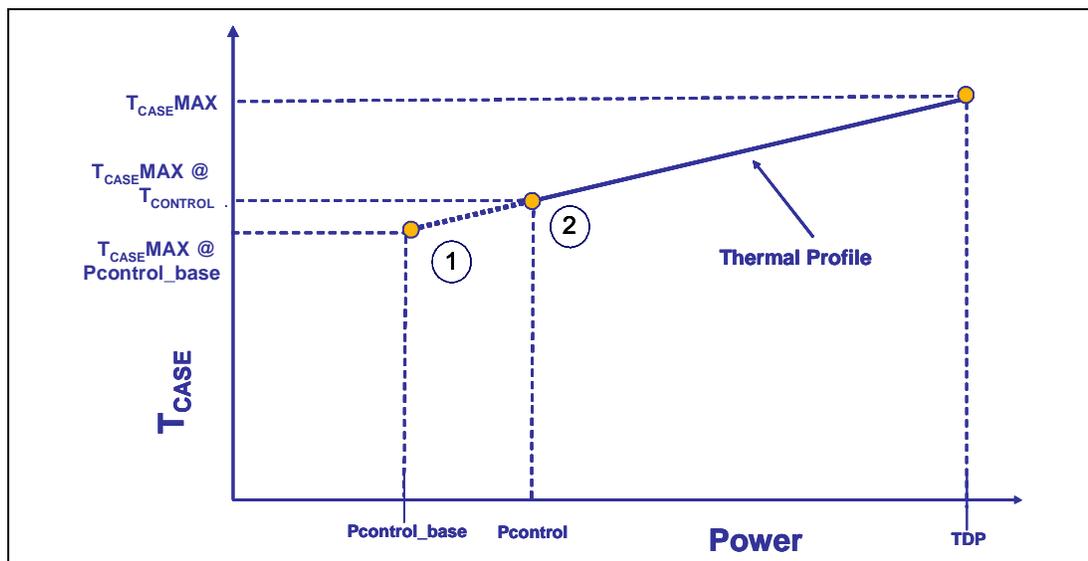
$$T_{\text{CONTROL}} = T_{\text{CONTROL_BASE}} + \text{Offset} \quad \text{Equation 2}$$

Where:

- $T_{\text{CONTROL_BASE}}$ = A fixed base value defined for a given processor generation as published in the processor datasheet.
- Offset = A value programmed into each processor during manufacturing that can be obtained by reading the IA32_TEMPERATURE_TARGET MSR. This is a static and a unique value.

The $T_{\text{CONTROL_BASE}}$ value for this processor is 50°C. The Offset value, which depends on several factors (i.e. leakage current) can be any number between 0 and $(T_{\text{CASE_MAX}} - T_{\text{CONTROL_BASE}})$. Figure 2-4 depicts the interaction between the Thermal Profile and T_{CONTROL} for an Offset value that is greater than 0 (i.e. T_{CONTROL} greater than $T_{\text{CONTROL_BASE}}$).

Figure 2-4. T_{CONTROL} and Thermal Profile Interaction



Since T_{CONTROL} is a processor diode temperature value, an equivalent T_{CASE} temperature must be determined to plot the $T_{\text{CASE_MAX}} @ T_{\text{CONTROL}}$ point on the Thermal Profile graph. Location 1 on the Thermal Profile represents a T_{CASE} value corresponding to an Offset of 0 (the theoretical minimum for the given processor family). Any Offset value greater than 0 moves the point where the Thermal Profile must be met upwards, as shown by location 2 on the graph. If the diode temperature is less than T_{CONTROL} , the case temperature is permitted to exceed the Thermal Profile, but the diode temperature must remain at or below T_{CONTROL} . In other words, there is no T_{CASE} specification for the processor at power levels less than P_{control} . The thermal solution for the processor must be able to keep the processor's T_{CASE} at or below the T_{CASE} values defined by the

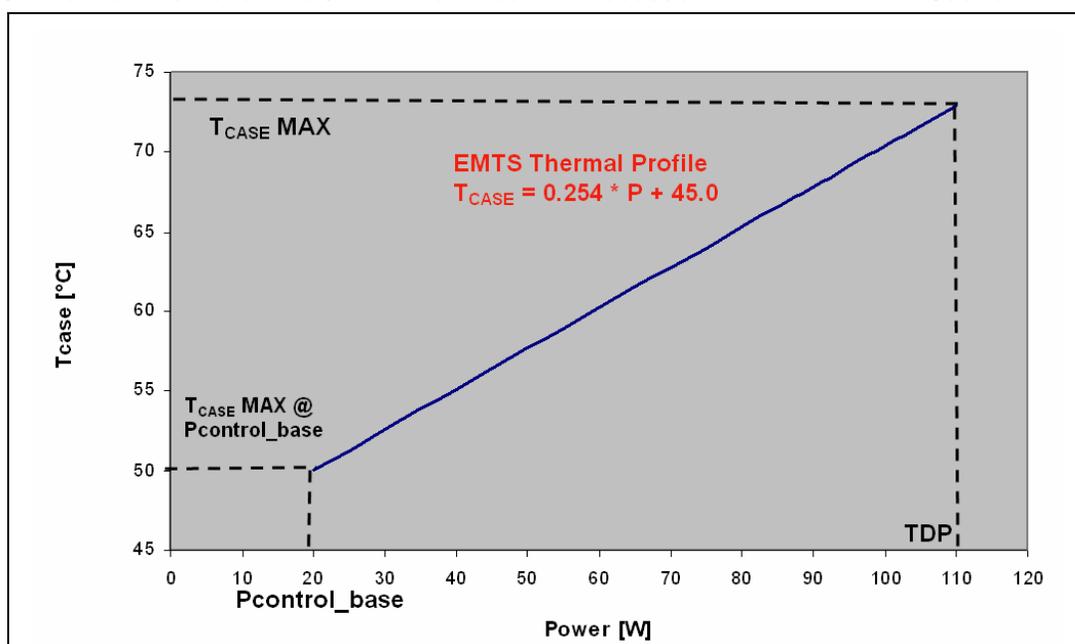
Thermal Profile between the $T_{CASE}^{MAX}@T_{CONTROL}$ and T_{CASE}^{MAX} points at the corresponding power levels.

Refer to Section 2.3.1 for the implementation of the $T_{CONTROL}$ value in support of fan speed control (FSC) design to achieve better acoustic performance.

2.2.3 Performance Targets

The Thermal Profile specification for this processor is published in the *64-bit Intel® Xeon™ Processor MP with 1 MB L2 Cache Datasheet*. The Thermal Profile specification is shown as a reference in the subsequent discussions.

Figure 2-5. Thermal Profile for the 64-bit Intel® Xeon™ Processor MP with 1 MB L2 Cache



NOTE: The thermal specification shown in this graph is for reference only. Refer to the *64-bit Intel® Xeon™ Processor MP with 1 MB L2 Cache Datasheet* for the Thermal Profile specification. In case of conflict, the data information in the datasheet supersedes any data in this figure.

Table 2-2 describes thermal performance targets for the processor cooling solution enabled by Intel.

Table 2-2. Performance Target Table (Sheet 1 of 2)

Parameter	Minimum	Maximum	Unit	Notes
Ψ_{CA}		0.299	°C/W	Mean + 3 σ (non-uniform heating)
Pressure Drop		0.15	In. H ₂ O	
Altitude		Sea-level		Heatsink designed at 0 meters
Airflow		23	CFM	Airflow through the heatsink fins
TIM-2 Dispense Weight		400	mg	Shin-Etsu®G751. Dispense weight is an approximate target.
TIM-2 Compressive Load	33 147	50 222	lbf N	Generated by the cooling solution.

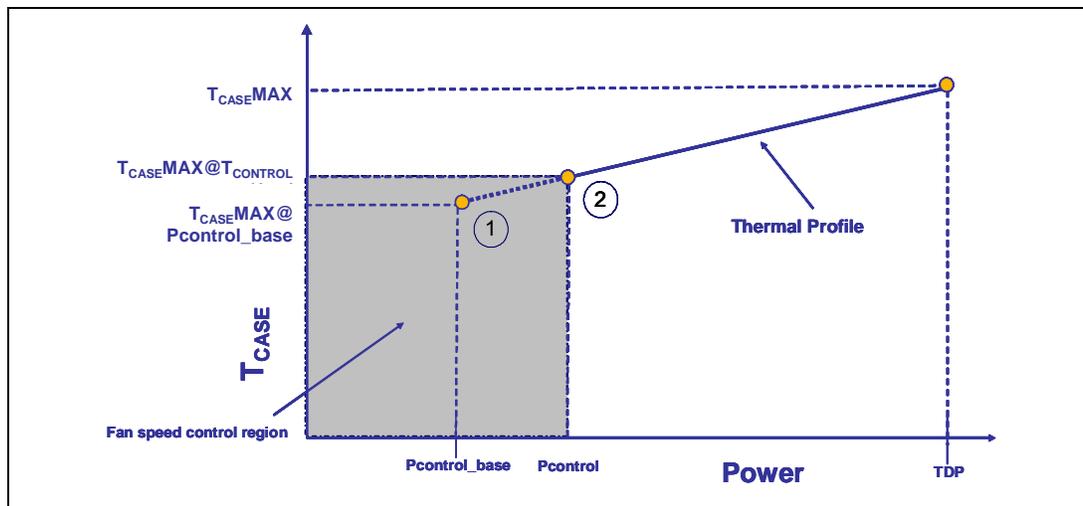
Table 2-2. Performance Target Table (Sheet 2 of 2)

Parameter	Minimum	Maximum	Unit	Notes
T_{CASE_MAX}		73	°C	In case of conflict, datasheet supercedes TMDG.
T_{CASE_MAX} @ $P_{control_base}$		50	°C	$P_{control_base} = 20\text{ W}$
T_{LA}		40	°C	
TDP		110	W	In case of conflict, datasheet supercedes TMDG.

2.3 Characterizing Cooling Solution Performance Requirements

2.3.1 Fan Speed Control

Fan speed control (FSC) techniques to reduce system level acoustic noise are a common practice in server designs. The fan speed is one of the parameters that determine the amount of airflow provided to the thermal solution. Additionally, airflow is proportional to a thermal solution's performance, which consequently determines the T_{CASE} of the processor at a given power level. Since the T_{CASE} of a processor is an important parameter in the long-term reliability of a processor, the FSC implemented in a system directly correlates to the processor's ability to meet the Thermal Profile and hence the long-term reliability requirements. For this purpose, the parameter called $T_{CONTROL}$ as explained in Section 2.2.2, is to be used in FSC designs to ensure that the long-term reliability of the processor is met while keeping the system level acoustic noise down. Figure 2-6 depicts the relationship between $T_{CONTROL}$ and FSC methodology.

Figure 2-6. $T_{CONTROL}$ and Fan Speed Control


Once the $T_{CONTROL}$ value is determined as explained earlier, the thermal diode temperature reading from the processor can be compared to this $T_{CONTROL}$ value. A fan speed control scheme can be implemented as described in Table 2-3 without compromising the long-term reliability of the processor.

Table 2-3. Fan Speed Control, T_{CONTROL} and T_{DIODE} Relationship

Condition	FSC Scheme
$T_{\text{DIODE}} \leq T_{\text{CONTROL}}$	FSC can adjust fan speed to maintain $T_{\text{DIODE}} \leq T_{\text{CONTROL}}$ (low acoustic region).
$T_{\text{DIODE}} > T_{\text{CONTROL}}$	FSC should adjust fan speed to keep T_{CASE} at or below the Thermal Profile specification (increased acoustic region).

There are many different ways of implementing fan speed control, including FSC based on processor ambient temperature; FSC based on processor thermal diode temperature (T_{DIODE}) or a combination of the two. If FSC is based only on the processor ambient temperature, low acoustic targets can be achieved under low ambient temperature conditions. However, the acoustics cannot be optimized based on the behavior of the processor temperature. If FSC is based only on the thermal diode, sustained temperatures above T_{CONTROL} , drives fans to maximum RPM. If FSC is based both on ambient and thermal diode, ambient temperature can be used to scale the fan RPM controlled by the thermal diode. This would result in an optimal acoustic performance. Regardless of which scheme is employed, system designers must ensure that the Thermal Profile specification is met when the processor diode temperature exceeds the T_{CONTROL} value for a given processor.

2.3.2 Processor Thermal Characterization Parameter Relationships

The idea of a “thermal characterization parameter”, Ψ (psi), is a convenient way to characterize the performance needed for the thermal solution and to compare thermal solutions in identical conditions (heating source, local ambient conditions). A thermal characterization parameter is convenient in that it is calculated using total package power, whereas actual thermal resistance, θ (theta), is calculated using actual power dissipated between two points. Measuring actual power dissipated by the heatsink is difficult since some of the power is dissipated via heat transfer into the socket and board. Be aware, however, of the limitations of lumped parameters such as Ψ when it comes to a real design. Heat transfer is a three-dimensional phenomenon that can rarely be accurately and easily modeled by lump values.

The case-to-local ambient thermal characterization parameter value (Ψ_{CA}) is used as a measure of the thermal performance of the overall thermal solution that is attached to the processor package. It is defined by the following equation, and measured in units of $^{\circ}\text{C}/\text{W}$:

$$\Psi_{\text{CA}} = (T_{\text{CASE}} - T_{\text{LA}}) / \text{TDP} \quad \text{Equation 3}$$

Where:

- Ψ_{CA} = Case-to-local ambient thermal characterization parameter ($^{\circ}\text{C}/\text{W}$).
- T_{CASE} = Processor case temperature ($^{\circ}\text{C}$).
- T_{LA} = Local ambient temperature in chassis at processor ($^{\circ}\text{C}$).
- P_{D} = TDP dissipation (W) (assumes all power dissipates through the integrated heat spreader (IHS)).

The case-to-local ambient thermal characterization parameter of the processor, Ψ_{CA} , is comprised of Ψ_{CS} , the TIM thermal characterization parameter, and of Ψ_{SA} , the sink-to-local ambient thermal characterization parameter:

$$\Psi_{\text{CA}} = \Psi_{\text{CS}} + \Psi_{\text{SA}} \quad \text{Equation 4}$$

Where:

Ψ_{CS} = Thermal characterization parameter of the TIM ($^{\circ}\text{C}/\text{W}$).

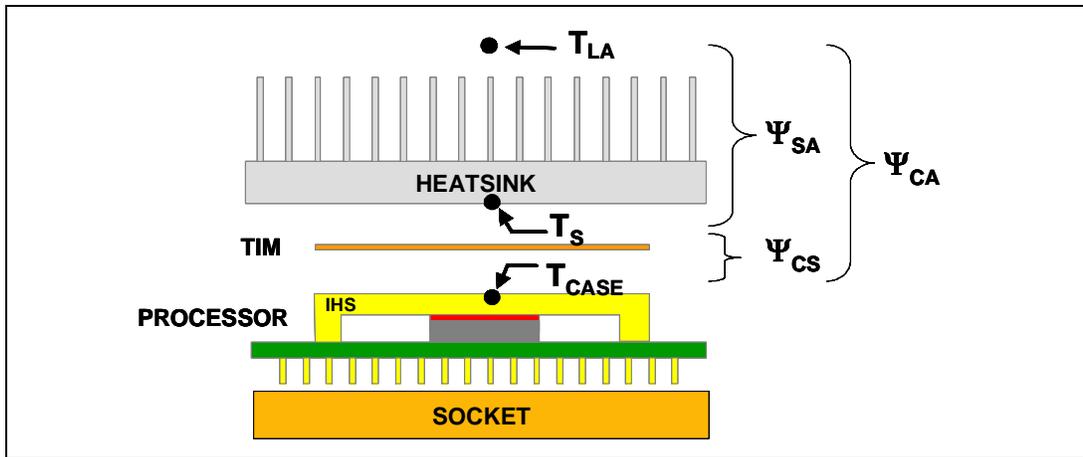
Ψ_{SA} = Thermal characterization parameter from heatsink-to-local ambient ($^{\circ}\text{C}/\text{W}$).

Ψ_{CS} is strongly dependent on the thermal conductivity and thickness of the TIM between the heatsink and IHS.

Ψ_{SA} is a measure of the thermal characterization parameter from the bottom of the heatsink to the local ambient air. Ψ_{SA} is dependent on the heatsink material, thermal conductivity, and geometry. It is also strongly dependent on the air velocity through the fins of the heatsink and the local ambient temperature surrounding the heatsink.

Figure 2-7 illustrates the combination of the different thermal characterization parameters.

Figure 2-7. Processor Thermal Characterization Parameter Relationships



Example

The cooling performance, Ψ_{CA} , is then defined using the principle of thermal characterization parameter described above:

- Define a target case temperature $T_{CASE-MAX}$ and corresponding TDP at a target frequency, F , given in the processor datasheet.
- Define a target local ambient temperature at the processor, T_{LA} .

Since the processor thermal specifications ($T_{CASE-MAX}$ and TDP) can vary with the processor frequency, it may be important to identify the worse case (lowest Ψ_{CA}) for a targeted chassis (characterized by T_{LA}) to establish a design strategy such that a given heatsink can cover a given range of processor frequencies.

The following provides an illustration of how one might determine the appropriate performance targets. The example power and temperature numbers used here are not related to any Intel processor thermal specifications, and are for illustrative purposes only.

Assume the datasheet TDP is 85 W and the case temperature specification is 68 $^{\circ}\text{C}$ for a given frequency. Assume as well that the system airflow has been designed such that the local processor ambient temperature is 45 $^{\circ}\text{C}$. The following could be calculated using equation 1 from above for the given frequency:

$$\Psi_{CA} = (T_{CASE} - T_{LA}) / TDP = (68 - 45) / 85 = 0.27 \text{ }^{\circ}\text{C/W}$$

To determine the required heatsink performance, a heatsink solution provider would need to determine Ψ_{CS} performance for the selected TIM and mechanical load configuration. If the heatsink solution was designed to work with a TIM material performing at $\Psi_{CS} \leq 0.05 \text{ }^{\circ}\text{C/W}$, solving for equation 2 from above, the performance of the heatsink would be:

$$\Psi_{SA} = \Psi_{CA} - \Psi_{CS} = 0.27 - 0.05 = 0.22 \text{ }^{\circ}\text{C/W}$$

If the local processor ambient temperature is assumed to be 40°C , the same calculation can be carried out to determine the new case-to-ambient thermal resistance:

$$\Psi_{CA} = (T_{CASE} - T_{LA}) / TDP = (68 - 40) / 85 = 0.33 \text{ }^{\circ}\text{C/W}$$

It is evident from the above calculations that, a reduction in the local processor ambient temperature has a significant positive effect on the case-to-ambient thermal resistance requirement.

2.3.3 Chassis Thermal Design Considerations

2.3.3.1 Chassis Thermal Design Capabilities and Improvements

One of the critical parameters in thermal design is the local ambient temperature assumption of the processor. Keeping the external chassis temperature fixed, internal chassis temperature rise is the only component that can affect the processor local ambient temperature. Every degree gained at the local ambient temperature directly translates into a degree relief in the processor case temperature.

Given the thermal targets for the processor, it is extremely important to optimize the chassis design to minimize the air temperature rise upstream to the processor (T_{RISE}), hence minimizing the processor local ambient temperature. Please refer to Appendix B.

The heat generated by components within the chassis must be removed to provide an adequate operating environment for both the processor and other system components. Moving air through the chassis brings in air from the external ambient environment and transports the heat generated by the processor and other system components out of the system. The number, size and relative position of fans, vents and other heat generating components determine the chassis thermal performance, and the resulting ambient temperature around the processor. The size and type (passive or active) of the thermal solution and the amount of system airflow can be traded off against each other to meet specific system design constraints. Additional constraints are board layout, spacing, component placement, and structural considerations that limit the thermal solution size.

In addition to passive heatsinks, fan heatsinks and system fans, other solutions exist for cooling integrated circuit devices. For example, ducted blowers, heat pipes and liquid cooling are all capable of dissipating additional heat. Due to their varying attributes, each of these solutions may be appropriate for a particular system implementation.

To develop a reliable, cost-effective thermal solution, thermal characterization and simulation should be carried out at the entire system level, accounting for the thermal requirements of each component. In addition, acoustic noise constraints may limit the size, number, placement, and types of fans that can be used in a particular design.

2.4 Thermal/Mechanical Reference Design Considerations

2.4.1 Heatsink Solutions

2.4.1.1 Heatsink Design Considerations

To remove the heat from the processor, three basic parameters should be considered:

- **The area of the surface on which the heat transfer takes place** - Without any enhancements, this is the surface of the processor package IHS. One method used to improve thermal performance is by attaching a heatsink to the IHS. A heatsink can increase the effective heat transfer surface area by conducting heat out of the IHS and into the surrounding air through fins attached to the heatsink base.
- **The conduction path from the heat source to the heatsink fins** - Providing a direct conduction path from the heat source to the heatsink fins and selecting materials with higher thermal conductivity typically improves heatsink performance. The length, thickness, and conductivity of the conduction path from the heat source to the fins directly impact the thermal performance of the heatsink. In particular, the quality of the contact between the package IHS and the heatsink base has a higher impact on the overall thermal solution performance as processor cooling requirements become strict. Thermal interface material (TIM) is used to fill in the gap between the IHS and the bottom surface of the heatsink, and thereby improves the overall performance of the thermal stackup (IHS-TIM-Heatsink). With extremely poor heatsink interface flatness or roughness, TIM may not adequately fill the gap. The TIM thermal performance depends on its thermal conductivity as well as the pressure load applied to it. Refer to Section 2.4.2 for further information on the TIM between the IHS and the heatsink base.
- **The heat transfer conditions on the surface on which heat transfer takes place** - Convective heat transfer occurs between the airflow and the surface exposed to the flow. It is characterized by the local ambient temperature of the air, T_{LA} , and the local air velocity over the surface. The higher the air velocity over the surface, the resulting cooling is more efficient. The nature of the airflow can also enhance heat transfer via convection. Turbulent flow can provide improvement over laminar flow. In the case of a heatsink, the surface exposed to the flow includes the fin faces and the heatsink base.

An active heatsink typically incorporates a fan that helps manage the airflow through the heatsink.

Passive heatsink solutions require in-depth knowledge of the airflow in the chassis. Typically, passive heatsinks see slower air speed. Therefore these heatsinks are typically larger (and heavier) than active heatsinks due to the increase in fin surface required to meet a required performance. As the heatsink fin density (the number of fins in a given cross-section) increases, the resistance to the airflow increases: it is more likely that the air will travel around the heatsink instead of through it, unless air bypass is carefully managed. Using air-ducting techniques to manage bypass area is an effective method for maximizing airflow through the heat sink fins.

2.4.2 Thermal Interface Material

TIM application between the processor IHS and the heatsink base is generally required to improve thermal conduction from the IHS to the heatsink. Many thermal interface materials can be pre-applied to the heatsink base prior to shipment from the heatsink supplier and allow direct heatsink attach, without the need for a separate TIM dispense or attach process in the final assembly factory.

All thermal interface materials should be sized and positioned on the heatsink base in a way that ensures the entire processor IHS area is covered. It is important to compensate for heatsink-to-processor attach positional alignment when selecting the proper TIM size.

When pre-applied material is used, it is recommended to have a protective application tape over it. This tape must be removed prior to heatsink installation.

The TIM performance is susceptible to degradation (i.e. grease breakdown) during the useful life of the processor due to the temperature cycling phenomena. For this reason, the measured T_{CASE} value of a given processor can decrease over time depending on the type of TIM material.

2.4.3 Summary

In summary, considerations in heatsink design include:

- The local ambient temperature T_{LA} at the heatsink, airflow (CFM), the power being dissipated by the processor, and the corresponding maximum T_{CASE} . These parameters are usually combined in a single lump cooling performance parameter, Ψ_{CA} (case to air thermal characterization parameter). More information on the definition and the use of Ψ_{CA} is given in Section 1.4 and Section 2.3.2.
- Heatsink interface (to IHS) surface characteristics, including flatness and roughness.
- The performance of the TIM used between the heatsink and the IHS.
- Surface area of the heatsink.
- Heatsink material and technology.
- Development of airflow entering and within the heatsink area.
- Physical volumetric constraints placed by the system.

2.4.4 Assembly Overview of the Intel Reference Thermal Mechanical Design

The 2U+ cooling solution consists of the following components:

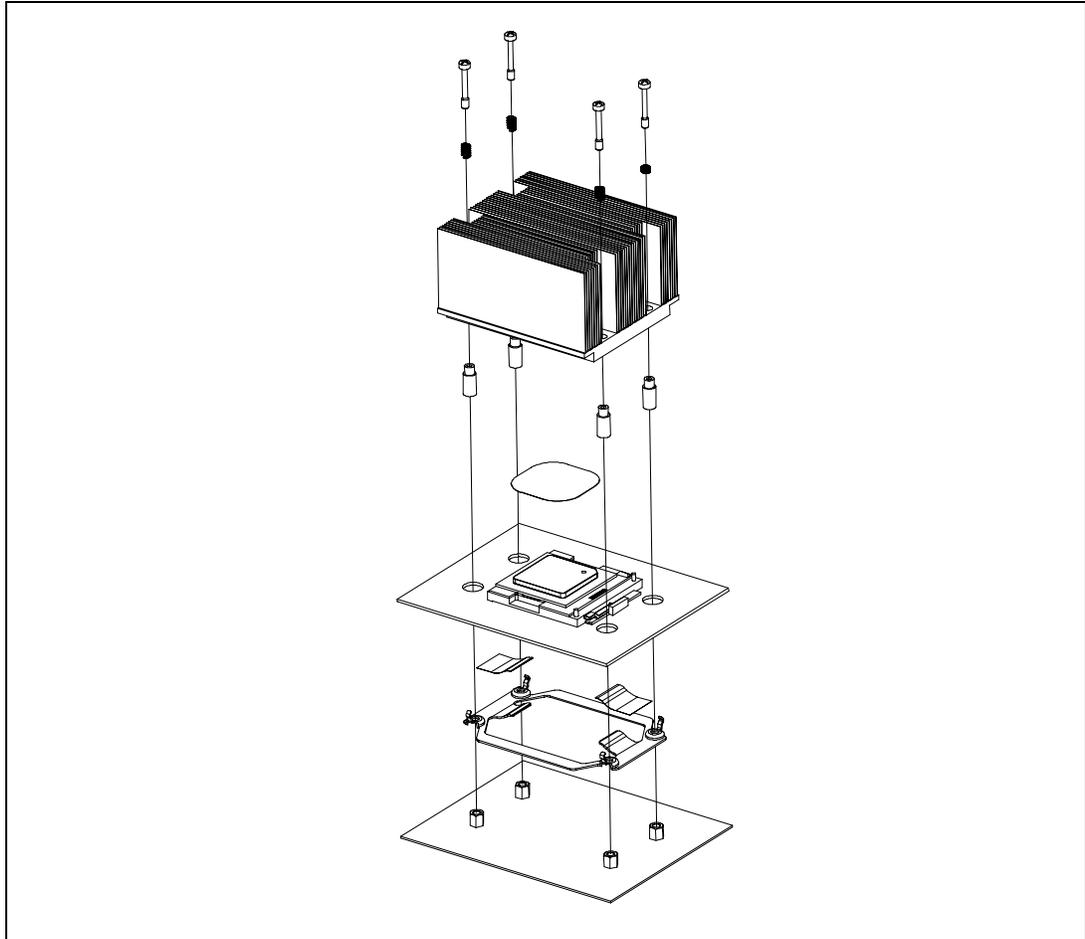
- Heatsink (with captive standoff and screws)
- Thermal Interface Material (TIM-2)
- Hat Spring

2.4.4.1 Geometric Envelope

The baseboard keepout zones on the primary and secondary sides and height restrictions under the enabling component region are shown in detail in Appendix A. The overall volumetric keep in zone encapsulates the processor, socket, and the entire thermal/mechanical enabling solution.

2.4.4.2 Assembly Drawing

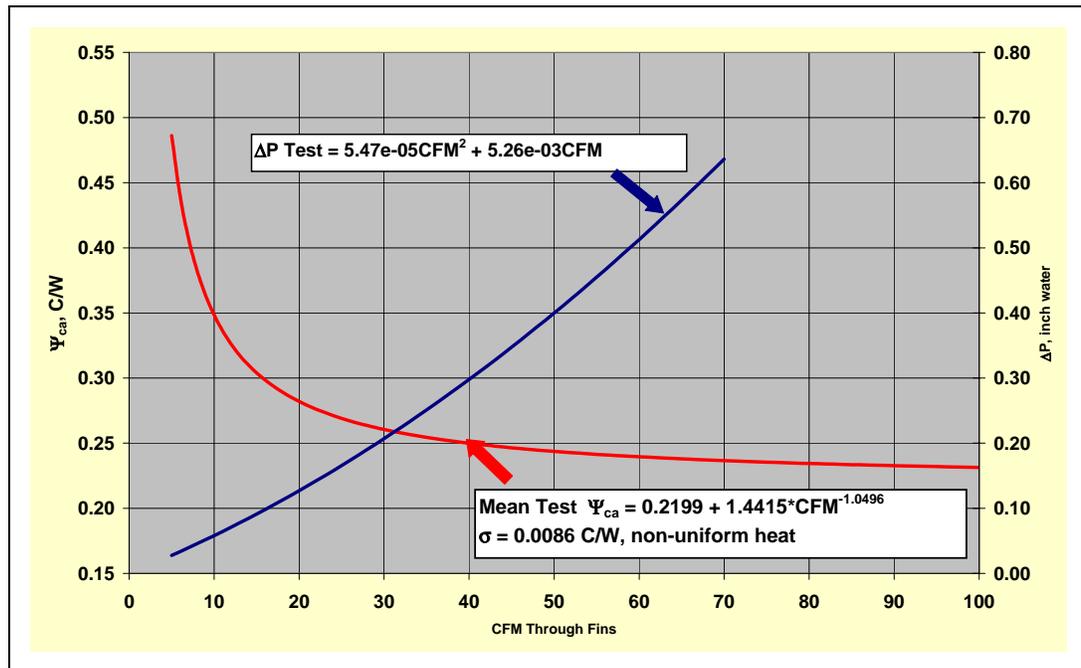
Figure 2-8. Exploded View of Cooling Solution Thermal Solution Components



2.4.5 Thermal Solution Performance Characteristics

The optimization of the cooling solution heatsink for thermal performance is completed and Figure 2-9 shows the thermal performance and the pressure drop through fins of the heatsink versus the airflow provided. The best-fit equations for these curves are also provided to make it easier for users to determine the desired value without any error associated with reading the graph.

Figure 2-9. 2U+ Cooling Solution Heatsink Thermal Performance



If other custom heatsinks are intended for use with the 64-bit Intel Xeon processor MP with 1 MB L2 cache, they must support the following interface control requirements to be compatible with the reference mechanical components:

- **Requirement 1:** Heatsink assembly must stay within the volumetric keep-in.
- **Requirement 2:** Maximum mass and center of gravity:

Current maximum heatsink mass is 1000 grams [2.2 lbs] and the maximum center of gravity 3.81 cm [1.5 in.] above the bottom of the heatsink base.

- **Requirement 3:** Maximum and minimum compressive load:

Any custom thermal solution design should meet the loading specification as documented within this document, and should refer to the datasheet for specific details on package loading specifications.

2.4.6 Structural Considerations of Cooling Solution

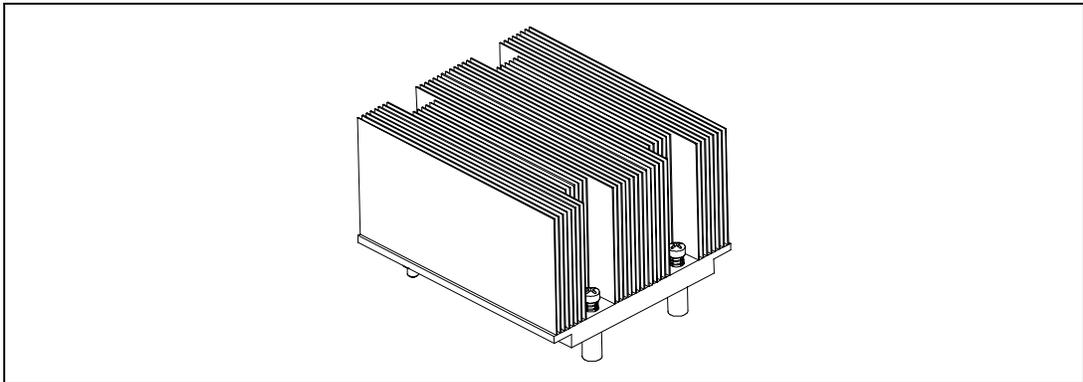
As Intel explores methods of keeping thermal solutions within the air-cooling space, the mass of the thermal solutions is increasing. Due to the flexible nature (and associated large deformation) of baseboard-only attachments, Intel reference solutions, such as cooling solution, are now commonly using direct chassis attach (DCA) as the mechanical retention design. The mass of the new thermal solutions is large enough to require consideration for structural support and stiffening on the chassis. Intel has published a best know method (BKM) document that provides specific structural guidance for designing DCA thermal solutions. The document is titled *Chassis Strength and Stiffness Measurement and Improvement Guidelines for Direct Chassis Attach Solutions*.

2.4.7 Components Overview

2.4.7.1 Heatsink with Captive Screws and Standoffs

The cooling solution reference heatsink uses snapped-fin technology for its design. It consists of a copper base and copper fins with Shin-Etsu* G751 thermal grease as the TIM. The mounting screws and standoffs are also made captive to the heatsink base for ease of handling and assembly as shown in Figure 2-10.

Figure 2-10. Isometric View of the 2U+ Cooling Solution Heatsink



NOTE: Refer to Appendix A for more detailed mechanical drawings of the heatsink.

The function of the standoffs is to provide a bridge between the chassis and the heatsink for attaching and load carrying. When assembled, the heatsink is rigid against the top of the standoff, and the standoff is rigid to a chassis standoff with the hat spring firmly sandwiched between the two. In dynamic loading situations the standoffs carry much of the heatsink load, especially in lateral conditions, when compared to the amount of load transmitted to the processor package. As such, it is comprised of steel. The distance from the bottom of the heatsink to the bottom of the standoff is 1.02 cm [0.402 in.].

The function of the screw is to provide a rigid attach method to sandwich the entire cooling solution assembly together, activating the hat spring under the baseboard, and thus providing the TIM preload. A screw is an inexpensive, low profile solution that does not negatively impact the thermal performance of the heatsink due to air blockage. Any fastener (i.e. head configuration) can be used as long as it is of steel construction; the head does not interfere with the heatsink fins, and is of the correct length of 1.27 cm [0.50 in.].

Although the cooling solution heatsink fits into the legacy volumetric keep-in, it has a larger footprint due to the elimination of retention mechanism and clips used in the older enabled thermal/mechanical components. This allows the heatsink to grow its base and fin dimensions, further improving the thermal performance. A drawback of this enlarged size and use of copper for both the base and fins is the increased weight of the heatsink. The cooling solution heatsink is estimated to weigh twice as much as previous heatsinks used with Intel Xeon processors. However, the retention scheme employed by cooling solution is designed to support heavy heatsinks (approximately up to 1000 grams) in cases of shock, vibration and installation as explained in Appendix D.

2.4.7.2 Thermal Interface Material (TIM-2)

A TIM must be applied between the package and the heatsink to ensure thermal conduction. The cooling solution reference design uses Shin-Etsu* G751 thermal grease.

The recommended grease dispenses weight to ensure full coverage of the processor IHS is given below. For an alternate TIM, full coverage of the entire processor IHS is recommended.

Table 2-4. Recommended Thermal Grease Dispense Weight

Processor	Recommended Thermal Grease	Dispense Weight (mg)
64-bit Intel® Xeon™ Processor MP with 1 MB L2 Cache	Shin-Etsu* G751	400

It is recommended that you use thermally conductive grease as the TIM requires special handling and dispense guidelines. The following guidelines apply to Shin-Etsu G751 thermal grease. The use of a semi-automatic dispensing system is recommended for high volume assembly to ensure an accurate amount of grease is dispensed on top of the IHS prior to assembly of the heatsink. A typical dispense system consists of an air pressure and timing controller, a hand held output dispenser, and an actuation foot switch. Thermal grease in cartridge form is required for dispense system compatibility. A precision scale with an accuracy of ± 5 mg is recommended to measure the correct dispense weight and set the corresponding air pressure and duration. The IHS surface should be free of foreign materials prior to grease dispense

Additional recommendations include recalibrating the dispense controller settings after any two hour pause in grease dispense. The grease should be dispensed just prior to heatsink assembly to prevent any degradation in material performance. Finally, the thermal grease should be verified to be within its recommended shelf life before use.

The cooling solution reference solution is designed to apply a compressive load of up to 222 N [50 lbf] on the TIM to improve the thermal performance.

2.4.7.3 Hat Spring

The hat spring, which is attached on the secondary side of the baseboard, is made from 0.80 mm [0.0315 in.] thick 301 stainless steel half hard. Any future versions of the spring will be made from a similar material. The hat spring has four embosses (called “hats”) which, when assembled, rest on the top of the chassis standoffs. The hat spring is located between the chassis standoffs and the heatsink standoffs. The purpose of the hat spring is to provide compressive preload at the TIM interface when the baseboard is pushed down upon it. This spring does not function as a clip of any kind. The two tabs on the spring are used to provide the necessary compressive preload for the TIM when the whole solution is assembled. The tabs make contact on the secondary side of the baseboard. In order to avoid damage to the contact locations on the baseboard, the tabs will be insulated with a 0.127 mm [0.005 in.] thick Kapton* tape (or equivalent). Figure 2-11 shows an isometric view of the hat spring design.

Figure 2-11. Hat Spring Isometric View

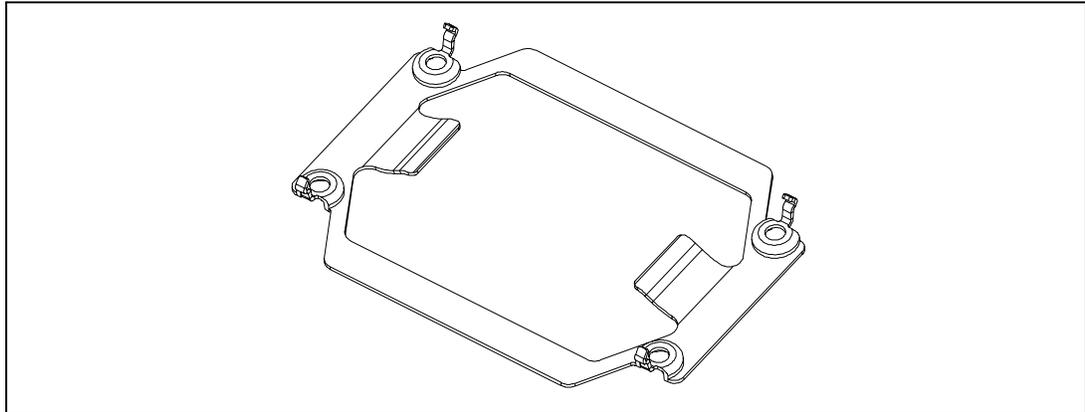
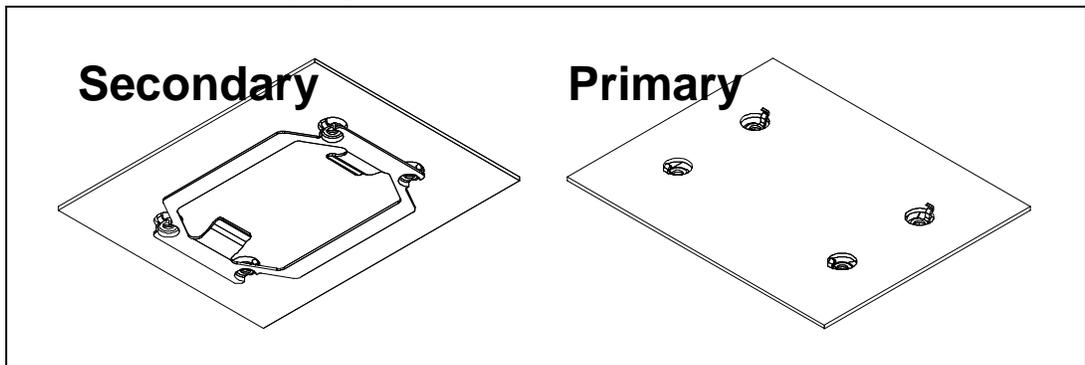


Figure 2-12. Isometric View of Hat Spring Attachment to the Base Board



Please refer to Appendix A for more detailed mechanical drawings of the hat spring. Also, the baseboard keepout requirements shown in Appendix A must be met to use this hat spring design.

§



A Mechanical Drawings

The mechanical drawings included in this appendix. These drawings refer to the thermal mechanical enabling components for the 64-bit Intel Xeon processor MP with 1 MB L2 cache.

Note: Intel reserves the right to make changes and modifications to the design as necessary.

Table A-1. Mechanical Drawing List

Drawing Description	Figure Number
2U Cooling Solution Heatsink (Sheet 1 of 4)	Figure A-1
2U Cooling Solution Heatsink (Sheet 2 of 4)	Figure A-2
2U Cooling Solution Heatsink (Sheet 3 of 4)	Figure A-3
2U Cooling Solution Heatsink (Sheet 4 of 4)	Figure A-4
Cooling Solution Hat Spring (Sheet 1 of 3)	Figure A-5
Cooling Solution Hat Spring (Sheet 2 of 3)	Figure A-6
Cooling Solution Hat Spring (Sheet 3 of 3)	Figure A-7
Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 1 of 5)	Figure A-8
Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 2 of 5)	Figure A-9
Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 3 of 5)	Figure A-10
Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 4 of 5)	Figure A-11
Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 5 of 5)	Figure A-12

Figure A-1. 2U Cooling Solution Heatsink (Sheet 1 of 4)

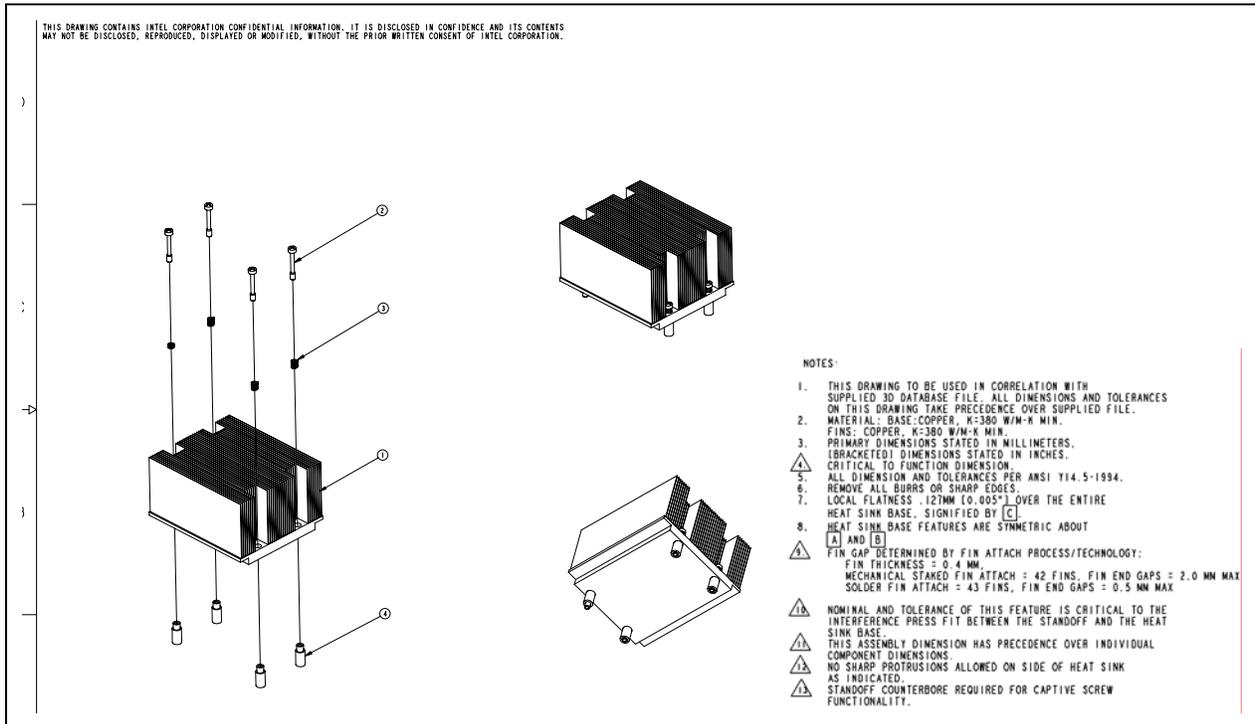


Figure A-2. 2U Cooling Solution Heatsink (Sheet 2 of 4)

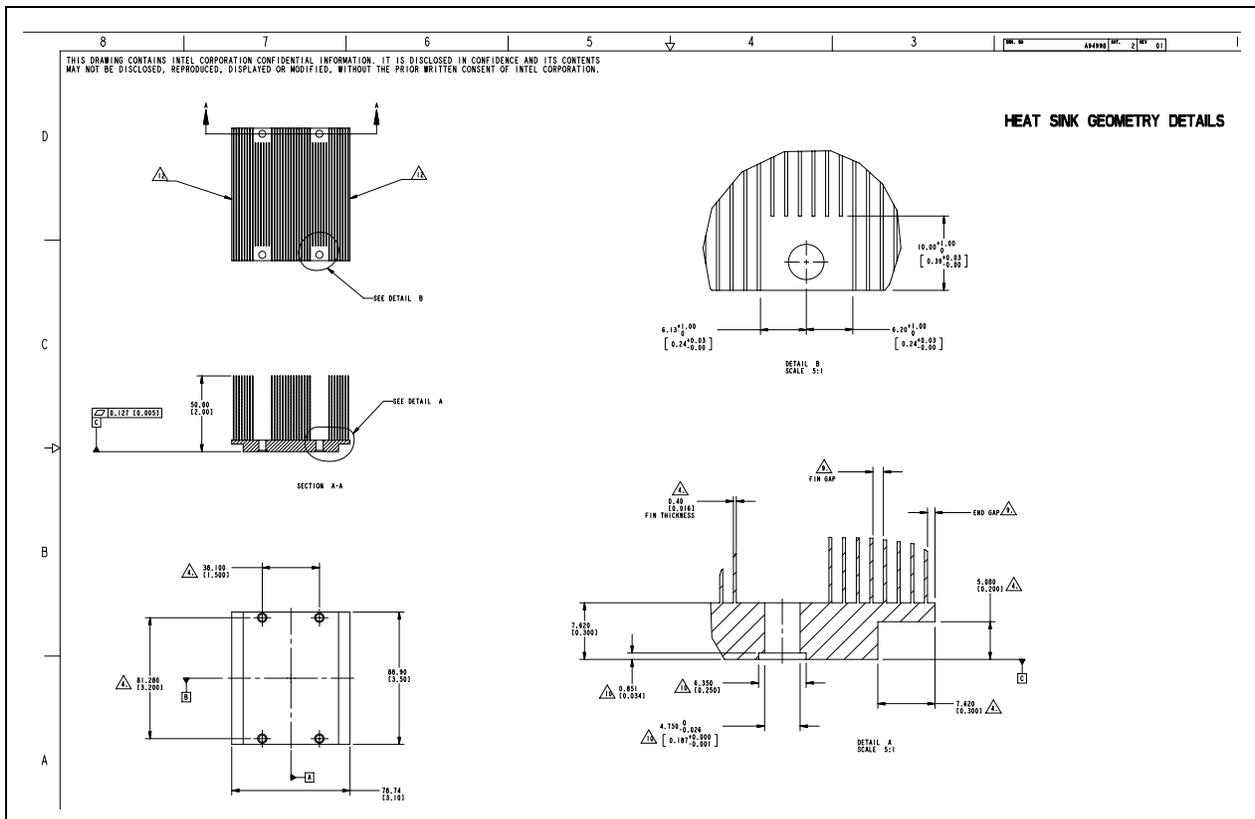


Figure A-3. 2U Cooling Solution Heatsink (Sheet 3 of 4)

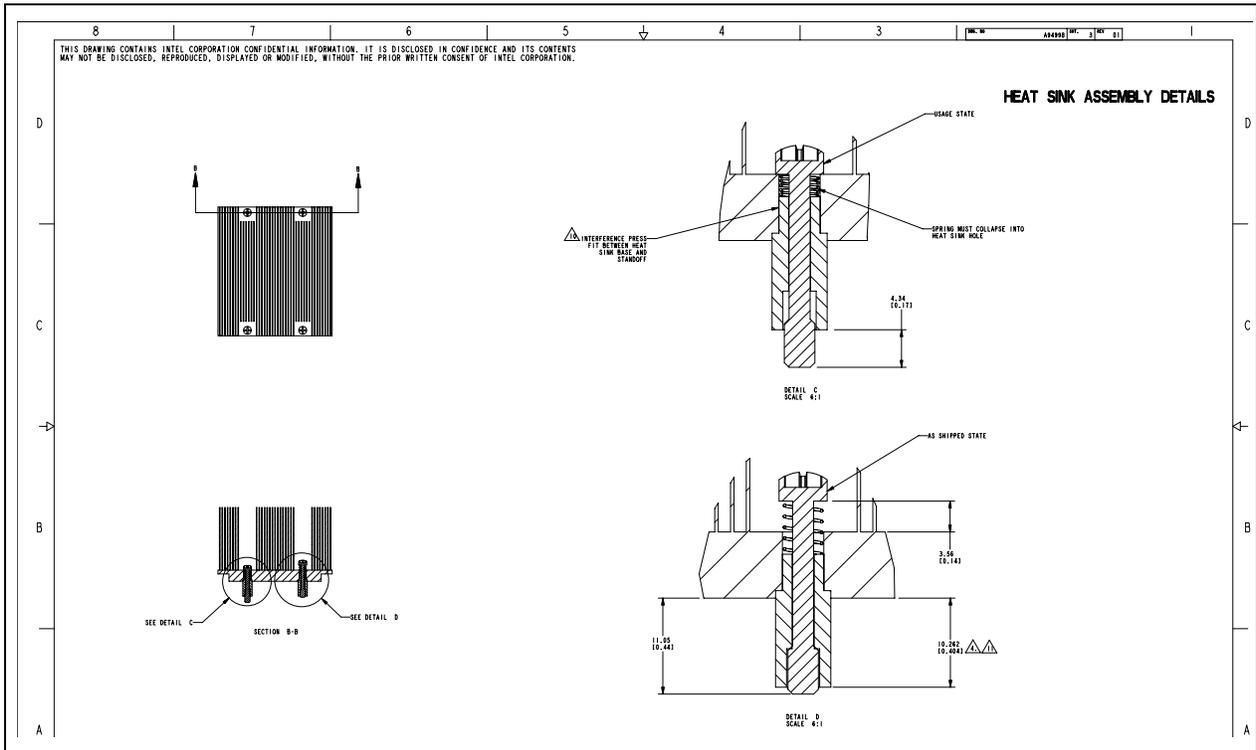


Figure A-4. 2U Cooling Solution Heatsink (Sheet 4 of 4)

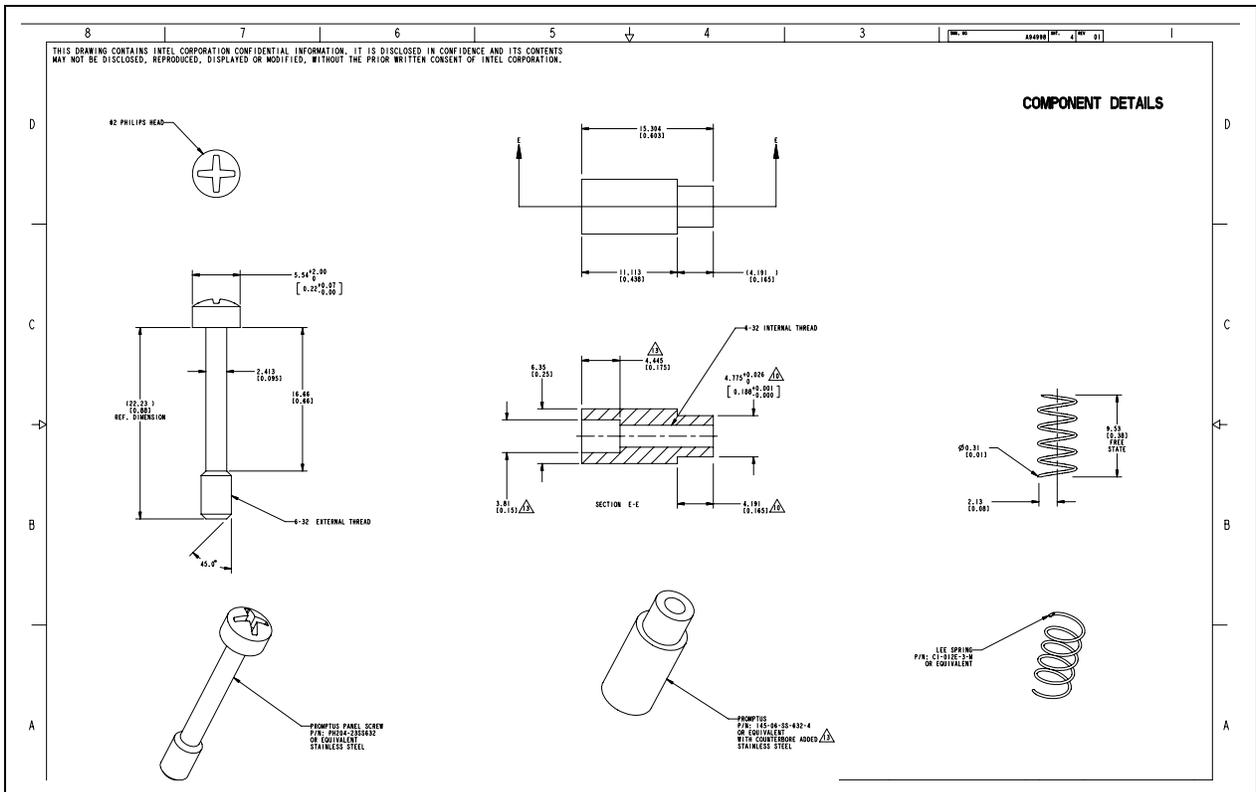


Figure A-5. Cooling Solution Hat Spring (Sheet 1 of 3)

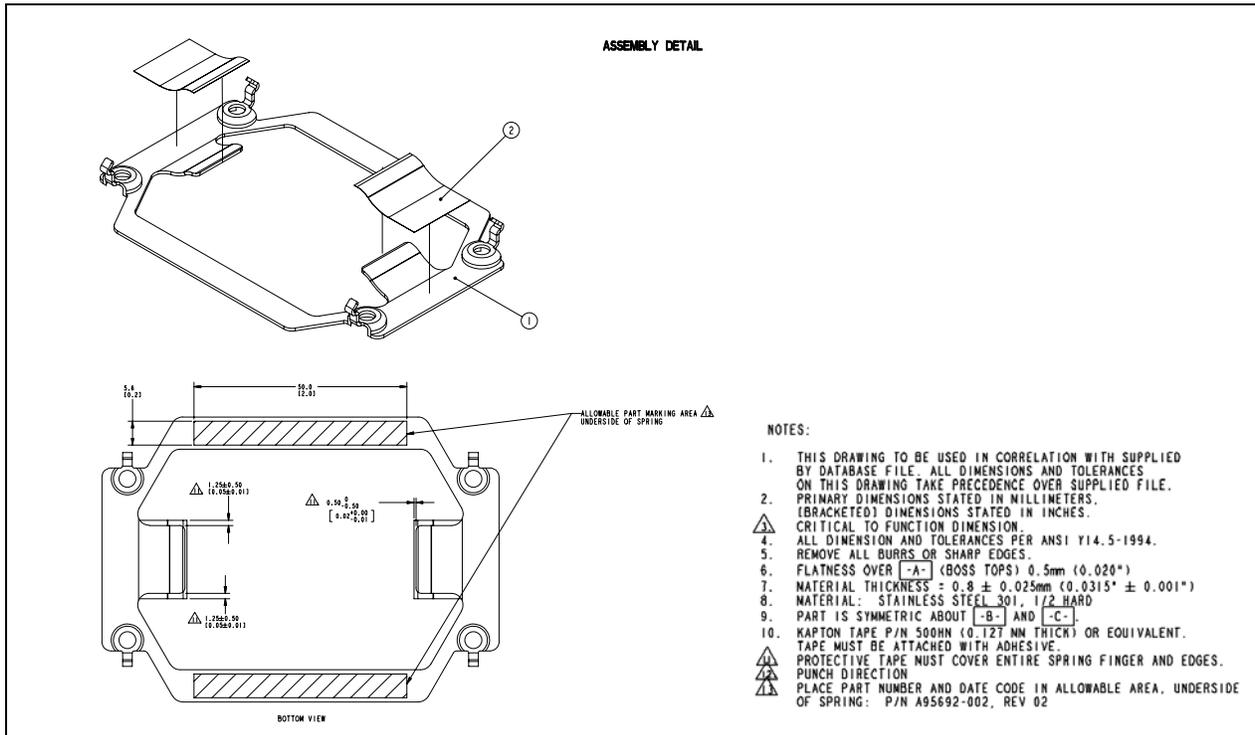


Figure A-6. Cooling Solution Hat Spring (Sheet 2 of 3)

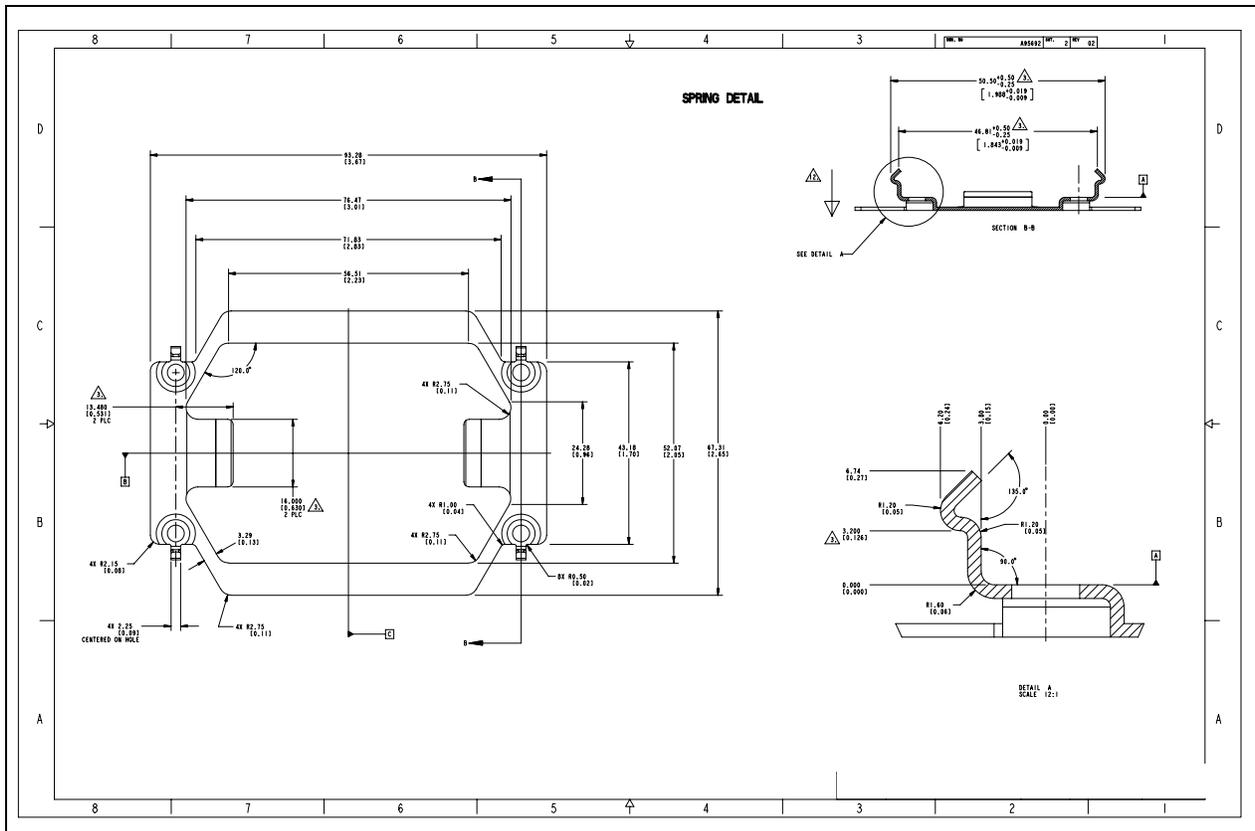


Figure A-7. Cooling Solution Hat Spring (Sheet 3 of 3)

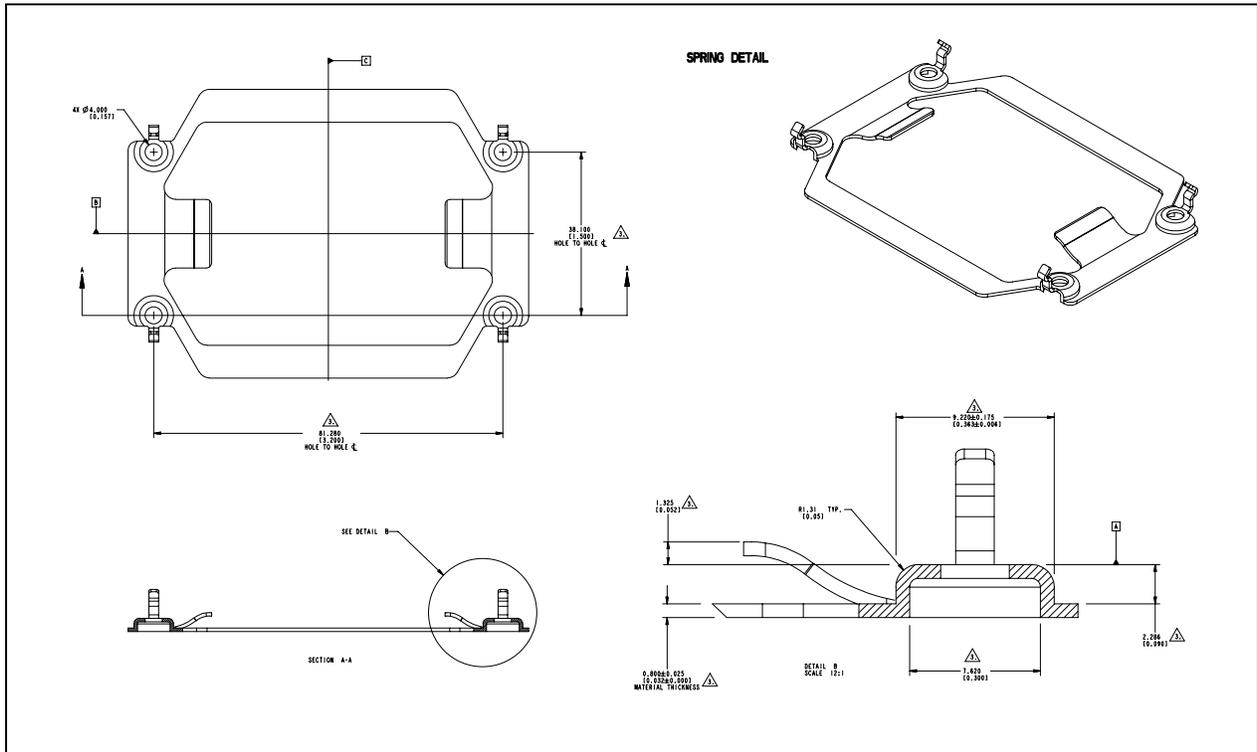


Figure A-8. Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 1 of 5)

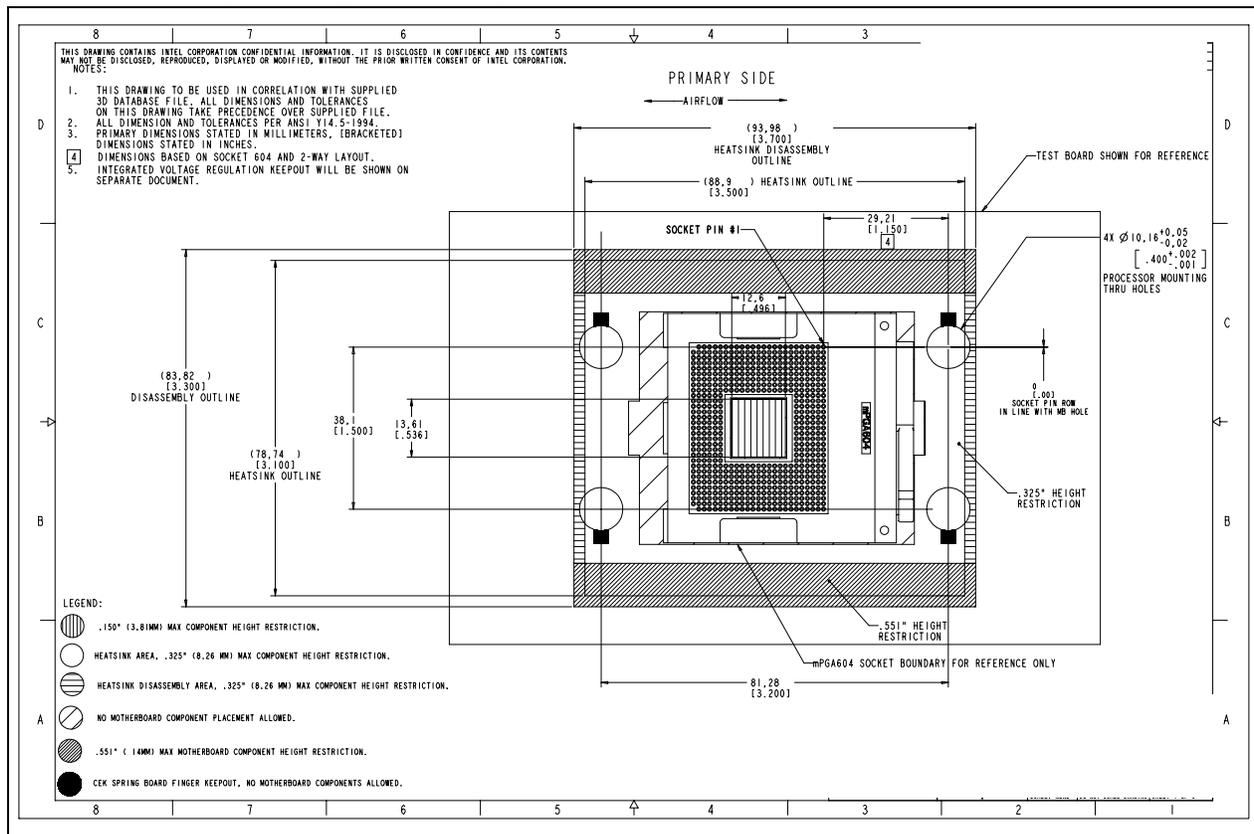


Figure A-10. Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 3 of 5)

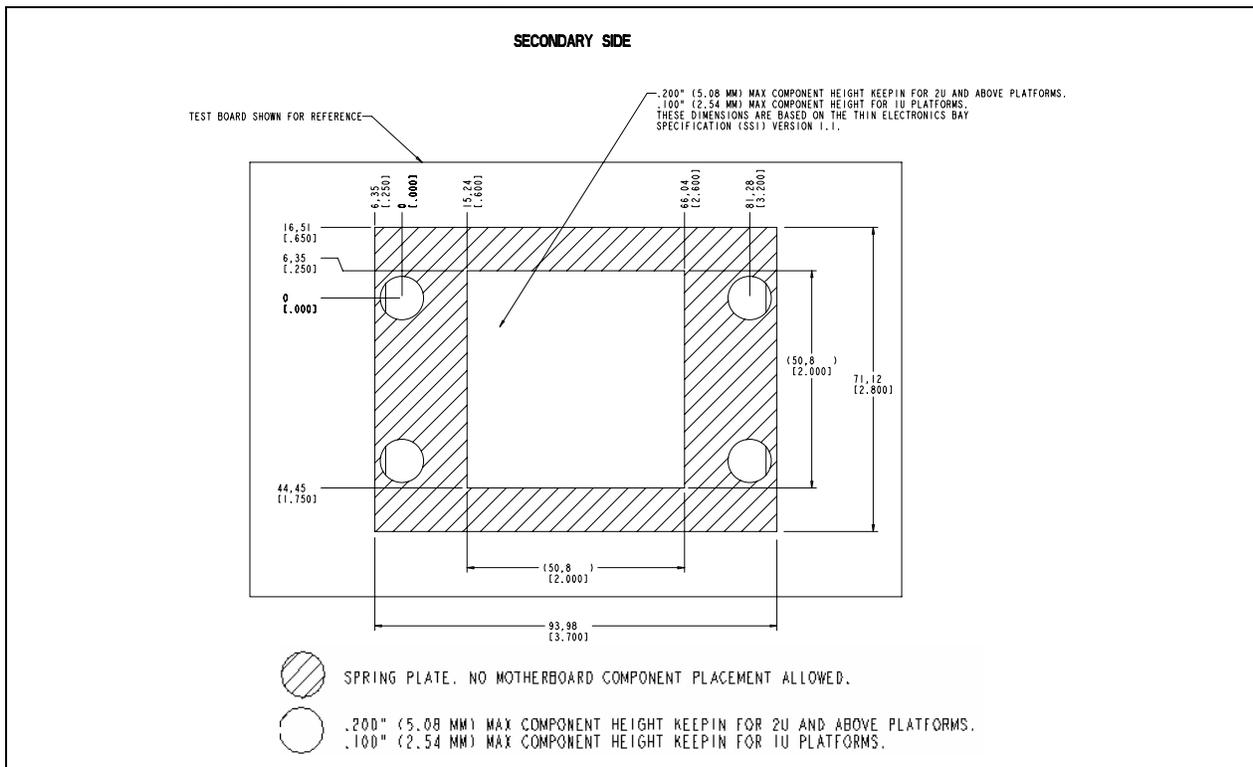


Figure A-11. Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 4 of 5)

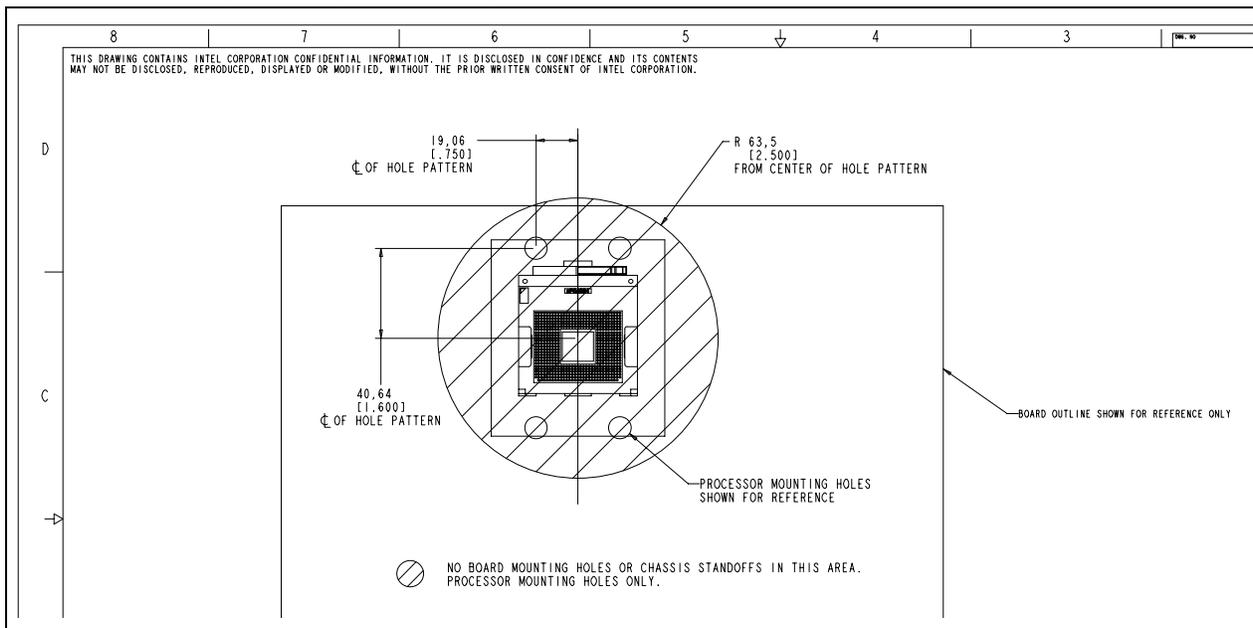
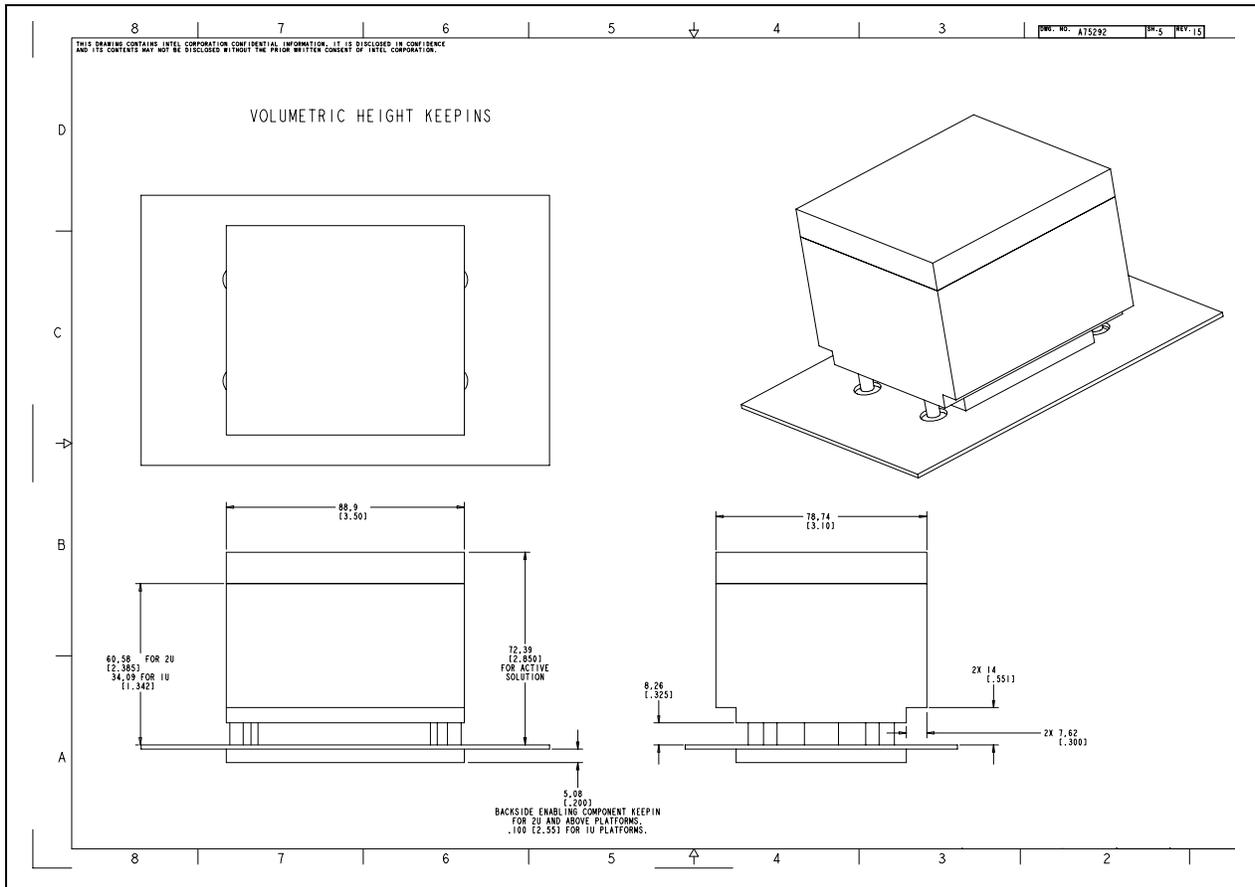


Figure A-12. Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 5 of 5)



§



B Testing Methods

B.1 Case Measurement

Processor cooling performance is determined by measuring the case temperature using a thermocouple and then applying the corresponding correction offset to this measurement. For case temperature measurements, the attach method outlined in this section is recommended for mounting a thermocouple.

Special care is required when measuring case temperature (TC) to ensure an accurate temperature measurement. Thermocouples are often used to measure TC. When measuring the temperature of a surface that is at a different temperature from the surrounding local ambient air, errors could be introduced in the measurements. The measurement errors could be caused by poor thermal contact between the thermocouple junction and the surface of the integrated heat spreader, heat loss by radiation, convection, by conduction through thermocouple leads, or by contact between the thermocouple cement and the heatsink base. To minimize these measurement errors, the approach outlined in the next section is recommended.

B.2 Supporting Test Equipment

To apply the reference thermocouple attach procedure, it is recommended to use the equipment (or equivalent).

Table B-1. Test Equipment

Item	Description	Part Number
Measurement and Output		
Microscope	Olympus* Light microscope or equivalent.	SZ-40
Digital Multi Meter	Digital Multi Meter for resistance measurement.	Not Available
Test Fixture(s)		
Micromanipulator (See Note below)	Micromanipulator set from YOU* Ltd. Or equivalent Mechanical 3D arm with needle (not included) to maintain TC bead location during the attach process.	YOU-3
Miscellaneous Hardware		
Loctite* 498 Adhesive	Super glue w/thermal characteristics.	49850
Adhesive Accelerator	Loctite* 7452 for fast glue curing.	18490
Kapton* Tape	For holding thermocouple in place or equivalent.	Not Available
Thermocouple	Omega*, 36 gauge, "T" Type.	5SRTC-TT-36-72
Calibration and Control		
Ice Point Cell	Omega*, stable 0 C temperature source for calibration and offset.	TRCIII
Hot Point Cell	Omega*, temperature source to control and understand meter slope gain.	CL950-A-110

NOTE: Three axes set consists of (1ea. U-31CF), (1ea. UX-6-6), (1ea. USM6) and (1ea. UPN-1). More information available at: http://www.narishige.co.jp/you_ltd/english/products/set/you-set.htm#3

B.3 Thermal Calibration and Controls

It is recommended that full and routine calibration of temperature measurement equipment be performed before attempting to perform temperature case measurement of processors. Intel recommends checking the meter probe set against known standards. This should be done at 0 °C (using ice bath or other stable temperature source) and at an elevated temperature, around 80 °C (using an appropriate temperature source).

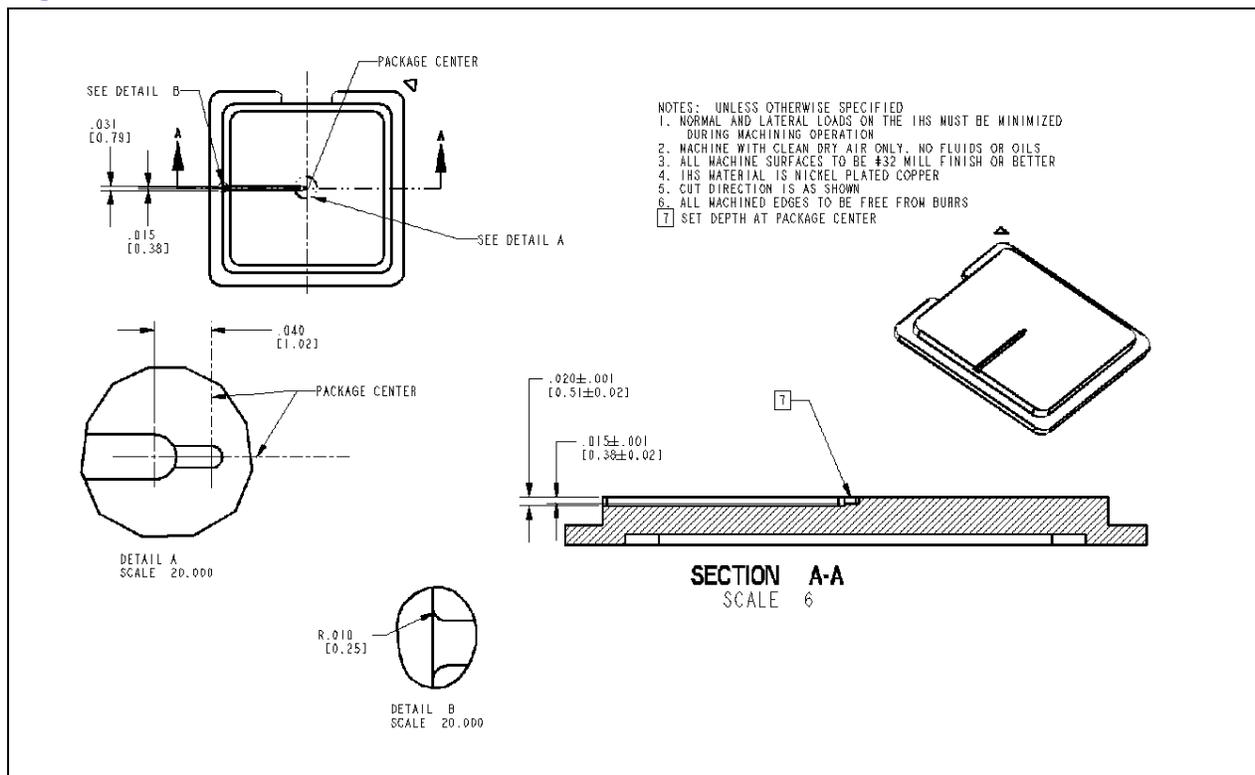
Wire gauge and length also should be considered as some less expensive measurement systems are heavily impacted by impedance. There are numerous resources available throughout the industry to assist with implementation of proper controls for thermal measurements.

1. It is recommended to follow company standard procedures and wear safety items like glasses for cutting the IHS and gloves for chemical handling.
2. Ask your Intel field sales representative if you need assistance to groove and/or install a thermocouple according to the reference process.

B.4 IHS Groove

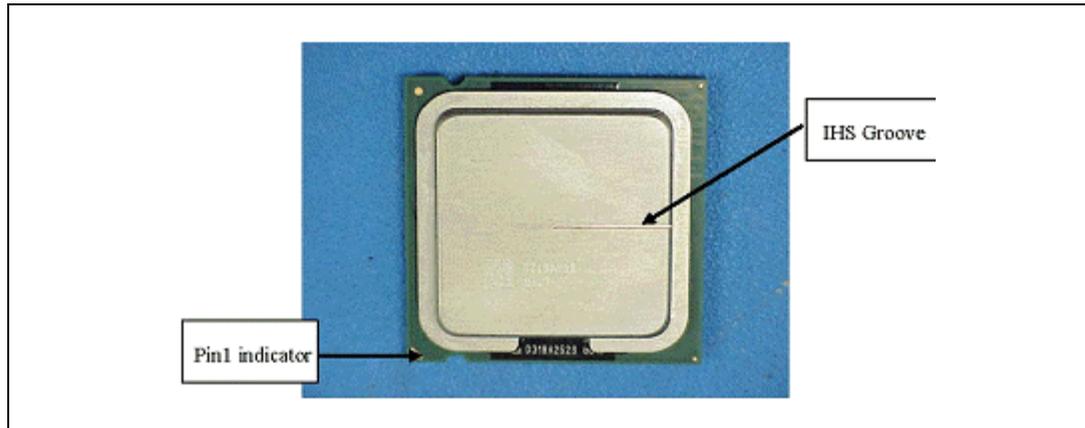
Cut a groove in the package IHS according to the drawing.

Figure B-1. IHS Groove



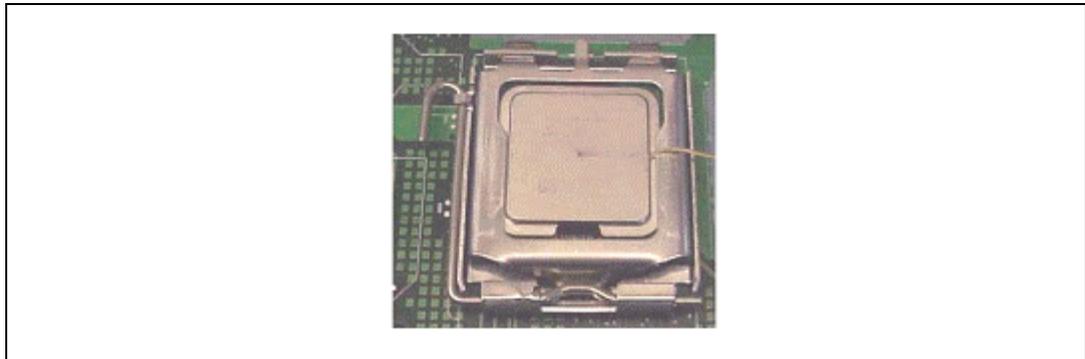
NOTE: In case of conflicts in dimensions, the processor datasheet supersedes this information.

Figure B-2. Groove to Pin Indicator



When the processor is installed in the socket, the groove is perpendicular to the socket load lever, and on the opposite side of the lever.

Figure B-3. IHS Groove

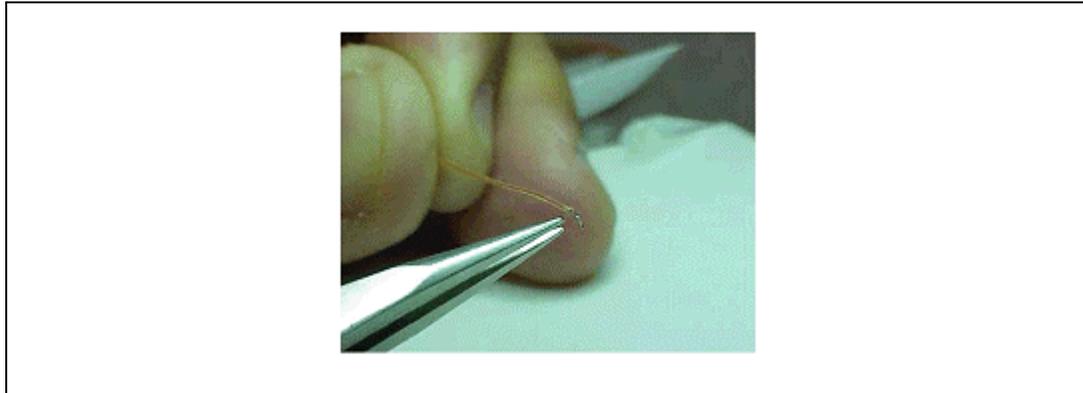


Select a machine shop that is capable of holding drawing specified tolerances. IHS channel geometry is critical for repeatable placement of the thermocouple bead, ensuring precise thermal measurements. The specified dimensions minimize the impact of the groove on the IHS under the socket load. A larger groove may cause the IHS to warp under the socket load such that it does not represent the performance of an ungrooved IHS on production packages.

Note: Inspect parts for compliance to specifications before accepting from machine shop.

B.5 Thermocouple Conditioning and Preparation

1. Use a calibrated thermocouple.
2. Measure the thermocouple resistance by holding both wires on one probe and the tip of thermocouple to the other probe of the DMM (compare to thermocouple resistance specifications).
3. Straighten the wire for about 38 mm [1 inch] from the bead to place it inside the channel.
4. Bend the tip of the thermocouple at approximately 45 degree angle by about 0.8 mm [0.030 inch] from the tip.

Figure B-4. Bending Tip of Thermocouple

B.6 Thermocouple Attachment to the IHS

1. Clean groove with IPA and a lint free cloth removing all residues prior to thermocouple attachment.
2. Place the thermocouple wire inside the groove letting the exposed wire and bead extend about 3.2 mm [0.125 inch] past the end of groove. Secure it with Kapton tape.
3. Lift the wire at the middle of groove with tweezers and bend the front of wire to place the thermocouple in the channel ensuring the tip is in contact with the end of the channel grooved in the IHS.
4. Place the processor under the microscope unit to continue with process. It is also recommended to use a fixture (like processor tray or a plate) to help holding the unit in place for the rest of the attach process.
5. Press the wire down about 6mm [0.125"] from the thermocouple bead using the tweezers. Look in the microscope to perform this task. Place a piece of Kapton tape to hold the wire inside the groove.
6. Using the micromanipulator, install the needle near to the end of groove on top of thermocouple. Using the X, Y, and Z axes on the arm place the tip of needle on top of the thermocouple bead. Press down until the bead is seated at the end of groove on top of the step.
7. Measure resistance from thermocouple end wires (hold both wires to a DMM probe) to the IHS surface. This should be the same value as measured during the thermocouple conditioning.
8. Place a small amount of Loctite 498 adhesive in the groove where the bead is installed. Using a fine point device, spread the adhesive in the groove around the needle, the thermocouple bead and the thermocouple wires already installed in the groove during step 5 above. Be careful not to move the thermocouple bead during this step.
9. Measure the resistance from the thermocouple end wires again using the DMM and to ensure the bead is still properly contacting the IHS.

Figure B-5. Securing Thermocouple Wires with Kapton* Tape

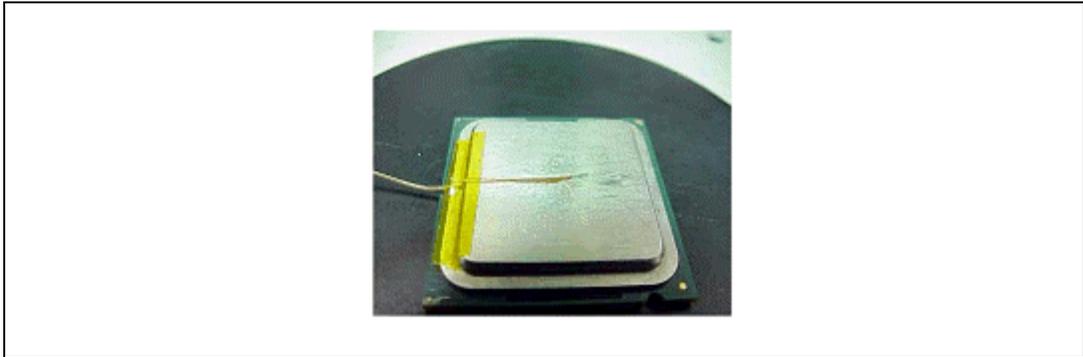


Figure B-6. Thermocouple Bead Placement

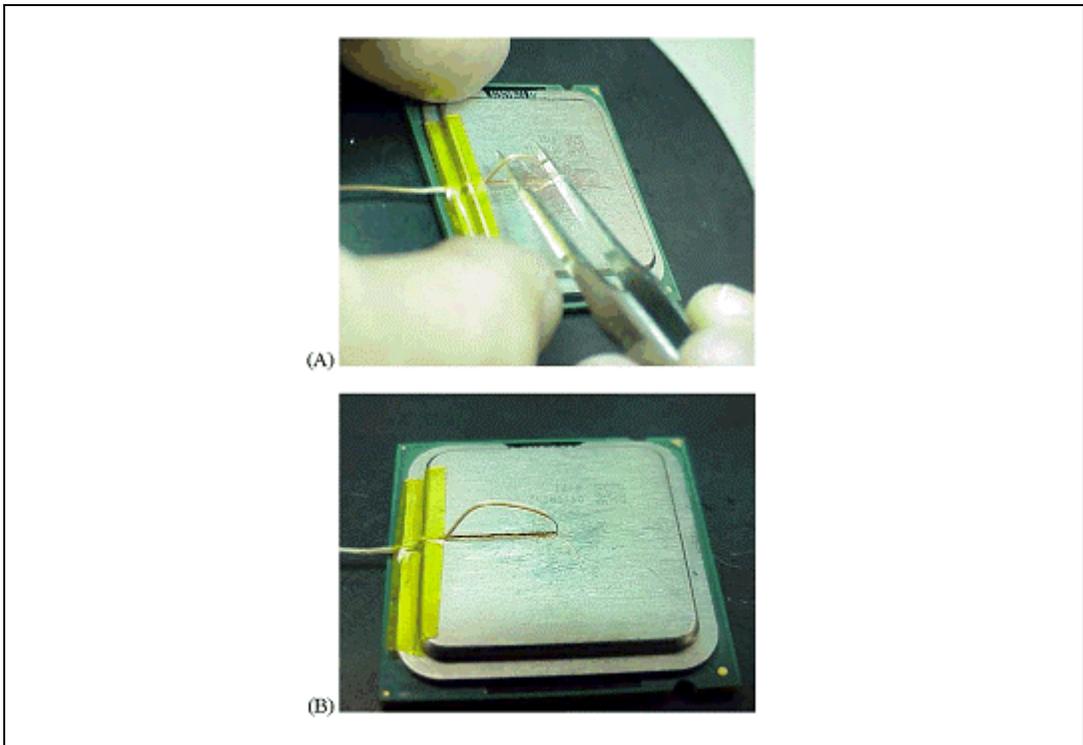


Figure B-7. Thermocouple Placement

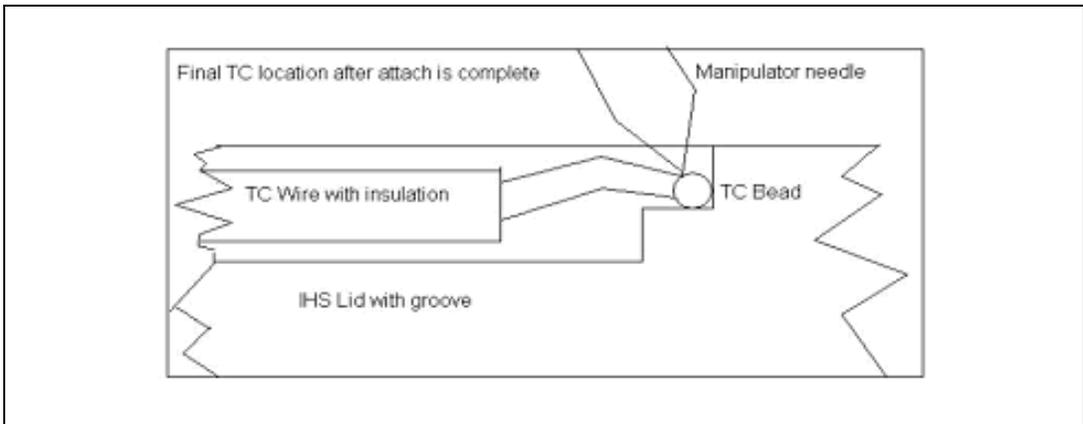


Figure B-8. 3D Micromanipulator to Secure Bead Location



Figure B-9. Measuring Resistance between Thermocouple and IHS

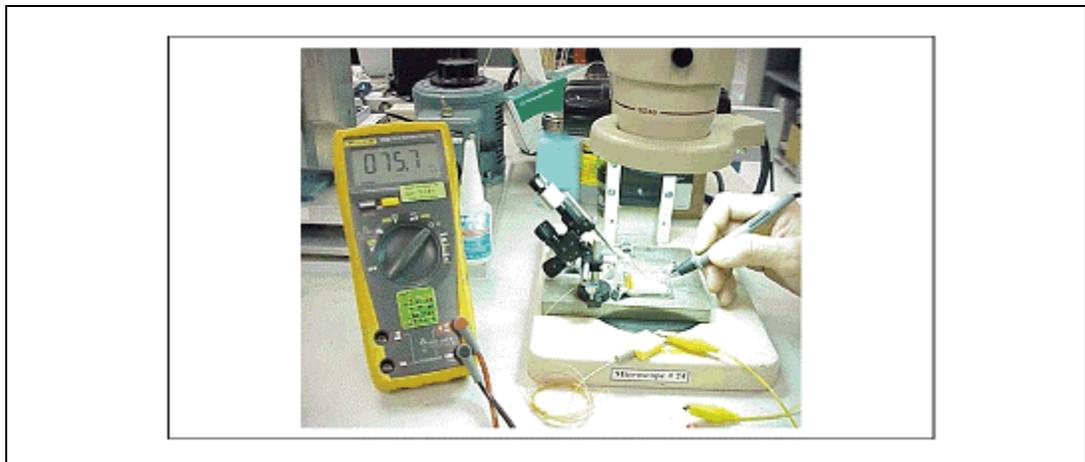
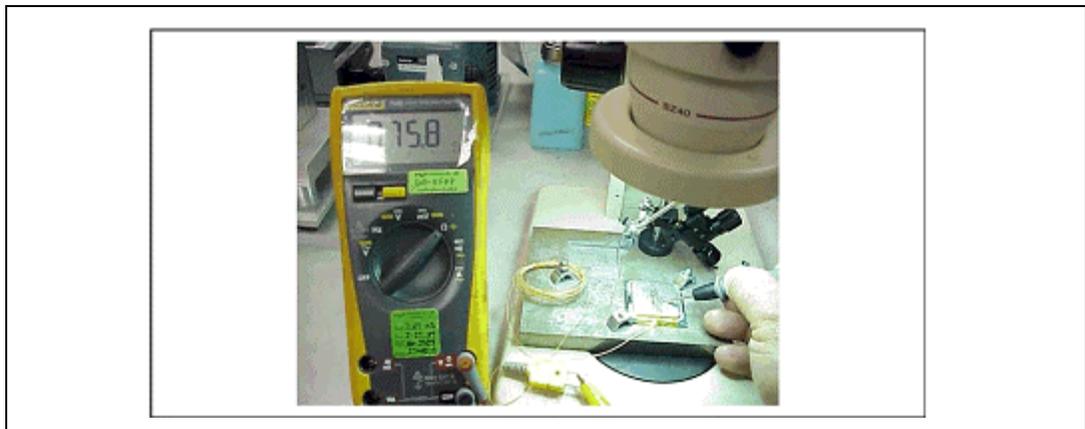


Figure B-10. Applying the Adhesive on the Thermocouple Bead



B.7 Curing Process

1. Let the thermocouple attach set in the open-air for at least 1/2 Hr. It is not recommended to use any curing accelerator like Loctite* Accelerator 7452 for this step, as rapid contraction of the adhesive during curing may weaken bead attach on the IHS.
2. Reconfirm electrical connectivity with DMM before removing the micromanipulator.
3. Remove the 3D Arm needle by holding down the processor unit and lifting the arm.
4. Remove the Kapton* tape, straighten the wire in the groove so it lays flat all the way to the end of the groove.
5. Using a blade to shave excess adhesive above the IHS surface.

Take usual precautions when using open blades.

1. Install new Kapton* tape to hold the thermocouple wire down and fill the rest of groove with adhesive. Make sure the wire and insulation is entirely within the groove and below the IHS surface.
2. Curing time for the rest of the adhesive in the groove can be reduced using Loctite* Accelerator 7452.
3. Repeat step 5 to remove any access adhesive to ensure flat IHS for proper mechanical contact to the heatsink surface.

B.8 Thermocouple Wire Management

Figure B-11. Thermocouple Wire Management in the Groove

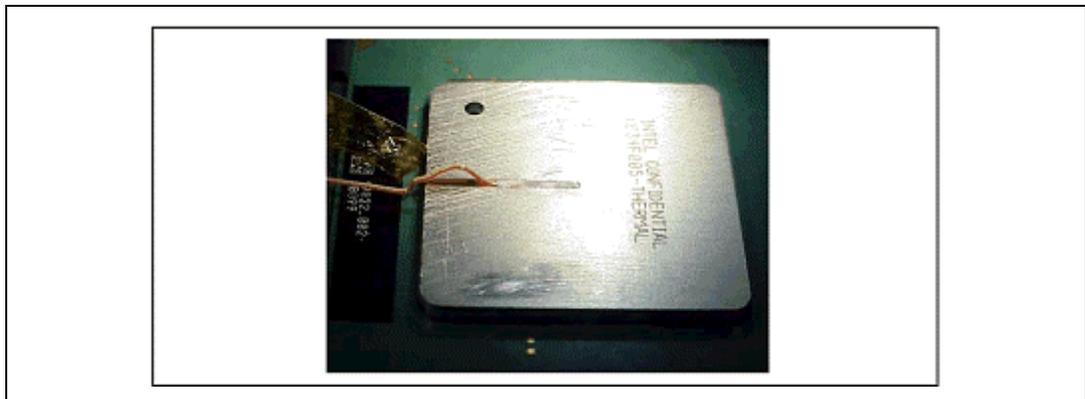


Figure B-12. Removing Excess Adhesive from the IHS

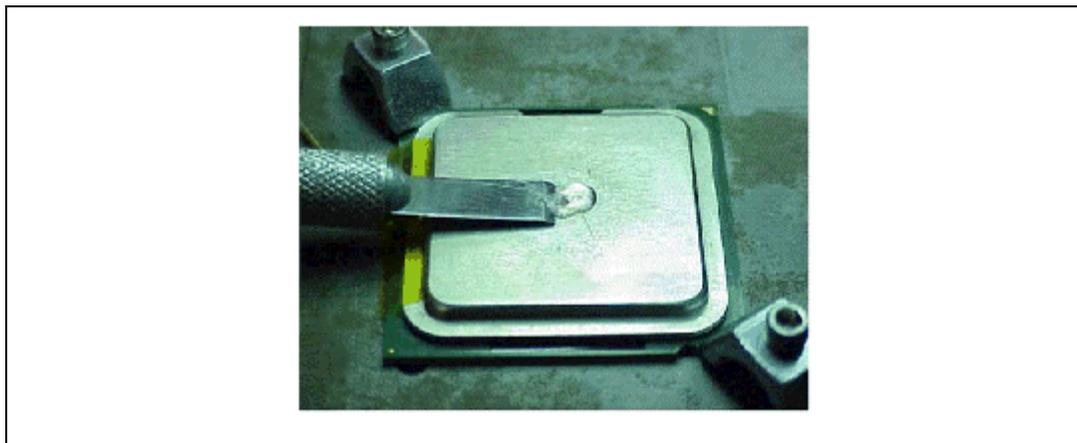
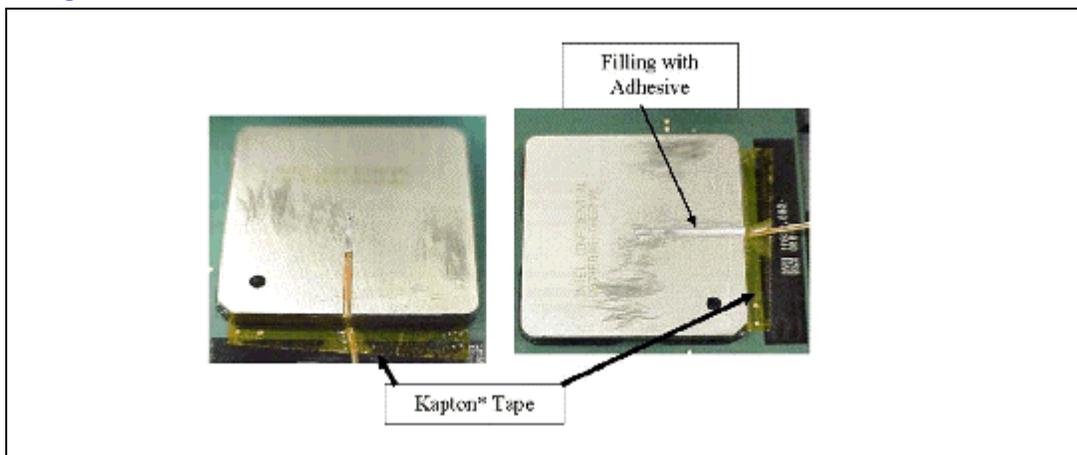


Figure B-13. Filling the Groove with Adhesive

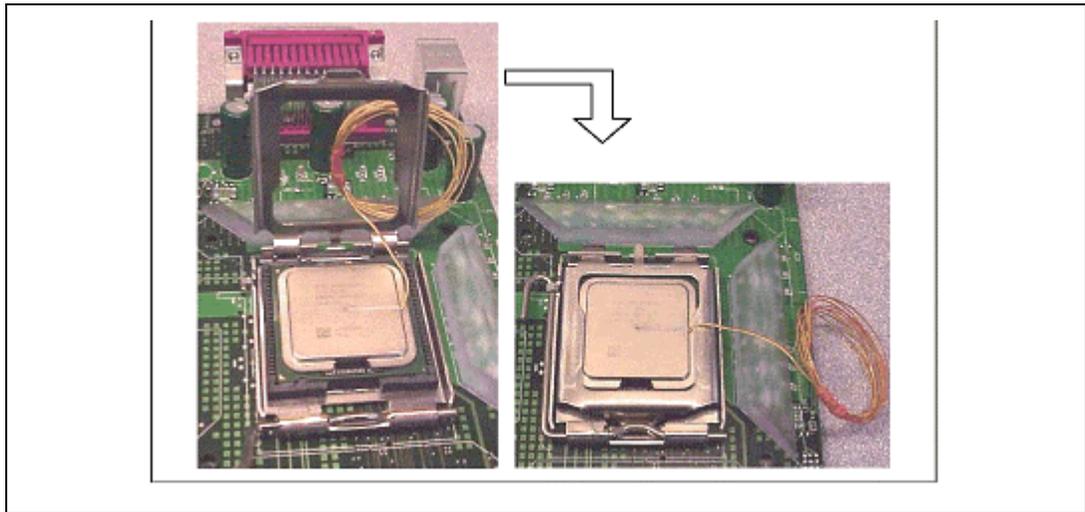


When installing the processor into the socket, make sure that the thermocouple wires exit above the load plate. Pinching the thermocouple wires between the load plate and the IHS will likely damage the wires.

Note: When thermocouple wires are damaged, the resulting reading may likely be wrong. For example, if there are any cuts into the wires insulation where the wires are pinched between the IHS and the load plate, the thermocouple wires can get in contact at this location. In that case, the temperature would be really measured will be measured on the edge of the IHS/socket load plate area. This temperature will likely be much lower than the temperature at the center of the IHS.

Prior to installing the heatsink, make sure that the thermocouple wires remain below the IHS top surface, by running a flat blade on top of the IHS for example.

Figure B-14. Thermocouple Wire Management



B.9 Local Air Thermocouple Placement

For passive heatsinks, two thermocouples will be placed 10 mm upstream of the processor heatsink. The thermocouples will be centered with respect to the height of the heatsink fins and evenly across the width of the heatsink.

For active heatsinks, four thermocouples will be placed on the fan inlet. These thermocouples will be mounted between 5 mm and 10 mm above the fan. The average of these measurements will be used to represent the local inlet temperature to the active heatsink.

Figure B-15. Local Air Thermocouple Placement for Passive Heatsinks

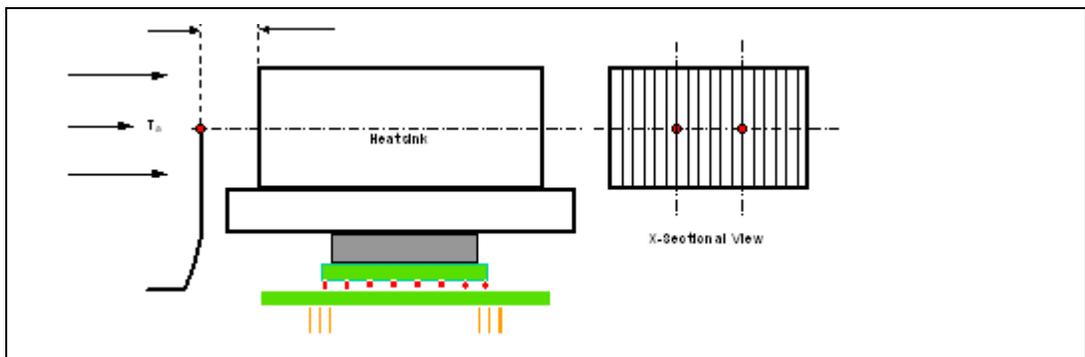


Figure B-16. Local Air Thermocouple Placement for Active Heatsinks (Side View)

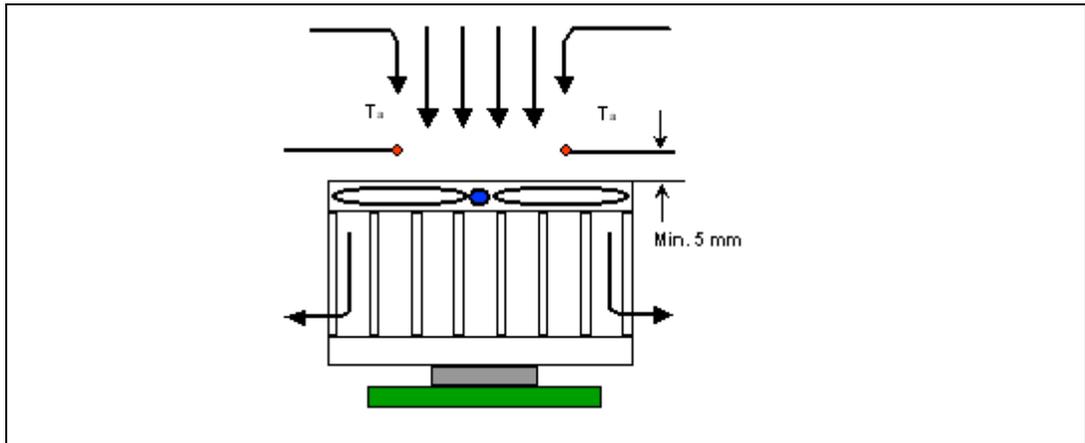
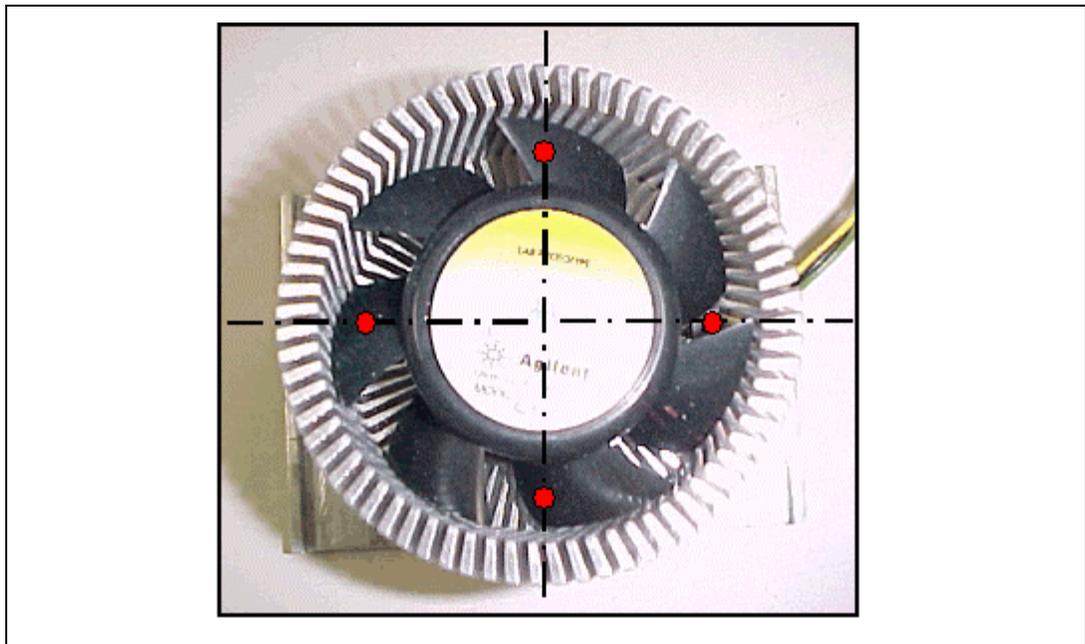


Figure B-17. Local Air Thermocouple Placement for Active Heatsinks (Plan View)





C Safety Requirements

Heatsink and attachment assemblies shall be consistent with the manufacture of units that meet the safety standards:

- UL Recognition-approved for flammability at the system level. All mechanical and thermal enabling components must be a minimum UL94V-2 approved.
- CSA Certification. All mechanical and thermal enabling components must have CSA certification.
- Heatsink fins must meet the test requirements of UL1439 for sharp edges.

§



D Quality and Reliability Requirements

D.1 Intel Verification Criteria for the Reference Designs

D.1.1 Reference Heatsink Thermal Verification

The Intel reference heatsinks will be verified within specific boundary conditions based on the methodology described in Appendix B.

The test results, for a number of samples, are reported in terms of a worst-case mean + 3σ value for thermal characterization parameter using real processors.

D.1.2 Environmental Reliability Testing

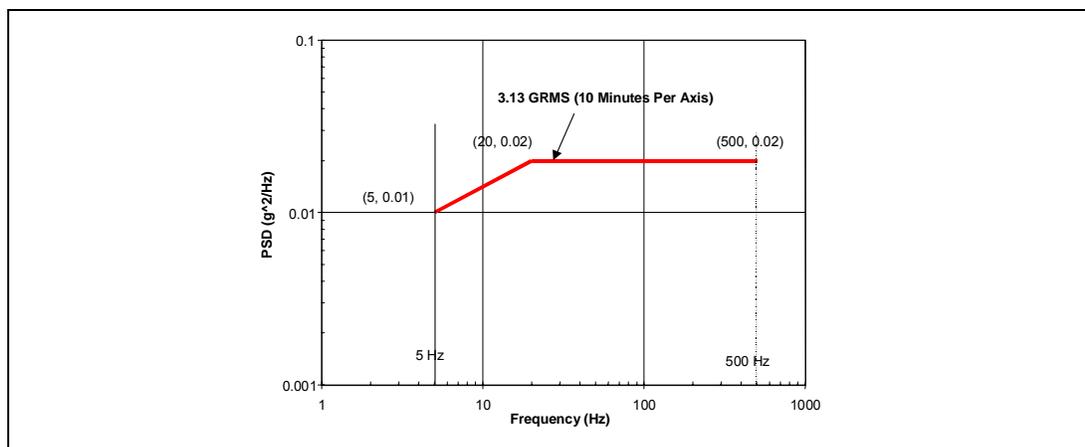
D.1.2.1 Structural Reliability Testing

Structural reliability tests consist of unpackaged, board-level vibration and shock tests of a given thermal solution in assembled state, as well as long-term reliability testing (temperature cycling, bake test). The thermal solution should be capable of sustaining thermal performance after these tests are conducted; however, the conditions of the tests outlined here may differ from the customers' system requirements.

D.1.2.2 Random Vibration Test Procedure

- Duration: 10 min/axis, 3 axes
- Frequency Range: 5 Hz to 500 Hz
- Power Spectral Density (PSD) Profile: 3.13 G RMS (refer to Figure D-1).

Figure D-1. Random Vibration PSD

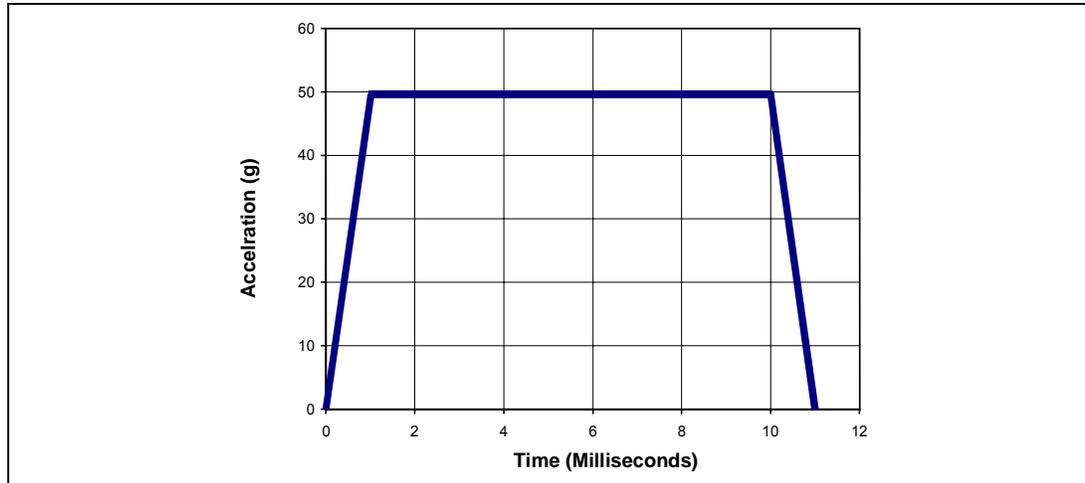


D.1.2.3 Shock Test Procedure

Recommended performance requirement for a baseboard:

- Quantity: 3 drops for + and – directions in each of 3 perpendicular axes (i.e. total 18 drops).
- Profile: 50 G trapezoidal waveform, 11 ms duration, 4.32 m/sec minimum velocity change.
- Setup: Mount sample board on test fixture.

Figure D-2. Shock Acceleration Curve



D.1.2.4 Recommended Test Sequence

Each test sequence should start with components (i.e. baseboard, heatsink assembly, etc.) that have not been previously submitted to any reliability testing.

The test sequence should always start with a visual inspection after assembly, and BIOS/Processor/memory test. The stress test should be then followed by a visual inspection and then BIOS/Processor/memory test.

D.1.2.5 Post-Test Pass Criteria

The post-test pass criteria are:

1. No significant physical damage to the heatsink and retention hardware.
2. Heatsink remains seated and its bottom remains mated flatly against the IHS surface. No visible gap between the heatsink base and processor IHS. No visible tilt of the heatsink with respect to the retention hardware.
3. No signs of physical damage on baseboard surface due to impact of heatsink.
4. No visible physical damage to the processor package.
5. Successful BIOS/Processor/memory test of post-test samples.
6. Thermal compliance testing to demonstrate that the case temperature specification can be met.

D.1.2.6 Recommended BIOS/Processor/Memory Test Procedures

This test is to ensure proper operation of the product before and after environmental stresses, with the thermal mechanical enabling components assembled. The test shall be conducted on a fully operational baseboard that has not been exposed to any battery of tests prior to the test being considered.

Testing setup should include the following components, properly assembled and/or connected:

- Appropriate system baseboard.
- Processor and memory.
- All enabling components, including socket and thermal solution parts.

The pass criterion is that the system under test shall successfully complete the checking of BIOS, basic processor functions and memory, without any errors. *Intel PC Diags* is an example of software that can be utilized for this test.

D.1.3 Material and Recycling Requirements

Material shall be resistant to fungal growth. Examples of non-resistant materials include cellulose materials, animal and vegetable based adhesives, grease, oils, and many hydrocarbons. Synthetic materials such as PVC formulations, certain polyurethane compositions (e.g. polyester and some polyethers), plastics which contain organic fillers of laminating materials, paints, and varnishes also are susceptible to fungal growth. If materials are not fungal growth resistant, then MIL-STD-810E, Method 508.4 must be performed to determine material performance.

Material used shall not have deformation or degradation in a temperature life test.

Any plastic component exceeding 25 grams must be recyclable per the European Blue Angel recycling standards.

§



E Enabled Suppliers Information

Component	Description	Development Suppliers	Supplier Contact Info
Cooling Solution Heatsink	Copper Fin, Copper Base	Fujikura* (stacked fin) CNDA 36187 Furukawa (crimped fin) CNDA 65755	Mechatronics* Steve Carlson 800-453-4569 x205 steve@mechatronics.com Furukawa America Katsu Mizushima (408) 232-9306 katsumizushima@mindspring.com
Thermal Interface Material	Grease	Shin-Etsu* G751 CNDA 75610	Donna Hartigan (480) 893-8898
Cooling Solution Spring	Stainless Steel 301, Kapton* Tape on Spring Fingers	ITW Fastex* CNDA 78538 AVC* CNDA 2085011 Foxconn* CNDA 11251	Ron Schmidt (847) 299-2222 rschmidt@itwfastex.com Felicia Lee 886-2-22996390 x144 felicia@avc.com.tw

§



F Processor Thermal Management Logic and Thermal Monitor Features

F.1 Thermal Management Logic and Thermal Monitor Feature

F.1.1 Processor Power Dissipation

An increase in processor operating frequency not only increases system performance, but also increases the processor power dissipation. The relationship between frequency and power is generalized in the following equation: $P = CV^2F$ (where P = power, C = capacitance, V = voltage, F = frequency). From this equation, it is evident that power increases linearly with frequency and with the square of voltage. In the absence of power saving technologies, ever increasing frequencies will result in processors with power dissipations in the hundreds of watts. Fortunately, there are numerous ways to reduce the power consumption of a processor, and Intel is aggressively pursuing low power design techniques. For example, decreasing the operating voltage, reducing unnecessary transistor activity, and using more power efficient circuits can significantly reduce processor power consumption.

An on-die thermal management feature called Thermal Monitor is available on the 64-bit Intel Xeon processor MP with 1 MB L2 cache. It provides a thermal management approach to support the continued increases in processor frequency and performance. By using a highly accurate on-die temperature sensing circuit and a fast acting temperature control circuit, the processor can rapidly initiate thermal management control. The Thermal Monitor can reduce cooling solution cost, by allowing designs to target TDP instead of maximum processor power.

F.1.2 Thermal Monitor Implementation

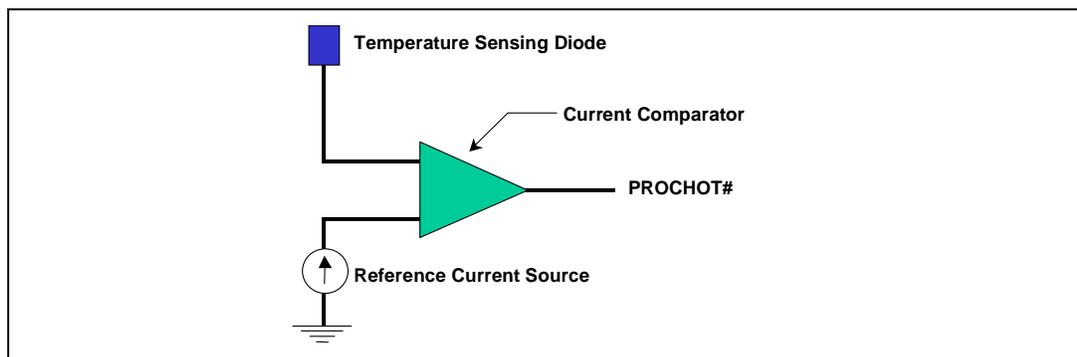
On the 64-bit Intel Xeon processor MP with 1 MB L2 cache, the Thermal Monitor is integrated into the processor silicon. The Thermal Monitor includes:

- An on-die temperature sensing circuit.
- An external output signal (PROCHOT#) that indicates the processor has reached its maximum operating temperature.
- An external input signal (FORCEPR#) that allows the platform to force a power reduction by the processor by activating the TCC.
- A TCC that can reduce processor temperature by rapidly reducing power consumption when the on-die temperature sensor indicates that it has reached the maximum operating point.
- Registers to determine the processor thermal status.

The processor temperature is determined through an analog thermal sensor circuit comprised of a temperature sensing diode, a factory calibrated reference current source, and a current comparator (see Figure F-1). A voltage applied across the diode induces a current flow that varies with

temperature. By comparing this current with the reference current, the processor temperature can be determined. The reference current source corresponds to the diode current when at the maximum permissible processor operating temperature. Processors are calibrated during manufacturing on a small sample set. Once configured, the processor temperature at which the PROCHOT# signal is asserted (trip point) is not re-configurable.

Figure F-1. Thermal Sensor Circuit

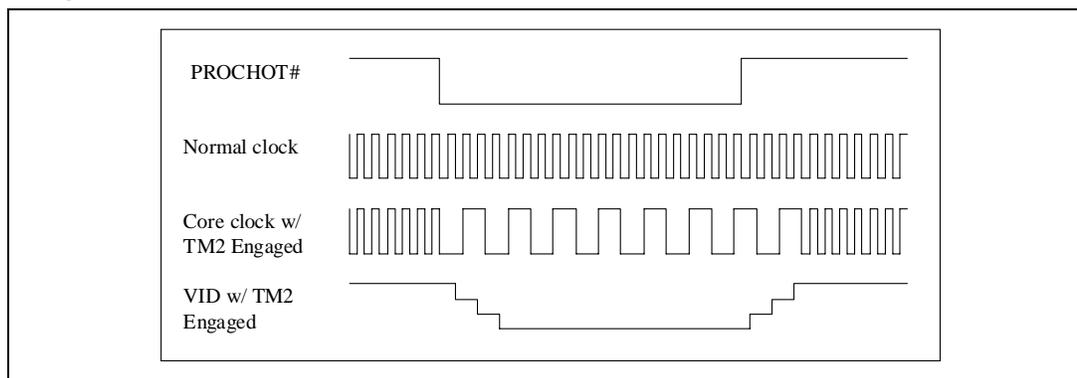


The PROCHOT# signal is available internally to the processor as well as externally. External indication of the processor temperature status is provided through the bus signal PROCHOT#. When the processor temperature reaches the trip point, PROCHOT# is asserted. When the processor temperature is below the trip point, PROCHOT# is de-asserted. Assertion of the PROCHOT# signal is independent of any register settings within the processor. It is asserted any time the processor die temperature reaches the trip point. The point where the TCC activates is set to the same temperature at which the processor is tested and at which PROCHOT# asserts.

The TCC portion of the Thermal Monitor must be enabled for the processor to operate within specifications. The Thermal Monitor's TCC, when active, lowers the processor temperature by reducing the power consumed by the processor. This is done by changing the duty cycle of the internal processor clocks, resulting in a lower effective frequency. When active, the TCC turns the processor clocks off and then back on with a predetermined duty cycle. The duty cycle is processor specific, and is fixed for a particular processor. The maximum time period the clocks are disabled is ~3 μ s, and is frequency dependent. Higher frequency processors will disable the internal clocks for a shorter time period. Figure F-2 illustrates the relationship between the internal processor clocks and PROCHOT#.

Performance counter registers, status bits in model specific registers (MSRs), and the PROCHOT# output pin are available to monitor and control the Thermal Monitor behavior.

Figure F-2. Concept for Clocks under Thermal Monitor Control



F.1.3 Operation and Configuration

To maintain compatibility with previous generations of processors, which have no integrated thermal logic, the TCC portion of Thermal Monitor is disabled by default. During the boot process, the BIOS must enable the TCC; or a software driver may do this after the operating system has booted. **Thermal Monitor or Thermal Monitor 2 feature must be enabled for the processor to remain within specification.**

The TCC feature can be configured and monitored in a number of ways. OEMs are expected to enable the TCC while using various registers and outputs to monitor the processor thermal status. The TCC is enabled by the BIOS setting a bit in an MSR (model specific register). Enabling the TCC allows the processor to maintain a safe operating temperature without the need for special software drivers or interrupt handling routines. When the TCC has been enabled, processor power consumption will be reduced within a few hundred clock cycles after the thermal sensor detects a high temperature, i.e. PROCHOT# assertion. The TCC and PROCHOT# transition to inactive once the temperature has been reduced below the thermal trip point, although a small time-based hysteresis has been included to prevent multiple PROCHOT# transitions around the trip point. External hardware can monitor PROCHOT# and generate an interrupt whenever there is a transition from active-to-inactive or inactive-to-active. PROCHOT# can also be configured to generate an internal interrupt which would initiate an OEM supplied interrupt service routine. Regardless of the configuration selected, PROCHOT# will consistently indicate the thermal status of the processor.

The TCC can also be activated manually using an “on-demand” mode.

F.1.4 Thermal Monitor 2

The 64-bit Intel Xeon processor MP with 1 MB L2 cache also supports an enhanced TCC that works in conjunction with the existing Thermal Monitor logic. This capability is known as Thermal Monitor 2. This improved TCC provides a more efficient means for limiting the processor temperature by reducing the power consumption within the processor.

When Thermal Monitor 2 is enabled, and a high temperature situation is detected, the enhanced TCC will be activated. The enhanced TCC causes the processor to adjust its operating frequency (bus-to-core multiplier) and input voltage identification (VID) value. This combination of reduced frequency and the lowering of VID results in a reduction in processor power consumption.

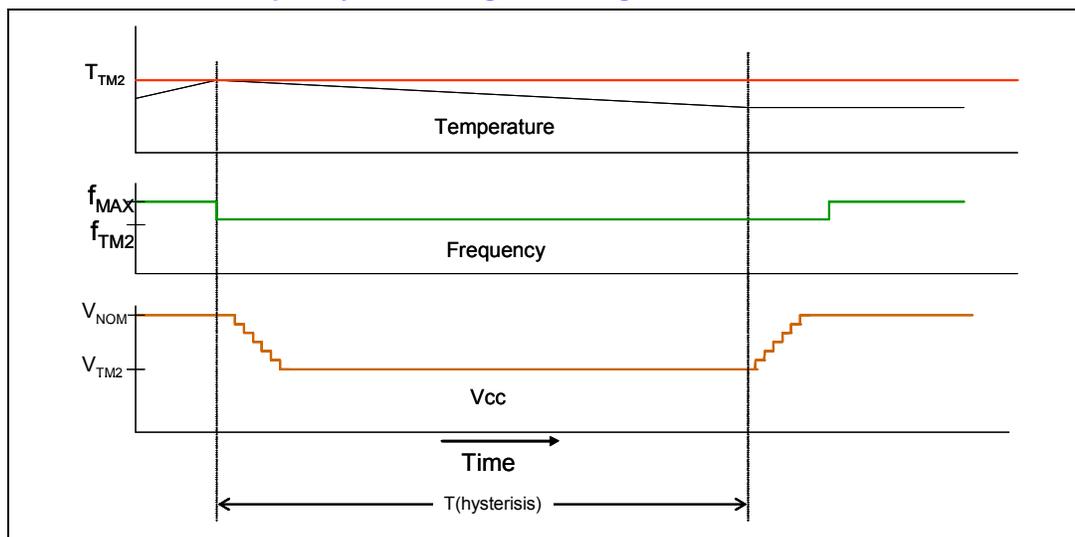
A processor enabled for Thermal Monitor 2 includes two operating points, each consisting of a specific operating frequency and voltage. The first operating point represents the normal operating condition for the processor. The second operating point consists of both a lower operating frequency and voltage.

When the TCC is activated, the processor automatically transitions to the new frequency. This transition occurs very rapidly (on the order of 5 microseconds). During the frequency transition, the processor is unable to service any bus requests, and consequently, all bus traffic is blocked during the frequency transition. Edge-triggered interrupts will be latched and kept pending until the processor resumes operation at the new frequency.

Once the new operating frequency is engaged, the processor will transition to the new core operating voltage by issuing a new VID code to the voltage regulator. The voltage regulator must support dynamic VID changes in order to support Thermal Monitor 2. During the voltage change, it will be necessary to transition through multiple VID codes to reach the target operating voltage. Each step will be one VID table entry (i.e. 12.5 mV steps). The processor continues to execute instructions during the voltage transition. Operation at the lower voltage reduces both the dynamic and leakage power consumption of the processor. Once the processor has sufficiently cooled, and the time based

hysteresis period has expired, the operating frequency and voltage transition back to the normal system operating point. Transition of the VID code will occur first, in order to insure proper operation once the processor reaches its normal operating frequency. Refer to Figure F-3 for an illustration of this ordering.

Figure F-3. Thermal Monitor 2 Frequency and Voltage Ordering



F.1.5 System Considerations

The Thermal Monitor feature may be used in a variety of ways, depending upon the system design requirements and capabilities.

Note: Intel requires the TCC to be enabled for all 64-bit Intel Xeon processor MP with 1 MB L2 cache -based systems. At a minimum, the TCC provides an added level of protection against processor thermal solution failure.

A system designed to meet the TDP and T_{CASE} targets published in the processor datasheet greatly reduces the probability of real applications causing the TCC to activate under normal operating conditions. Systems that do not meet these specifications could be subject to more frequent activation of the TCC depending upon ambient air temperature and application power profile. Moreover, if a system is significantly under designed, there is a risk that the Thermal Monitor feature will not be capable of maintaining a safe operating temperature and the processor could shutdown and signal THERMTRIP#.

For information regarding THERMTRIP#, refer to Appendix F.1.7.2 and to the processor datasheet

F.1.6 Operating System and Application Software Considerations

The Thermal Monitor feature and its TCC work seamlessly with ACPI compliant operating systems and those utilizing hardware based timing routines. The Thermal Monitor feature is transparent to application software since the processor bus snooping, ACPI timer, and interrupts are active at all times.

Activation of the TCC during a non-ACPI aware operating system boot process may result in incorrect calibration of operating system software timing loops. This is also the case with operating

systems that utilize execution based timing routines. The BIOS must disable the TCC prior to boot and then the operating system or BIOS must enable the TCC after the operating system boot process completes.

Intel has worked with the major operating system vendors to ensure support for non-execution based operating system calibration loops and ACPI support for the Thermal Monitor feature.

F.1.7 Legacy Thermal Management Capabilities

In addition to Thermal Monitor, the 64-bit Intel Xeon processor MP with 1 MB L2 cache supports the same thermal management features originally introduced with the Intel® Pentium® III Xeon™ processor. These features include the on-die thermal diode and THERMTRIP# signal for indicating catastrophic thermal failure.

F.1.7.1 On-Die Thermal Diode

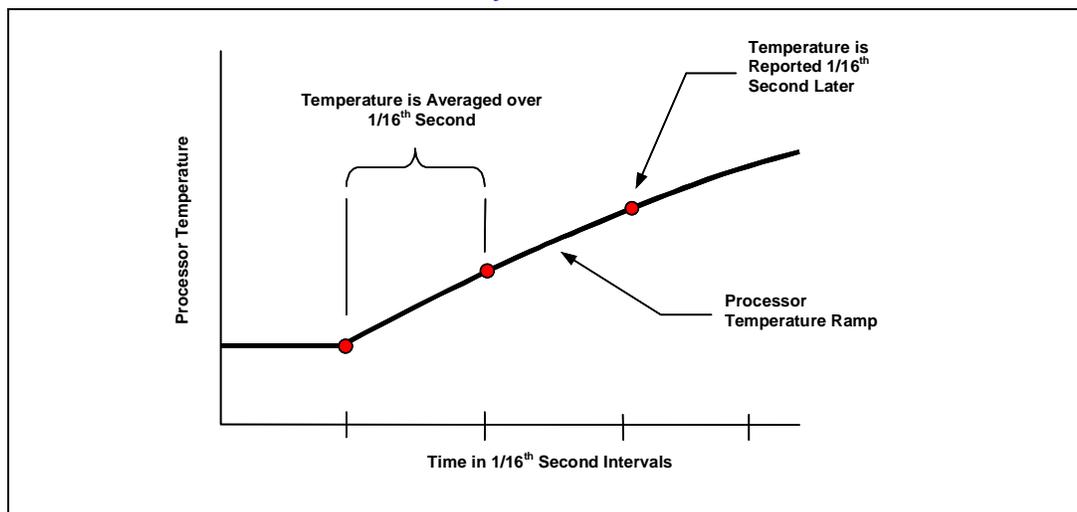
There are two independent thermal diodes in the 64-bit Intel Xeon processor MP with 1 MB L2 cache. One is the on-die thermal diode and the other is in the temperature sensor used for the Thermal Monitor and for THERMTRIP#. The Thermal Monitor's temperature sensor and the on-die thermal diode are independent and isolated devices with no direct correlation to one another. Circuit constraints and performance requirements prevent the Thermal Monitor's temperature sensor and the on-die thermal diode from being located at the same place on the silicon. The temperature distribution across the die may result in significant temperature differences between the on-die thermal diode and the Thermal Monitor's temperature sensor. This temperature variability across the die is highly dependent on the application being run. As a result, it is not possible to predict the activation of the TCC by monitoring the on-die thermal diode.

System integrators that plan on using the thermal diode for system or component level fan control need to be aware of the potential for rapid changes in processor power consumption as the executing workload changes. Variable performance thermal solutions that fail to react quickly to changing workloads may experience TCC activation or worst yet, result in automatic shutdown via THERMTRIP# (refer to Appendix F.1.7.2 for more information on THERMTRIP). One example of this situation is as follows: A fan control scheme slows the fans such that the processor is operating very near the thermal trip point while executing a relatively low power workload. The start of a higher power application creates a sudden increase in power consumption and elevates the temperature of the processor above the trip point, causing the TCC to activate. The power reduction resulting from TCC activation slows the rate of temperature increase, but is not sufficient to clamp the temperature, due to inadequate thermal solution performance at reduced fan speed. As a result, the temperature continues to slowly increase. The fan is then sped up to compensate for the change in processor workload but reacts too slowly to prevent the processor from shutting down due to THERMTRIP# activation.

High temperature change rates on-die can also limit the ability to accurately measure the on-die thermal diode temperature. As a result, the on-die thermal diode should not be relied upon to warn of processor cooling system failure or predict the onset of the TCC. An illustration of this is as follows. Many thermal diode sensors report temperatures a maximum of 8 times per second. Within the 1/8th (0.125 sec.) second time period, the temperature is averaged over 1/16th of a second. In a scenario where the silicon temperature ramps at 50°C/sec, or approximately 6°C/0.125 sec, the processor will be ~4.5°C above the temperature reported by the thermal sensor. Change in diode temperature averaged over 1/16th seconds = ~1.5°C; temperature reported 1/16th second later at 1/8th second when the actual processor temperature would be 6°C higher (see Figure F-4).

The on-die thermal diode can be used with an external device (thermal diode sensor) to monitor long-term temperature trends. By averaging this data information over long time periods (hours/days vs. min/sec), it may be possible to derive a trend of the processor temperature. Analysis of this information could be useful in detecting changes in the system environment that may require attention. Design characteristics and usage models of the thermal diode sensors are described in datasheets available from the thermal diode sensor manufacturers.

Figure F-4. On-Die Thermal Diode Sensor Time Delay



F.1.7.2 THERMTRIP# Signal Pin

In the event of a catastrophic cooling failure, the processor will automatically shut down when the silicon temperature has reached its operating limit. At this point the system bus signal THERMTRIP# signal goes active and power must be removed from the processor. THERMTRIP# stays active until RESET# has been initiated. THERMTRIP# activation is independent of processor activity and does not generate any bus cycles.

F.1.7.3 FORCEPR# Signal Pin

The 64-bit Intel Xeon processor MP with 1 MB L2 cache provides a means for system hardware to force activation of the TCC. One possible usage model would be to use this capability to protect the voltage regulator from overheating in order to avoid a catastrophic shutdown. Refer to the appropriate platform design guidelines and voltage regulator design guidelines for implementation details. The use of the FORCEPR# signal pin requires that BIOS code enable the signal's recognition via an MSR.

F.1.8 Cooling System Failure Warning

If desired, the system may be designed to cool the maximum processor power. In this situation, it may be useful to use the PROCHOT# signal as an indication of cooling system failure. Messages could be sent to the system administrator to warn of the cooling failure, while the TCC would allow the system to continue functioning or allow a graceful system shutdown. If no thermal management action is taken, the silicon temperature may exceed the operating limits, causing THERMTRIP# to activate and shut down the processor. Regardless of the system design requirements or thermal solution ability, the Thermal Monitor feature must still be enabled to ensure proper processor operation.

§

