Intel® Celeron® Processor 440 for Embedded Applications

Thermal Design Guide

October 2007

Order Number: 317998-001US

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.

UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

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 information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on 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 processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See http://www.intel.com/products/processor_number for details.

BunnyPeople, Intel Celeron processor 440Celeron, Celeron Inside, Centrino, Centrino logo, Core Inside, FlashFile, i960, InstantIP, Intel, Intel logo, Intel386, Intel486, Intel740, IntelDX2, IntelDX4, IntelSX2, Intel Core, Intel Inside, Intel Inside logo, Intel. Leap ahead., Intel. Leap ahead. logo, Intel NetBurst, Intel NetMerge, Intel NetStructure, Intel SingleDriver, Intel SpeedStep, Intel StrataFlash, Intel Viiv, Intel vPro, Intel XScale, Itanium, Itanium Inside, MCS, MMX, Oplus, OverDrive, PDCharm, Pentium, Pentium Inside, skoool, Sound Mark, The Journey Inside, VTune, Xeon, and Xeon Inside are trademarks of Intel Corporation in the U.S. and other countries.

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

Copyright © 2007, Intel Corporation. All rights reserved.

Contents

Figures

Tables

Revision History

Date	Revision	Description
October 2007	001	Initial release

1.0 Introduction

1.1 Document Goals and Scope

1.1.1 Importance of thermal management

The objective of thermal management is to ensure that the temperatures of all components in a system are maintained within their functional temperature range. Within this temperature range, a component is expected to meet its specified performance. Operation outside the functional temperature range can degrade system performance, cause logic errors or cause component and/or system damage. Temperatures exceeding the maximum operating limit of a component may result in irreversible changes in the operating characteristics of this component.

In a system environment, the processor temperature is a function of both system and component thermal characteristics. The system level thermal constraints consist of the local ambient air temperature and airflow over the processor as well as the physical constraints at and above the processor. The processor temperature depends in particular on the component power dissipation, the processor package thermal characteristics, and the processor thermal solution.

All of these parameters are affected by the continued push of technology to increase processor performance levels and packaging density (more transistors). As operating frequencies increase and packaging size decreases, the power density increases while the thermal solution space and airflow typically become more constrained or remains the same within the system. The result is an increased importance on system design to ensure that thermal design requirements are met for each component, including the processor, in the system.

1.1.2 Document Goals

Depending on the type of system and the chassis characteristics, new system and component designs may be required to provide adequate cooling for the processor. The goal of this document is to provide an understanding of these thermal characteristics and discuss guidelines for meeting the thermal requirements imposed on single processor systems using the Intel Celeron processor 440.

The concepts given in this document are applicable 1U, PICMG1.3 and Embedded ATX form factor. For ATX and BTX form factor thermal solution, please refer to *Intel® Celeron® Processor 440 Thermal Mechanical Design Guide* as listed in the reference document list.

1.1.3 Document Scope

In this document when a reference is made to “the processor” it is intended that this includes the Intel® Celeron® Processor 440 for Embedded Applications unless it is otherwise specified. If needed for clarity, the specific processor will be listed.

This design guide supports the Intel Celeron processor 440.

In this document, when a reference is made to the “EMTS”, the reader should refer to the *Intel® Celeron® Processor 440 EMTS*. If needed for clarity, the specific processor EMTS will be referenced.

Chapter 2.0 discusses package thermal mechanical requirements to design a thermal solution for the Processor in the context of embedded applications.

Chapter 3.0 discusses the thermal solution considerations and metrology recommendations to validate a processor thermal solution.

Chapter 4.0 addresses the benefits of the processor's integrated thermal management logic for thermal design.

Chapter 5.0 gives information on the Intel reference thermal solution for the processor.

Chapter 6.0 discusses the implementation of Intel® Quiet System Technology (Intel® QST).

The physical dimensions and thermal specifications of the processor that are used in this document are for illustration only. Refer to the EMTS for the product dimensions, thermal power dissipation and maximum case temperature. In case of conflict, the data in the EMTS supersedes any data in this document.

1.2 References

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

Document	Source/Reference Number
Intel® Celeron® Processor 400 Datasheet	316963
Intel® Celeron® Processor 400 Processor Thermal Mechanical Design Guide	316965
Intel® Core™2 Duo Desktop Processor, Intel® Pentium® Dual-Core Processor and Intel® Pentium® 4 Processor 6x1 Sequence Thermal and Mechanical Design Guidelines	317804
LGA775 Socket Mechanical Design Guide	302666
uATX SFF Design Guidance	http://www.formfactors.org/
Fan specification for 4-wire PWM Controlled Fans	http://www.formfactors.org/
ATX Thermal Design Suggestions	http://www.formfactors.org/
microATX Thermal Design Suggestions	http://www.formfactors.org/
Balanced Technology Extended (BTX) System Design Guide	http://www.formfactors.org/
Thermally Advantaged Chassis version 1.1	http://www3.intel.com/cd/channel/reseller/asmo-na/eng/products/53211.htm?iid=go_chassis

1.3 Definition of Terms

Table 1. DOT (Sheet 1 of 2)

Term	Description
T_A	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_C	The case temperature of the processor, measured at the geometric center of the topside of the IHS.
T_E	The ambient air temperature external to a system chassis. This temperature is usually measured at the chassis air inlets.
T_S	Heatsink temperature measured on the underside of the heatsink base, at a location corresponding to T_C.

Table 1. DOT (Sheet 2 of 2)

Term	Description
$T_{C\text{-}MAX}$	The maximum case temperature as specified in a component specification.
Ψ_{CA}	Case-to-ambient thermal characterization parameter (psi). A measure of thermal solution performance using total package power. Defined as (T_C – T_A) / Total Package Power. Note: Heat source must 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_C – T_S) / Total Package Power. Note: Heat source must be specified for Ψ measurements.
Ψ_{SA}	Sink-to-ambient thermal characterization parameter. A measure of heatsink thermal performance using total package power. Defined as (T_S – T_A) / Total Package Power. Note: Heat source must be specified for Ψ measurements.
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.
P_{MAX}	The maximum power dissipated by a semiconductor component.
TDP	Thermal Design Power: a power dissipation target based on worst-case applications. Thermal solutions should be designed to dissipate the thermal design power.
IHS	Integrated Heat Spreader: a thermally conductive lid integrated into a processor package to improve heat transfer to a thermal solution through heat spreading.
LGA775 Socket	The surface mount socket designed to accept the processors in the 775–Land LGA package.
ACPI	Advanced Configuration and Power Interface.
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. This design guide covers the requirements for the 2006 FMB (65W).
Thermal Monitor	A feature on the processor that attempts to keep the processor die temperature within factory specifications.
TCC	Thermal Control Circuit: Thermal Monitor uses the TCC to reduce die temperature by lowering effective processor frequency when the die temperature has exceeded its operating limits.
T_{DIODE}	Temperature reported from the on-die thermal diode.
FSC	Fan Speed Control: Thermal solution that includes a variable fan speed which is driven by a PWM signal and uses the on-die thermal diode as a reference to change the duty cycle of the PWM signal.
$T_{CONTROL}$	$T_{CONTROL}$ is the specification limit for use with the on-die thermal diode.
PWM	Pulse width modulation is a method of controlling a variable speed fan. The enabled 4 wire fans use the PWM duty cycle % from the fan speed controller to modulate the fan speed.
Health Monitor Component	Any standalone or integrated component that is capable of reading the processor temperature and providing the PWM signal to the 4 pin fan header.
BTX	Balanced Technology Extended.
TMA	Thermal Module Assembly. The heatsink, fan and duct assembly for the BTX thermal solution.

2.0 Processor Thermal/Mechanical Information

2.1 Mechanical Requirements

2.1.1 Processor Package

The processors covered in this document are packaged in a 775-Land LGA package that interfaces with the motherboard via a LGA775 socket. Please refer to the EMTS for detailed mechanical specifications.

The processor connects to the motherboard through a land grid array (LGA) surface mount socket. The socket contains 775 contacts arrayed about a cavity in the center of the socket with solder balls for surface mounting to the motherboard. The socket is named LGA775 socket. A description of the socket can be found in the *LGA775 Socket Mechanical Design Guide*.

The package includes an integrated heat spreader (IHS) that is shown in Figure 1 on page 10 for illustration only. Refer to the processor EMTS for further information. In case of conflict, the package dimensions in the processor EMTS supersedes dimensions provided in this document.

Figure 1. Package IHS Load Areas

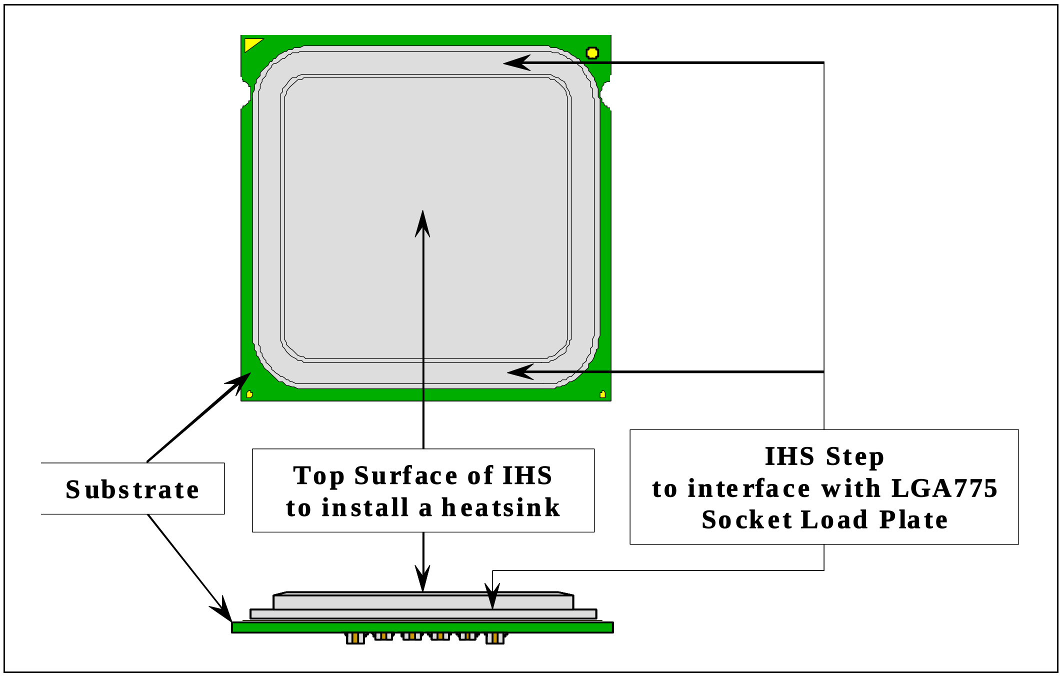

The primary function of the IHS is to transfer the non-uniform heat distribution from the die to the top of the IHS, out of which the heat flux is more uniform and spread over a larger surface area (not the entire IHS area). This allows more efficient heat transfer out of the package to an attached cooling device. The top surface of the IHS is designed to be the interface for contacting a heatsink.

The IHS also features a step that interfaces with the LGA775 socket load plate, as described in *LGA775 Socket Mechanical Design Guide*. The load from the load plate is distributed across two sides of the package onto a step on each side of the IHS. It is then distributed by the package across all of the contacts. When correctly actuated, the top surface of the IHS is above the load plate allowing proper installation of a heatsink on the top surface of the IHS. After actuation of the socket load plate, the seating plane of the package is flush with the seating plane of the socket. Package movement during socket actuation is along the Z direction (perpendicular to substrate) only. Refer to the *LGA775 Socket Mechanical Design Guide* for further information about the LGA775 socket.

The processor package has mechanical load limits that are specified in the processor EMTS. The specified maximum static and dynamic load limits should not be exceeded during their respective stress conditions. These include heatsink installation, removal, mechanical stress testing, and standard shipping conditions.

- When a compressive static load is necessary to ensure thermal performance of the thermal interface material between the heatsink base and the IHS, it should not exceed the corresponding specification given in the processor EMTS.
- When a compressive static load is necessary to ensure mechanical performance, it should remain in the minimum/maximum range specified in the processor EMTS.
- The heatsink mass can also generate additional dynamic compressive load to the package during a mechanical shock event. Amplification factors due to the impact force during shock must be taken into account in dynamic load calculations. The total combination of dynamic and static compressive load should not exceed the processor EMTS compressive dynamic load specification during a vertical shock. For example, with a 0.550 kg [1.2 lb.] heatsink, an acceleration of 50G during an 11 ms trapezoidal shock with an amplification factor of 2 results in approximately a 539 N [117 lbf] dynamic load on the processor package. If a 178 N [40 lbf] static load is also applied on the heatsink for thermal performance of the thermal interface material the processor package could see up to a 717 N [156 lbf]. The calculation for the thermal solution of interest should be compared to the processor EMTS specification.

No portion of the substrate should be used as a load-bearing surface.

Finally, the processor EMTS provides package handling guidelines in terms of maximum recommended shear, tensile and torque loads for the processor IHS relative to a fixed substrate. These recommendations should be followed in particular for heatsink removal operations.

2.1.2 Heatsink Attach

2.1.2.1 General Guidelines

There are no features on the LGA775 socket to directly attach a heatsink. A mechanism must be designed to attach the heatsink directly to the motherboard. In addition to holding the heatsink in place on top of the IHS; this mechanism plays a significant role in the robustness of the system in which it is implemented, in particular:

- Ensuring thermal performance of the thermal interface material (TIM) applied between the IHS and the heatsink. TIMs based on phase change materials are very sensitive to applied pressure: the higher the pressure, the better the initial performance. TIMs such as thermal greases are not as sensitive to applied pressure. Designs should consider a possible decrease in applied pressure over time due to potential structural relaxation in retention components.
- Ensuring system electrical, thermal, and structural integrity under shock and vibration events. The mechanical requirements of the heatsink attach mechanism depend on the mass of the heatsink and the level of shock and vibration that the system must support. The overall structural design of the motherboard and the system have to be considered when designing the heatsink attach mechanism. Their design should provide a means for protecting LGA775 socket solder joints. One of the strategies for mechanical protection of the socket is to use a preload and

high stiffness clip. This strategy is implemented by the reference design and described in document *Intel® Celeron® Processor 400 Sequence Thermal and Mechanical Design Guidelines*.

Note: Package pull-out during mechanical shock and vibration is constrained by the LGA775 socket load plate (refer to the *LGA775 Socket Mechanical Design Guide* for further information).

2.1.2.2 Heatsink Clip Load Requirement

The attach mechanism for the heatsink developed to support the processor should create a static preload on the package between **18 lbf** and **70 lbf** throughout the life of the product for designs compliant with the Intel reference design assumptions:

- 72 mm x 72 mm mounting hole span for ATX (refer to Figure 47)
- TMA preload vs. stiffness for BTX within the limits shown in *Intel® Celeron® Processor 400 Sequence Datasheet*.
- And no board stiffening device (backing plate, chassis attach, etc.).

The minimum load is required to protect against fatigue failure of socket solder joint in temperature cycling.

It is important to take into account potential load degradation from creep over time when designing the clip and fastener to the required minimum load. This means that, depending on clip stiffness, the initial preload at beginning of life of the product may be significantly higher than the minimum preload that must be met throughout the life of the product. For additional guidelines on mechanical design, in particular on designs departing from the reference design assumptions refer to Appendix A.

- For clip load metrology guidelines, refer to *Intel® Celeron® Processor 400 Sequence Thermal and Mechanical Design Guidelines*.

2.1.2.3 Additional Guidelines

In addition to the general guidelines given above, the heatsink attach mechanism for the processor should be designed to the following guidelines:

- *Holds the heatsink in place under mechanical shock and vibration events and applies force to the heatsink base to maintain desired pressure on the thermal interface material*. Note that the load applied by the heatsink attach mechanism must comply with the package specifications described in the processor EMTS. One of the key design parameters is the height of the top surface of the processor IHS above the motherboard. The IHS height from the top of board is expected to vary from 7.517 mm to 8.167 mm. This data is provided for information only, and should be derived from:
  - — The height of the socket seating plane above the motherboard after reflow, given in the *LGA775 Socket Mechanical Design Guide* with its tolerances.
  - — The height of the package, from the package seating plane to the top of the IHS, and accounting for its nominal variation and tolerances that are given in the corresponding processor EMTS.
- *Engages easily, and if possible, without the use of special tools*. In general, the heatsink is assumed to be installed after the motherboard has been installed into the chassis.
- *Minimizes contact with the motherboard surface during installation* and actuation to avoid scratching the motherboard.

2.2 Thermal Requirements

Refer to the EMTS for the processor thermal specifications. The majority of processor power is dissipated through the IHS. There are no additional components, e.g., BSRAMs, which generate heat on this package. The amount of power that can be dissipated as heat through the processor package substrate and into the socket is usually minimal.

The thermal limits for the processor are the Thermal Profile and $T_{CONTROL}$. The Thermal Profile defines the maximum case temperature as a function of power being dissipated. $T_{CONTROL}$ is a specification used in conjunction with the temperature reported by the digital thermal sensor and a fan speed control method. Designing to these specifications allows optimization of thermal designs for processor performance and acoustic noise reduction.

2.2.1 Processor Case Temperature

For the processor, the case temperature is defined as the temperature measured at the geometric center of the package on the surface of the IHS. For illustration, Figure 2 shows the measurement location for a 37.5 mm x 37.5 mm [1.474 in x 1.474 in] 775-Land LGA processor package with a 28.7 mm x 28.7 mm [1.13 in x 1.13 in] IHS top surface. Techniques for measuring the case temperature are detailed in Chapter 3.4, "Processor Case Temperature Measurement Guidelines."

Note: In case of conflict, the package dimensions in the processor EMTS supersedes dimensions provided in this document.

Figure 2. Processor Case Temperature Measurement Location

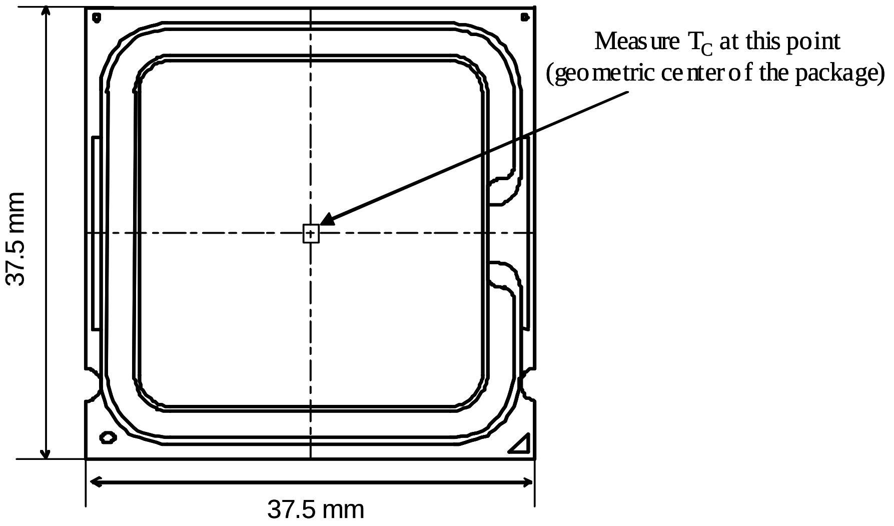

2.2.2 Thermal Profile

The Thermal Profile defines the maximum case temperature as a function of processor power dissipation. Refer to *Intel® Celeron® Processor 440 EMTS* for further information. Table 2 shows the thermal specification for the thermal design power (TDP) segment.

Table 2. Intel® Celeron® Processor 440 Thermal Specifications

Processor	FMB	TDP(W)	$T_{C\text{-}MAX}$	Notes
Intel® Celeron® Processor 440	2006	35	60.4 °C	1,2,3

Notes:

1. Thermal Design power (TDP) should be used for processor thermal solution design targets. The TDP is not the maximum power that the processor can dissipate.
2. FMB, or Flexible Motherboard, guidelines provide a design target for meeting future thermal requirements.
3. T_C and TDP values provided in this table are for reference only. Please contact your Intel field representative for any updates that could occur in the processor EMTS prior to the next revision of this document.

Thermal solutions that are intended to function in a multitude of systems and environments need to be designed for the worst-case thermal environment. The majority of ATX/BTX platforms are targeted to function in an environment that will have up to a 35C ambient temperature external to the system (Table 3).

Table 3. Intel® Celeron® Processor 440 Thermal Targets

Processor	FMB	TDP(W)	$T_{C\text{-}MAX}$	T_A	Max Ψ_{CA} (°C/W)	Notes
Intel® Celeron® Processor 440	2006	35	60.4 °C	35	0.73	1,2,3,4
Intel® Celeron® Processor 440	2006	35	60.4 °C	40	0.58	1,2,3,4

Notes:

1. Thermal Module solutions that meet the Max Ψ_{CA} requirement must also satisfy the thermal profile requirement.
2. T_C, TDP values provided in this table are for reference only. Please contact your Intel field representative for any updates that could occur in the processor EMTS prior to the next revision of this document.
3. To determine compliance to the thermal profile, a measurement of the actual processor power dissipation is required. The measured power is plotted on the Thermal Profile to determine the maximum case temperature. Using the example in Figure 3 for a processor dissipating 35W the maximum case temperature is 60.4°C. See the EMTS for the thermal profile. See the *Thermal Readiness for Performance FMB Platforms User's Guide* for specific details on power measurement.

Figure 3. Example Thermal Profile

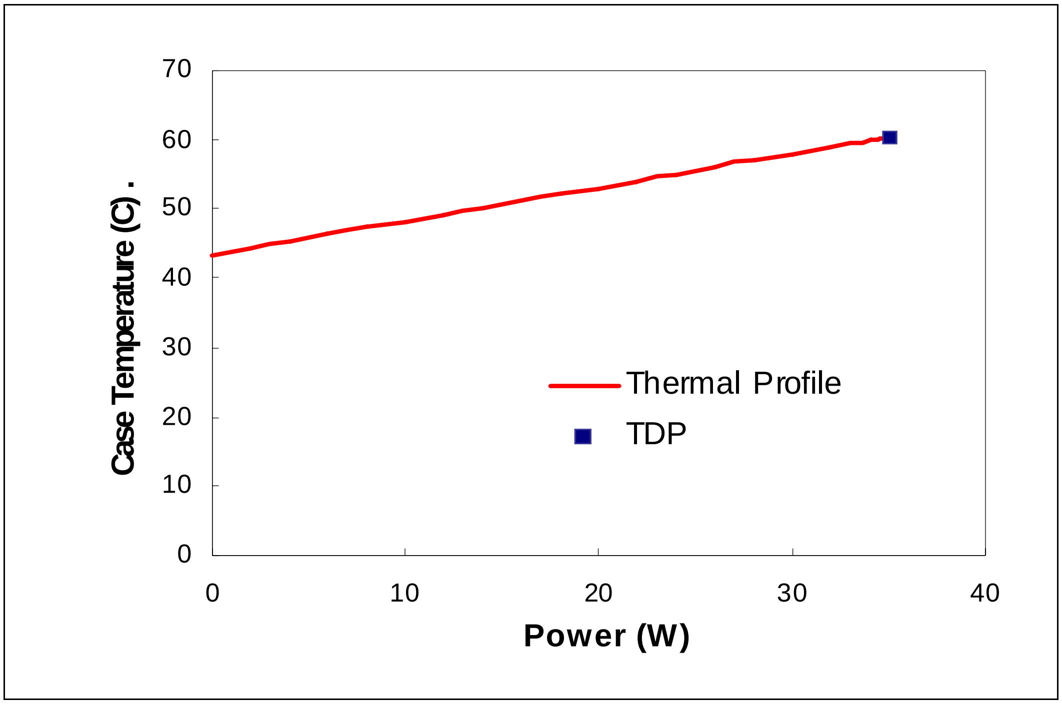

2.2.3 $T_{CONTROL}$

$T_{CONTROL}$ defines the maximum operating temperature for the digital thermal sensor when the thermal solution fan speed is being controlled by the digital thermal sensor. The $T_{CONTROL}$ parameter defines a very specific processor operating region where fan speed can be reduced. This allows the system integrator a method to reduce the acoustic noise of the processor cooling solution, while maintaining compliance to the processor thermal specification.

Note: The $T_{CONTROL}$ value for the Intel Celeron processor 440 is relative to the Thermal Control Circuit (TCC) activation set point which will be seen as 0 (zero) via the digital thermal sensor. As a result, the $T_{CONTROL}$ value will always be a negative number. See Chapter 4.0, "Thermal Management Logic and Thermal Monitor Feature" for the discussion the thermal management logic and features and Chapter 6.0, "Intel® Quiet System Technology (Intel® QST)."

The value of $T_{CONTROL}$ is driven by a number of factors. One of the most significant of these is the processor idle power. As a result, a processor with a high $T_{CONTROL}$ will dissipate more power than a part with lower value of $T_{CONTROL}$ when running the same application.

The value of $T_{CONTROL}$ is calculated such that regardless of the individual processor's $T_{CONTROL}$ value, the thermal solution should perform similarly. The higher power of some parts is offset by a higher value of $T_{CONTROL}$ in such a way that they should behave virtually the same acoustically. This is achieved in part by using the Ψ_{CA} vs. RPM and RPM vs. Acoustics (dBA) performance curves from the Intel enabled thermal solution. A thermal solution designed to meet the thermal profile should have similar acoustic performance for any value of $T_{CONTROL}$.

The value for $T_{CONTROL}$ is calculated by the system BIOS based on values read from a factory configured processor register. The result can be used to program a fan speed control component. See the appropriate EMTS for further details on reading the register and calculating $T_{CONTROL}$.

See Chapter 6.0, "Intel® Quiet System Technology (Intel® QST)" for details on implementing a design using $T_{CONTROL}$ and the Thermal Profile.

2.3 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 stricter. Thermal interface material (TIM) is used to fill in the gap between the IHS and the bottom surface of the heatsink, and thereby improve the overall performance of the stack-up (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 applied to it. Refer to Chapter 2.3.4, "Thermal Interface Material" and Appendix B, "Thermal Interface Management" for further information on TIM and on bond line management 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_A, and the local air velocity over the surface. The higher the air velocity over the surface, and the cooler the air, the more efficient is the resulting cooling. 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 in particular the fin faces and the heatsink base.

Active heatsinks typically incorporate 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 lower air speed. These heatsinks are therefore 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 travels around the heatsink instead of through it, unless air bypass is carefully managed. Using air-ducting techniques to manage bypass area can be an effective method for controlling airflow through the heatsink.

2.3.1 Heatsink Size

The size of the heatsink is dictated by height restrictions for installation in a system and by the real estate available on the motherboard and other considerations for component height and placement in the area potentially impacted by the processor heatsink. The height of the heatsink must comply with the requirements and recommendations published for the motherboard form factor of interest. Designing a heatsink to the recommendations may preclude using it in system adhering strictly to the form factor requirements, while still in compliance with the form factor documentation.

For the PICMG 1.3 server form factor, it is recommended to use:

- The PICMG 1.3 motherboard keep-out footprint definition and height restrictions for enabling components, defined for the platforms designed with the LGA775 socket in Appendix E, "Mechanical Drawings" of this design guide.
- The motherboard primary side height constraints are located at http://picmg.org/specifications.stm.
- For information regarding EmbATX the following document can be used http://www.intel.com/design/intarch/platforms/iaclient/eatx/index.htm.

The resulting space available above the motherboard is generally not entirely available for the heatsink. The target height of the heatsink must take into account airflow considerations (for fan performance for example) as well as other design considerations (air duct, etc.).

2.3.2 Heatsink Mass

With the need to push air cooling to better performance, heatsink solutions tend to grow larger (increase in fin surface) resulting in increased mass. The insertion of highly thermally conductive materials like copper to increase heatsink thermal conduction performance results in even heavier solutions. As mentioned in *Chapter 2.1, "Mechanical Requirements,"* the heatsink mass must take into consideration the package and socket load limits, the heatsink attach mechanical capabilities, and the mechanical shock and vibration profile targets. Beyond a certain heatsink mass, the cost of developing and implementing a heatsink attach mechanism that can ensure the system integrity under the mechanical shock and vibration profile targets may become prohibitive.

2.3.3 Package IHS Flatness

The package IHS flatness for the product is specified in the EMTS and can be used as a baseline to predict heatsink performance during the design phase.

Intel recommends testing and validating heatsink performance in full mechanical enabling configuration to capture any impact of IHS flatness change due to combined socket and heatsink loading. While socket loading alone may increase the IHS warpage, the heatsink preload redistributes the load on the package and improves the resulting IHS flatness in the enabled state.

2.3.4 Thermal Interface Material

Thermal interface material application between the processor IHS and the heatsink base is 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 thermal interface material 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 thermal interface material 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.

2.4 System Thermal Solution Considerations

2.4.1 Chassis Thermal Design Capabilities

The Intel reference thermal solution for PICMG 1.3 chassis assumes that the chassis delivers a maximum T_A of 38-40°C with 15-25 CFM of airflow at the inlet of the processor heatsink.

2.4.2 Improving Chassis Thermal Performance

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 and vents 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, acoustic requirements and structural considerations that limit the thermal solution size. For more information, refer to the *Thin Electronics Bay specification* at the following web site www.ssiforum.org.

In addition to passive heatsinks, fan heatsinks and system fans are other solutions that 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.

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 TDP instead of maximum power. Thermal Monitor attempts to protect the processor during sustained workload above TDP. Implementation options and recommendations are described in Chapter 4.0, "Thermal Management Logic and Thermal Monitor Feature."

2.4.3 Summary

In summary, considerations in heatsink design include:

- The local ambient temperature T_A at the heatsink, which is a function of chassis design.
- The thermal design power (TDP) of the processor, and the corresponding maximum T_C as calculated from the thermal profile. 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 Chapter 3.1, "Characterizing Cooling Performance Requirements."
- Heatsink interface to IHS surface characteristics, including flatness and roughness.

- The performance of the thermal interface material used between the heatsink and the IHS.
- The required heatsink clip static load, between 18 lbf to 70 lbf throughout the life of the product (Refer to Chapter 2.1.2.2, “Heatsink Clip Load Requirement” for further information).
- Surface area of the heatsink.
- Heatsink material and technology.
- Volume of airflow over the heatsink surface area.
- Development of airflow entering and within the heatsink area.
- Physical volumetric constraints placed by the system.

2.5 System Integration Considerations

Manufacturing with Intel® Components using 775–Land LGA Package and LGA775 Socket documentation provides Best Known Methods for all aspects LGA775 socket based platforms and systems manufacturing. Of particular interest for package and heatsink installation and removal is the *System Assembly* module. A video covering system integration is also available. Contact your Intel field sales representative for further information.

3.0 Thermal Metrology

This section discusses guidelines for testing thermal solutions, including measuring processor temperatures. In all cases, the thermal engineer must measure power dissipation and temperature to validate a thermal solution. To define the performance of a thermal solution, the “thermal characterization parameter”, Ψ (“psi”) will be used.

3.1 Characterizing Cooling Performance Requirements

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 situations (same heat source and local ambient conditions). The thermal characterization parameter is calculated using total package power.

Note: Heat transfer is a three-dimensional phenomenon that can rarely be accurately and easily modeled by a single resistance parameter like Ψ.

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 Equation 1, and measured in units of °C/W.

Equation 1.

$$\Psi_{CA} = (T_C - T_A) / P_D$$

Where:

Ψ_{CA} = Case-to-local ambient thermal characterization parameter (°C/W)

T_C = Processor case temperature (°C)

T_A = Local ambient temperature in chassis at processor (°C)

P_D = Processor total power dissipation (W) (assumes all power dissipates through the IHS)

The case-to-local ambient thermal characterization parameter of the processor, Ψ_{CA}, is comprised of Ψ_{CS}, the thermal interface material thermal characterization parameter, and of Ψ_{SA}, the sink-to-local ambient thermal characterization parameter as defined in Equation 2.

Equation 2.

$$\Psi_{CA} = \Psi_{CS} + \Psi_{SA}$$

Where:

Ψ_{CS} = Thermal characterization parameter of the thermal interface material (°C/W)

Ψ_{SA} = Thermal characterization parameter from heatsink-to-local ambient (°C/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.

Figure 4 illustrates the combination of the different thermal characterization parameters.

Figure 4. Processor Thermal Characterization Parameter Relationships

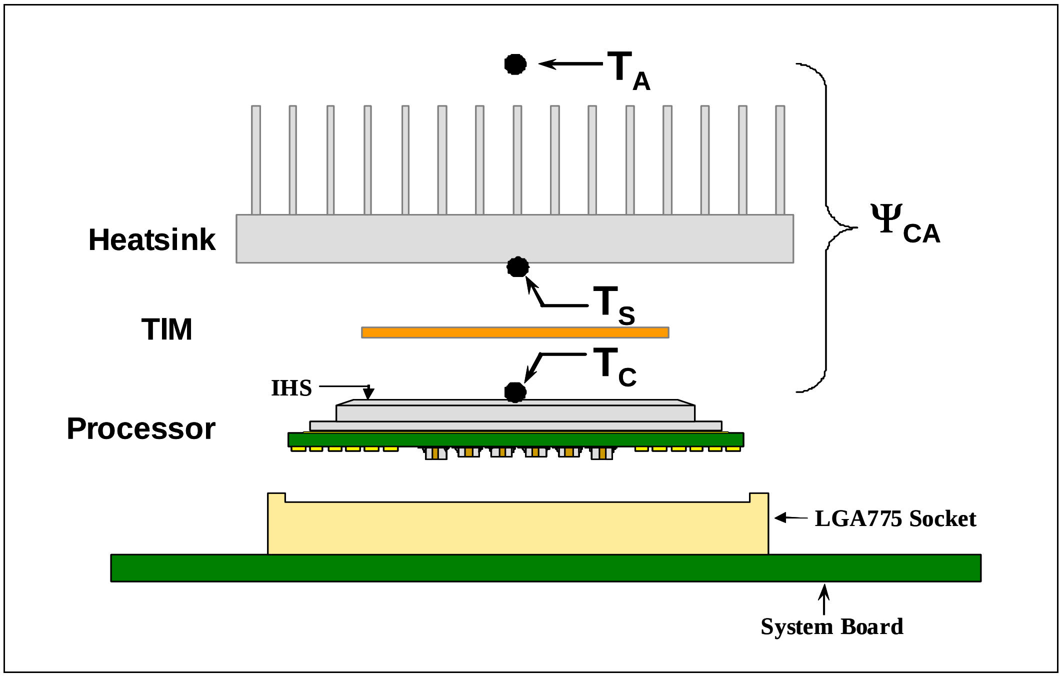

3.1.1 Example

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

- The case temperature $T_{C\text{-}MAX}$ and thermal design power TDP given in the processor Datasheet.
- Define a target local ambient temperature at the processor, T_A.

Since the processor thermal profile applies to all processor frequencies, it is important to identify the worst case (lowest Ψ_{CA}) for a targeted chassis characterized by T_A to establish a design strategy.

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 specific Intel processor thermal specifications, and are for illustrative purposes only.

Assume the TDP, as listed in the EMTS, is 35 W and the maximum case temperature from the thermal profile for 35W is 60.4 °C. Assume as well that the system airflow has been designed such that the local ambient temperature is 38°C. Then the following could be calculated using Equation 1:

$$\Psi_{CA} = (T_{C,} - T_A) / TDP = (60.4 – 38) / 35 = 0.64\ °C/W$$

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 were designed to work with a TIM material performing at $\Psi_{CS} \leq 0.10$ °C/W, solving for Equation 2, the performance of the heatsink would be:

$$\Psi_{SA} = \Psi_{CA} – \Psi_{CS} = 0.64 – 0.10 = 0.54\ °C/W$$

3.2 Processor Thermal Solution Performance Assessment

Thermal performance of a heatsink should be assessed using a thermal test vehicle (TTV) provided by Intel. The TTV is a stable heat source that the user can make accurate power measurements, whereas processors can introduce additional factors that can impact test results. In particular, the power level from actual processors varies significantly, even when running the maximum power application provided by Intel, due to variances in the manufacturing process. The TTV provides consistent power and power density for thermal solution characterization and results can be easily translated to real processor performance. Accurate measurement of the power dissipated by an actual processor is beyond the scope of this document. See the *Thermal Readiness for Performance FMB Platforms User's Guide* for further information.

Once the thermal solution is designed and validated with the TTV, it is strongly recommended to verify functionality of the thermal solution on real processors and on fully integrated systems. The Intel maximum power application enables steady power dissipation on a processor to assist in this testing. This application is called Maximum Power Program for the Intel Celeron processor 440. Contact your Intel Field Sales representative for a copy of the latest release of this application.

3.3 Local Ambient Temperature Measurement Guidelines

The local ambient temperature T_A is the temperature of the ambient air surrounding the processor. For a passive heatsink, T_A is defined as the heatsink approach air temperature. For an actively cooled heatsink, it is the temperature of inlet air to the active cooling fan.

It is worthwhile to determine the local ambient temperature in the chassis around the processor to understand the effect it may have on the case temperature.

T_A is best measured by averaging temperature measurements at multiple locations in the heatsink inlet airflow. This method helps reduce error and eliminate minor spatial variations in temperature. The following guidelines are meant to enable accurate determination of the localized air temperature around the processor during system thermal testing.

For **active heatsinks**, it is important to avoid taking a measurement in the dead flow zone that usually develops above the fan hub and hub spokes. Measurements should be taken at four different locations uniformly placed at the center of the annulus formed by the fan hub and the fan housing to evaluate the uniformity of the air temperature at the fan inlet. The thermocouples should be placed approximately 3 mm to 8 mm [0.1 to 0.3 in] above the fan hub vertically and halfway between the fan hub and the fan housing horizontally as shown in Figure 5 on page 23 (avoiding the hub spokes). Using an open bench to characterize an active heatsink can be useful, and usually ensures more uniform temperatures at the fan inlet. However, additional tests that include a solid barrier above the test motherboard surface can help evaluate the potential impact of the chassis. This barrier is typically clear Plexiglas*, extending at least 100 mm [4 in] in all directions beyond the edge of the thermal solution. Typical distance from the motherboard to the barrier is 81 mm [3.2 in]. For even more realistic airflow, the motherboard should be populated with significant elements like memory cards, a graphic card, and a chipset heatsink. If a barrier is used, the thermocouple can be taped directly to the barrier with a clear tape at the horizontal location as previously described, half way between the fan hub and the fan housing. If a variable speed fan is used, it may be useful to add a thermocouple taped to the barrier above the location of the temperature sensor used by the fan to check its speed setting against air temperature. When measuring T_A in a chassis with a live motherboard, add-in cards, and other system components, it is likely that the T_A measurements will reveal a highly non-uniform temperature distribution across the inlet fan section.

For **passive heatsinks**, thermocouples should be placed approximately 13 mm to 25 mm [0.5 to 1.0 in] away from processor and heatsink as shown in Figure 6. The thermocouples should be placed approximately 51 mm [2.0 in] above the baseboard. This placement guideline is meant to minimize the effect of localized hot spots from baseboard components.

Note: Testing an active heatsink with a variable speed fan can be done in a thermal chamber to capture the worst-case thermal environment scenarios. Otherwise, when doing a bench top test at room temperature, the fan regulation prevents the heatsink from operating at its maximum capability. To characterize the heatsink capability in the worst-case environment in these conditions, it is then necessary to disable the fan regulation and power the fan directly, based on guidance from the fan supplier.

Figure 5. Locations for Measuring Local Ambient Temperature, Active Heatsink

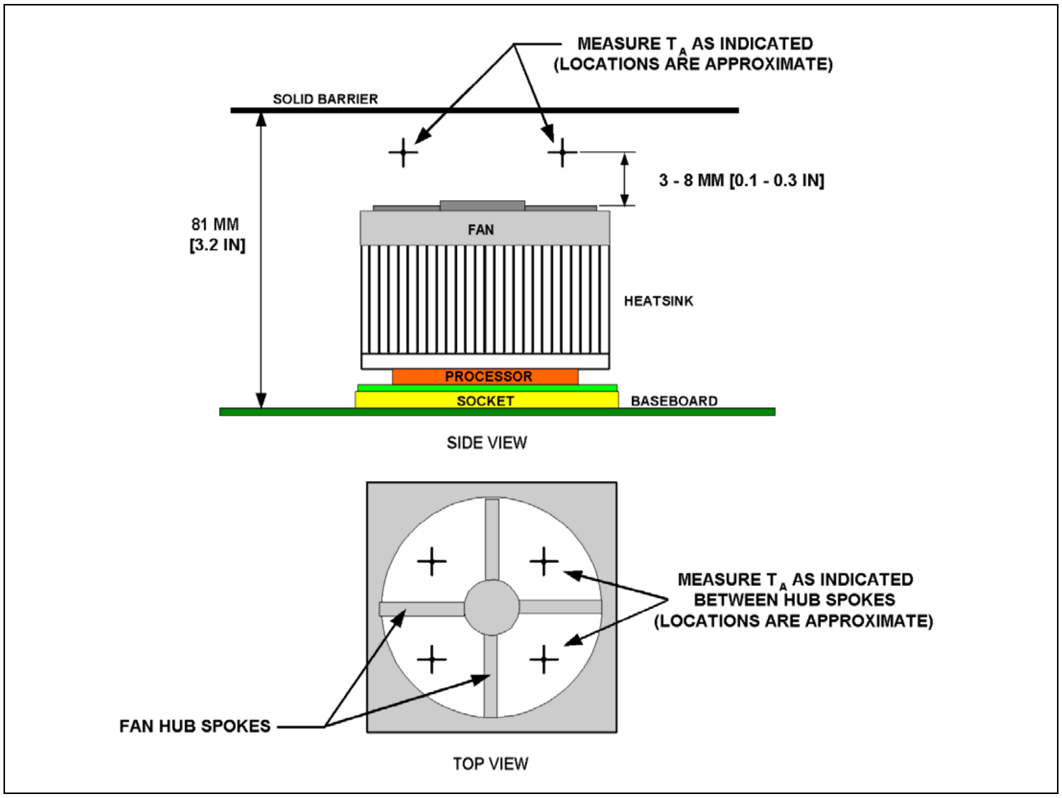

Note: Drawing Not to Scale

Figure 6. Locations for Measuring Local Ambient Temperature, Passive Heatsink

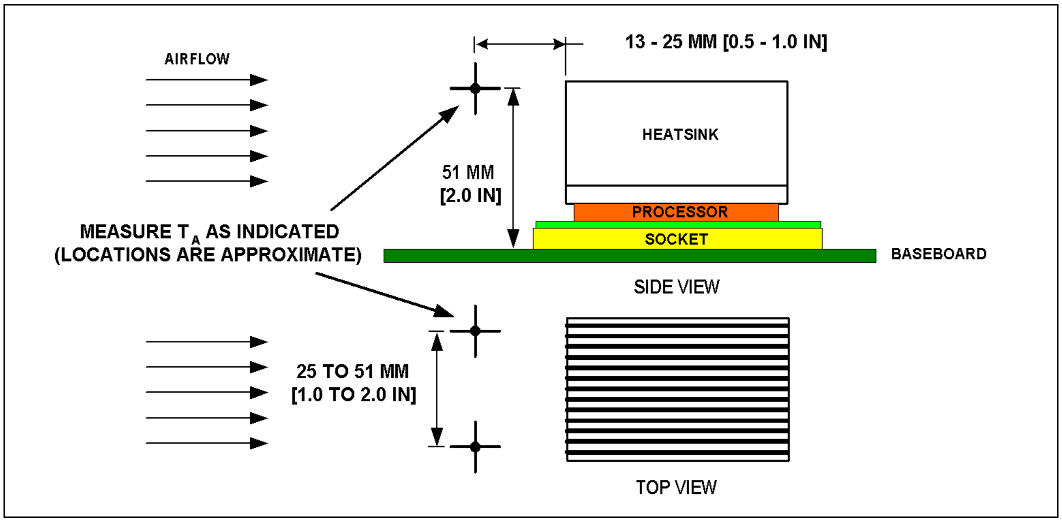

Note: Drawing Not to Scale

3.4 Processor Case Temperature Measurement Guidelines

To ensure functionality and reliability, the processor is specified for proper operation when T_C is maintained at or below the thermal profile as listed in the EMTS. The measurement location for T_C is the geometric center of the IHS. Figure 2 shows the location for T_C measurement.

Special care is required when measuring T_C to ensure an accurate temperature measurement. Thermocouples are often used to measure T_C. Before any temperature measurements are made, the thermocouples must be calibrated, and the complete measurement system must be routinely checked against known standards. 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 junction of the thermocouple 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.

Appendix C, “Case Temperature Reference Metrology” defines a reference procedure for attaching a thermocouple to the IHS of a 775-Land LGA processor package for T_C measurement. This procedure takes into account the specific features of the 775-Land LGA package and of the LGA775 socket for which it is intended.

4.0 Thermal Management Logic and Thermal Monitor Feature

4.1 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 processor. 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 Thermal Control Circuit (TCC), the processor can rapidly initiate thermal management control. The Thermal Monitor can reduce cooling solution cost, by allowing thermal designs to target TDP.

The processor also supports an additional power reduction capability known as Thermal Monitor 2 described in Chapter 4.2.3.

4.2 Thermal Monitor Implementation

The Thermal Monitor consists of the following components:

- A highly accurate on-die temperature sensing circuit.
- A bi-directional signal (PROCHOT#) that indicates if the processor has exceeded its maximum temperature or can be asserted externally to activate the Thermal

Control Circuit (TCC) (see Chapter 4.2.1, “PROCHOT# Signal” for more details on user activation of TCC via PROCHOT# signal).

- A Thermal Control Circuit that will attempt to reduce processor temperature by rapidly reducing power consumption when the on-die temperature sensor indicates that it has exceeded the maximum operating point.
- Registers to determine the processor thermal status.

4.2.1 PROCHOT# Signal

The primary function of the PROCHOT# signal is to provide an external indication the processor has exceeded its maximum operating temperature. While PROCHOT# is asserted, the TCC will be active. 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.

PROCHOT# can be configured via BIOS as an output or bi-directional signal. As an output, PROCHOT# will go active when the processor temperature of either core exceeds its maximum operating temperature. This indicates the TCC has been activated. As an input, assertion of PROCHOT# will activate the TCC for both cores. The TCC will remain active until the system de-asserts PROCHOT#.

The temperature at which the PROCHOT# signal goes active is individually calibrated during manufacturing. The power dissipation of each processor affects the set point temperature. The temperature where PROCHOT# goes active roughly parallels the thermal profile. Once configured, the processor temperature at which the PROCHOT# signal is asserted is not re-configurable.

One application is the thermal protection of voltage regulators (VR). System designers can create a circuit to monitor the VR temperature and activate the TCC when the temperature limit of the VR is reached. By asserting PROCHOT# (pulled-low) or FORCEPR#, which activates the TCC, the VR can cool down as a result of reduced processor power consumption. Bi-directional PROCHOT# can allow VR thermal designs to target maximum sustained current instead of maximum current. Systems should still provide proper cooling for the VR, and rely on bi-directional PROCHOT# signal only as a backup in case of system cooling failure.

Note: A thermal solution designed to meet the thermal profile targets should rarely experience activation of the TCC as indicated by the PROCHOT# signal going active.

4.2.2 Thermal Control Circuit

The Thermal Control Circuit portion of the Thermal Monitor must be enabled for the processor to operate within specifications. The Thermal Monitor’s TCC, when active, will attempt to lower the processor temperature by reducing the processor power consumption. In the original implementation of thermal monitor, 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 µs. This time period is frequency dependent and higher frequency processors will disable the internal clocks for a shorter time period. Figure 7 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 the Thermal Monitor behavior.

Figure 7. Concept for Clocks under Thermal Monitor Control

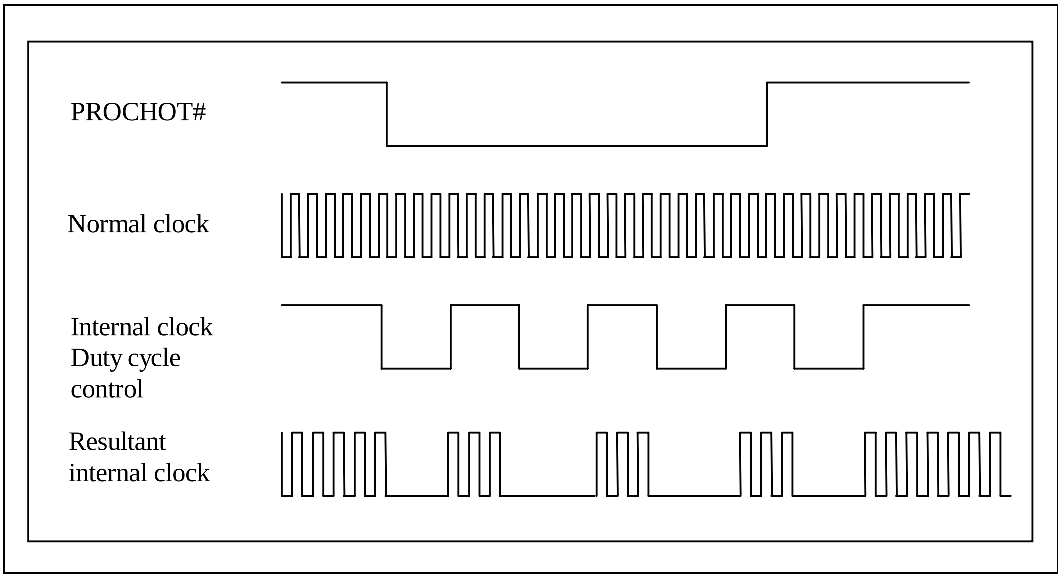

4.2.3 Thermal Monitor 2

The processor supports an enhanced Thermal Control Circuit. In conjunction with the existing Thermal Monitor logic, this capability is known as Thermal Monitor 2. This enhanced TCC provides an efficient means of reducing the power consumption within the processor and limiting the processor temperature.

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 (by dropping the bus-to-core multiplier to its minimum available value) and input voltage identification (VID) value. This combination of reduced frequency and 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, all bus traffic is blocked. 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 VID transitions 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 the power consumption of the processor, providing a temperature reduction.

Once the processor has sufficiently cooled, and a minimum activation time has expired, the operating frequency and voltage transition back to the normal system operating point. Transition of the VID code will occur first to insure proper operation once the processor reaches its normal operating frequency. Refer to Figure 8 on page 27 for an illustration of this ordering.

Figure 8. Thermal Monitor 2 Frequency and Voltage Ordering

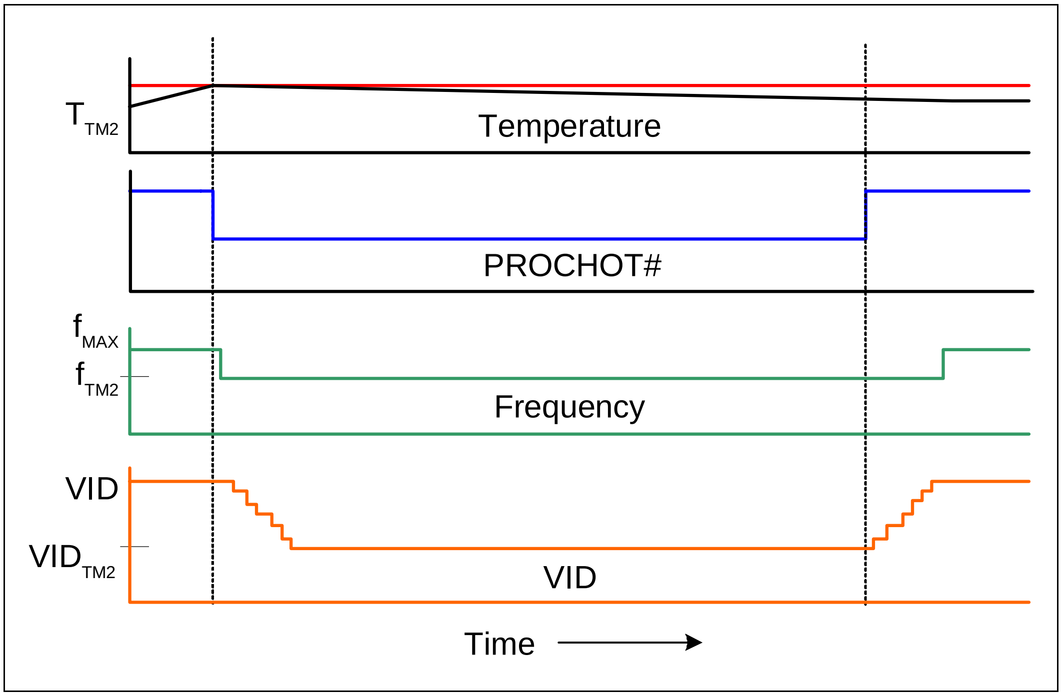

Refer to the EMTS for further information on Thermal Monitor 2.

4.2.4 Operation and Configuration

To maintain compatibility with previous generations of processors, which have no integrated thermal logic, the Thermal Control Circuit portion of Thermal Monitor is disabled by default. During the boot process, the BIOS must enable the Thermal Control Circuit.

Note: Thermal Monitor must be enabled to ensure proper processor operation.

The Thermal Control Circuit feature can be configured and monitored in a number of ways. OEMs are required to enable the Thermal Control Circuit while using various registers and outputs to monitor the processor thermal status. The Thermal Control Circuit is enabled by the BIOS setting a bit in an MSR (model specific register). Enabling the Thermal Control Circuit allows the processor to attempt to maintain a safe operating temperature without the need for special software drivers or interrupt handling routines. When the Thermal Control Circuit has been enabled, processor power consumption will be reduced after the thermal sensor detects a high temperature, i.e. PROCHOT# assertion. The Thermal Control Circuit and PROCHOT# transitions 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 always indicate the thermal status of the processor.

The power reduction mechanism of thermal monitor can also be activated manually using an “on-demand” mode. Refer to Chapter 4.2.5, “On-Demand Mode” for details on this feature.

4.2.5 On-Demand Mode

For testing purposes, the thermal control circuit may also be activated by setting bits in the ACPI MSRs. The MSRs may be set based on a particular system event (e.g., an interrupt generated after a system event), or may be set at any time through the operating system or custom driver control thus forcing the thermal control circuit on. This is referred to as “on-demand” mode. Activating the thermal control circuit may be useful for thermal solution investigations or for performance implication studies. When using the MSRs to activate the on-demand clock modulation feature, the duty cycle is configurable in steps of 12.5%, from 12.5% to 87.5%.

For any duty cycle, the maximum time period the clocks are disabled is ~3 µs. This time period is frequency dependent, and decreases as frequency increases. To achieve different duty cycles, the length of time that the clocks are disabled remains constant, and the time period that the clocks are enabled is adjusted to achieve the desired ratio. For example, if the clock disable period is 3 µs, and a duty cycle of ¼ (25%) is selected, the clock on time would be reduced to approximately 1 µs [on time (1 µs) ÷ total cycle time (3 + 1) µs = ¼ duty cycle]. Similarly, for a duty cycle of 7/8 (87.5%), the clock on time would be extended to 21 µs [21 ÷ (21 + 3) = 7/8 duty cycle].

In a high temperature situation, if the thermal control circuit and ACPI MSRs (automatic and on-demand modes) are used simultaneously, the fixed duty cycle determined by automatic mode would take precedence.

Note: On-demand mode can not activate the power reduction mechanism of Thermal Monitor 2.

4.2.6 System Considerations

Intel requires the Thermal Monitor and Thermal Control Circuit to be enabled for all processors. The thermal control circuit is intended to protect against short term thermal excursions that exceed the capability of a well designed processor thermal solution. Thermal Monitor should not be relied upon to compensate for a thermal solution that does not meet the thermal profile up to the thermal design power (TDP).

Each application program has its own unique power profile, although the profile has some variability due to loop decisions, I/O activity and interrupts. In general, computer intensive applications with a high cache hit rate dissipate more processor power than applications that are I/O intensive or have low cache hit rates.

The processor TDP is based on measurements of processor power consumption while running various high power applications. This data 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 is used to derive the TDP targets published in the processor EMTS.

A system designed to meet the thermal profile at TDP and $T_{C\text{-}MAX}$ values published in the processor EMTS greatly reduces the probability of real applications causing the thermal control circuit to activate under normal operating conditions. Systems that do not meet these specifications could be subject to more frequent activation of the thermal control circuit 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 shut down and signal THERMTRIP#.

For information regarding THERMTRIP#, refer to the processor EMTS and to Chapter 4.2.8, “THERMTRIP# Signal” of this Thermal Design Guidelines.

4.2.7 Operating System and Application Software Considerations

The Thermal Monitor feature and its thermal control circuit work seamlessly with ACPI compliant operating systems. The Thermal Monitor feature is transparent to application software since the processor bus snooping, ACPI timer, and interrupts are active at all times.

Refer to the *BIOS Writer’s Guide* for specific programming details on the thermal control circuit enabling sequence.

4.2.8 THERMTRIP# Signal

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# goes active and power must be removed from the processor. THERMTRIP# activation is independent of processor activity and does not generate any bus cycles. Refer to the processor EMTS for more information about THERMTRIP#.

The temperature where the THERMTRIP# signal goes active is individually calibrated during manufacturing. The temperature where THERMTRIP# goes active is roughly parallel to the thermal profile and greater than the PROCHOT# activation temperature. Once configured, the temperature at which the THERMTRIP# signal is asserted is neither re-configurable nor accessible to the system.

4.2.9 Cooling System Failure Warning

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 thermal control circuit would allow the system to continue functioning or allow a normal 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.

4.2.10 Digital Thermal Sensor

The Intel Celeron processor 440 uses Thermal Diode and Digital Thermal Sensor (DTS) as the on-die sensor to use for fan speed control (FSC). This section will document the on-die thermal diode used in the digital thermal sensor. The DTS is monitoring the same sensor that activates the TCC (see Chapter 4.2.2, “Thermal Control Circuit”). Readings from the DTS are relative to the activation of the TCC. The DTS value where TCC activation occurs is 0 (zero).

The DTS can be accessed by two methods. The first is via a MSR. The value read via the MSR is an unsigned number of degrees C away from TCC activation. The second method which is expected to be the primary method for FSC is via the PECI interface. The value of the DTS when read via the PECI interface is always negative and again is degrees C away from TCC activation.

A $T_{CONTROL}$ value will be provided for use with DTS. The usage model for $T_{CONTROL}$ with the DTS is the same as with the on-die thermal diode:

- If the Digital thermal sensor is less than $T_{CONTROL}$, the fan speed can be reduced.

- If the Digital thermal sensor is greater than or equal to $T_{CONTROL}$, then T_C must be maintained at or below the Thermal Profile for the measured power dissipation.

The calculation of $T_{CONTROL}$ is slightly different from previous product. There is no base value to sum with the T_{OFFSET} located in the same MSR as used in previous processors. The BIOS only needs to read the T_{OFFSET} MSR and provide this value to the fan speed control device.

Figure 9. $T_{CONTROL}$ for Digital Thermal Sensor

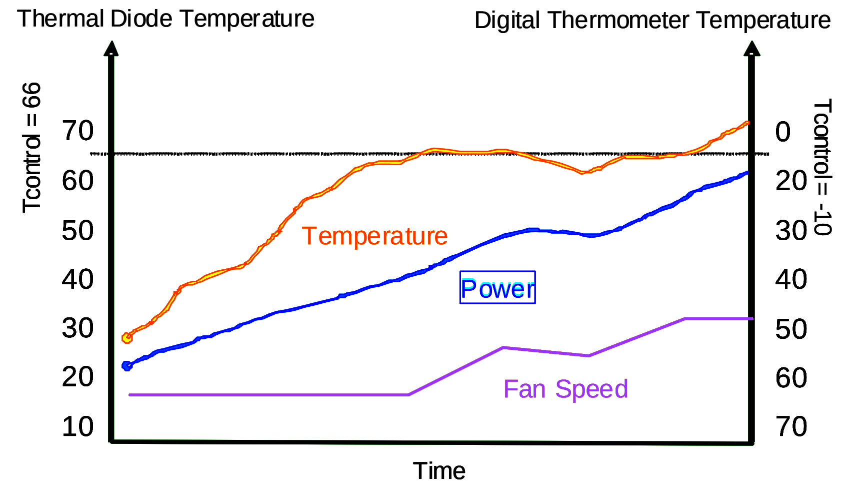

Multiple digital thermal sensors can be implemented within the package without adding a pair of signal pins per sensor as required with the thermal diode. The digital thermal sensor is easier to place in thermally sensitive locations of the processor than the thermal diode. This is achieved due to a smaller foot print and decreased sensitivity to noise. Since the DTS is factory set on a per-part basis, there is no need for the health monitor components to be updated at each processor family.

Note: Intel Celeron processor 440 does not have an on-die thermal diode. The $T_{CONTROL}$ in the MSR is relevant only to the DTS.

4.2.11 Platform Environmental Control Interface (PECI)

The PECI interface is a proprietary single wire bus between the processor and the chipset or other health monitoring device. At this time, the digital thermal sensor is the only data being transmitted. For an overview of the PECI interface see PECI Feature Set Overview. For additional information on the PECI see the *Intel® Celeron® Processor 440 EMTS*.

The PECI bus is available on pin G5 of the LGA 775 socket. Intel chipsets beginning with the ICH8 have included the PECI host controller. The PECI interface and the Manageability Engine, embedded in the Intel 965 Express chipset family, are key elements to the Intel QST, see Chapter 6 and the *Intel® Quiet System Technology (Intel® QST) Configuration and Tuning Manual*.

Intel has worked with many vendors that provide fan speed control devices to provide PECI host controllers. Consult the local representative for your preferred vendor for their product plans and availability.

5.0 Intel Reference Thermal Solution

5.1 Thermal Solution Requirements

The thermal performance required for the heatsink is determined by calculating the case-to-ambient thermal characterization parameter, Ψ_{CA}, as explained in Chapter 3.1. This is a basic thermal engineering parameter that may be used to evaluate and compare different thermal solutions in similar boundary conditions. For the Intel Celeron processor 440, an example of how Ψ_{CA} is calculated is shown in Equation 3.

Equation 3. Case-to-Ambient Thermal Characterization Parameter

$$\Psi_{CA} = \frac{T_{C\max}(^{\circ}C) - T_{LA}(^{\circ}C)}{TDP(W)} = \frac{60.4^{\circ}C - 38^{\circ}C}{35W} = 0.640\frac{^{\circ}C}{W}$$

In this calculation, $T_{C\ max}$ and TDP are taken from the thermal profile specification in the *Intel® Celeron® Processor 440 EMTS*. It is important to note that in this calculation, the $T_{C\ max}$ and TDP are constant, while Ψ_{CA} will vary according to the local ambient temperature (T_{LA}).

Table 4 shows an example of required thermal characterization parameters for the thermal solution at various T_{LA}s. This table uses the $T_{C\ max}$ and TDP from the *Intel® Celeron® Processor 440 EMTS*. These numbers are subject to change and in case of conflict the specifications in the processor EMTS supersede the $T_{C\ MAX}$ and TDP specifications in this document.

Table 4. Thermal Characterization Parameter at various TLAs

Intel® Celeron® Processor 440		**Required Ψ_{CA} (°C/W) of Thermal Solution at T_{LA} = (°C)**			
TDP (W)	T_{cMax} (°C)	44.1	40	35	30
35	60.4	0.466	0.583	0.726	0.869

Figure 10 on page 32 further illustrates the required thermal characterization parameter for the Intel Celeron processor 440 at various operating ambient temperatures. The thermal solution design must have a Ψ_{CA} less than the values shown for the given local ambient temperature.

Figure 10. Thermal Characterization Parameters for Various Operating Conditions

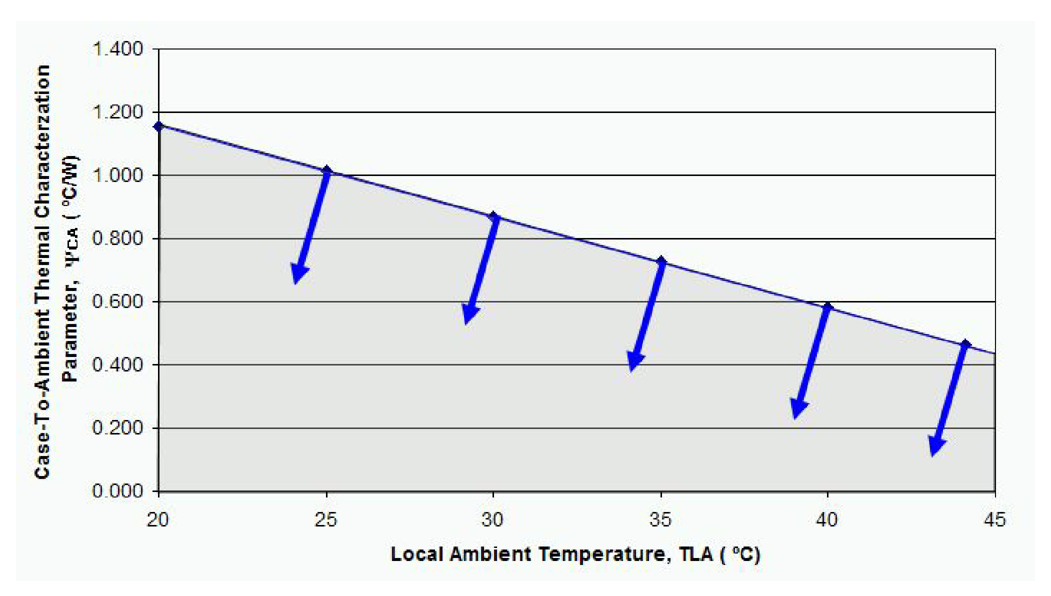

5.2 PICMG 1.3 Form Factor

Thermal solution design for the Intel Celeron processor 440 in the PICMG 1.3 form factor is very challenging. Due to limited volume for the heatsink (mainly in direction of heatsink height) and the available amount of airflow, system designers may have to make some tradeoffs in the system boundary condition requirements (i.e. maximum T_{LA}, acoustic requirements, etc.) in order to meet the processor's thermal requirements. The entire thermal solution from the heatsink design, chassis configuration and airflow source needs to be optimized for server systems in order to obtain the best performing solution.

Intel has worked with a third party vendor to enable a heatsink design for the Intel Celeron processor 440 for the PICMG 1.3 form factor. This design was optimized for the PICMG 1.3 form factor within the available volume for the thermal solution. The motherboard component keep-ins can be seen in Figure 47, "1U or PICMG 1.3 Motherboard Keep-out Footprint Definition and Height Restrictions for Enabling Components, Primary Side" on page 66 and Figure 48, "1U or PICMG 1.3 Motherboard Keep-out, Secondary Side" on page 67.

This solution requires 100% of the airflow to be ducted through the heatsink fins in order to prevent heatsink bypass. It is a copper base and copper fin heatsink that is attached to the motherboard with the use of a backplate. This solution is shown in Figure 11 on page 33.

Figure 11. 1U or PICMG 1.3 Copper Heatsink

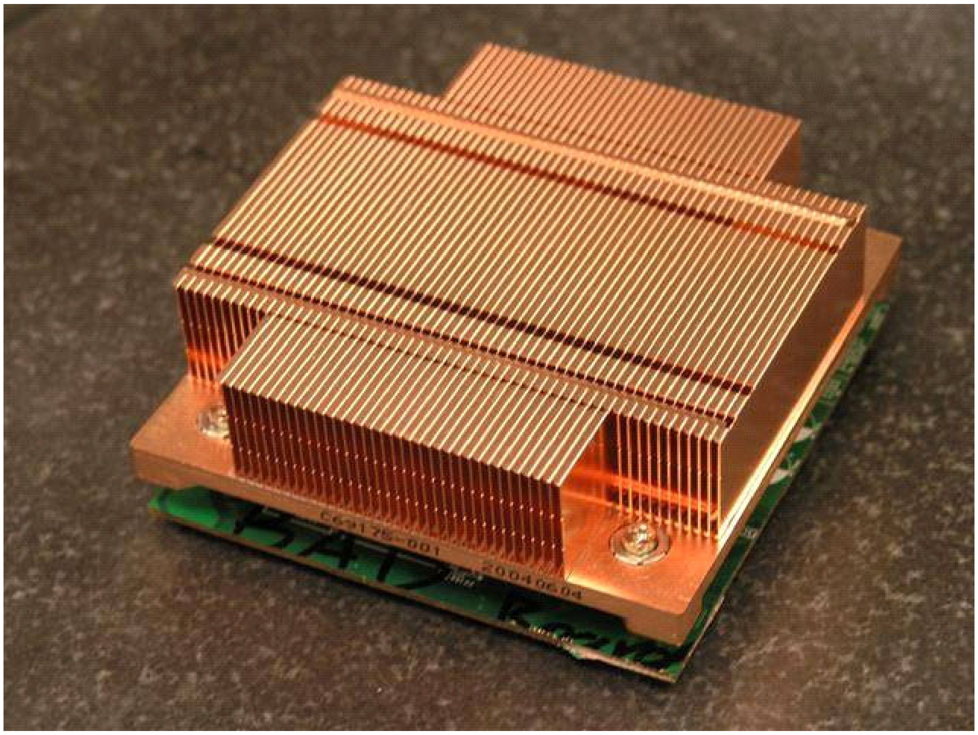

Based on lab test data, the case-to-ambient (Ψ_{CA}) performance of heatsink was found to be 0.356 °C/W with 18 CFM of airflow through the heatsink fins. The estimated performance for additional airflows is shown in Figure 12 on page 34.

Figure 12. 1U or PICMG 1.3 Heatsink Performance

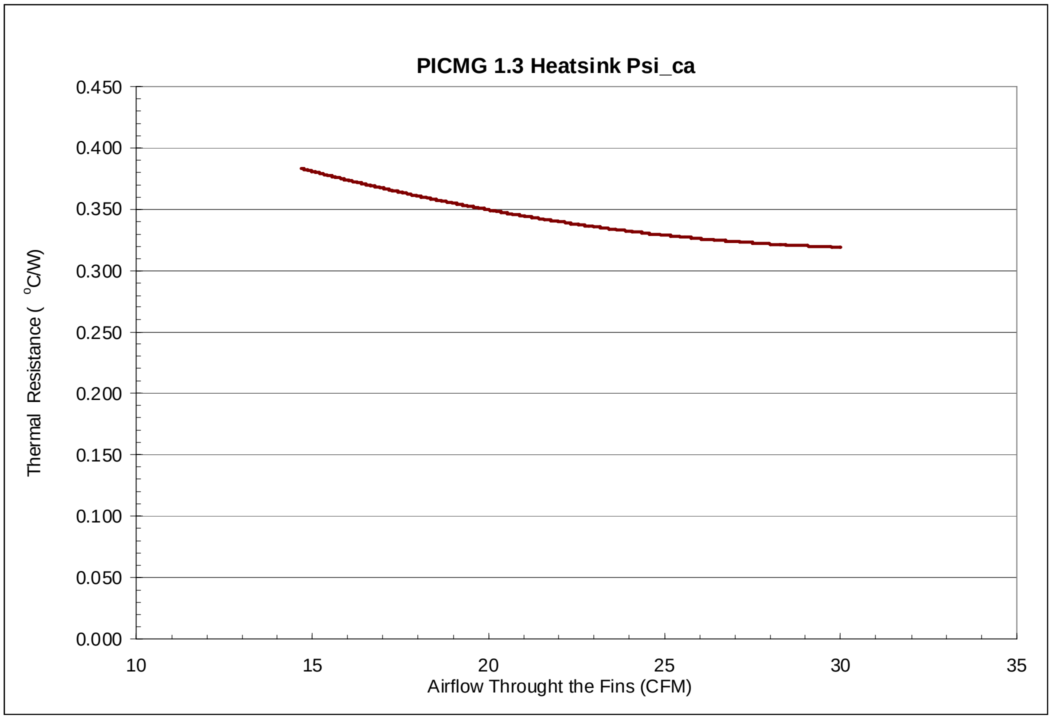

The performance of the heatsink could improve with more airflow, however the final intended thermal solution including, heatsink, airflow source, TIM, and attach mechanism must be validated by system integrators.

Developers who wish to design thermal solutions for the Processor, need to ensure that it meets the processor thermal specifications as stated in the processor EMTS and follow the recommended motherboard component keep-out as shown in Figure 47, "1U or PICMG 1.3 Motherboard Keep-out Footprint Definition and Height Restrictions for Enabling Components, Primary Side" on page 66. This keep-out will ensure that the processor thermal solution will not interfere with the voltage regulator components. In addition to this, a thermal solution design must meet the maximum component heights as specified by the PICMG 1.3 (http://picmg.org/specifications.stm). It should be noted that due to the vertical orientation of the heatsink there might be some stresses in the board due to the heatsink weight.

5.3 ATX/BTX form factors

For information regarding the Intel Thermal/Mechanical Reference Design thermal solution and design criteria for the ATX and BTX form factors, refer to the *Intel® Core™2 Duo Desktop Processor, Intel® Pentium® Dual-Core Processor and Intel® Pentium® 4 Processor 6x1 Sequence Thermal and Mechanical Design Guidelines (TMDG)*.

5.4 Embedded ATX form factor

Intel has worked with a third party vendor to enable for using Intel® Celeron® Processor 440 processor in Embedded ATX form factor. It is an active thermal solution made with Al6063 and copper as center core for improving heat transfer. Because of the lower profile of the active solution made for Embedded ATX form factor, it can fit for most 1U height constraint form factor too. Please refer to Embedded ATX form factor specification for form factor details. The active fan used for this active solution is a 70mm x 70mm DC Axial Fan at 15mm thickness. Air will be flow towards motherboard direction. See Figure 13 for details.

Based on lab test, the reference solution is capable of achieving a thermal resistance of 0.32C/W at air intake temperature of 49C, as tested on a 35W thermal test vehicle with full fan speed.

Figure 13. Embedded ATX Thermal Solution

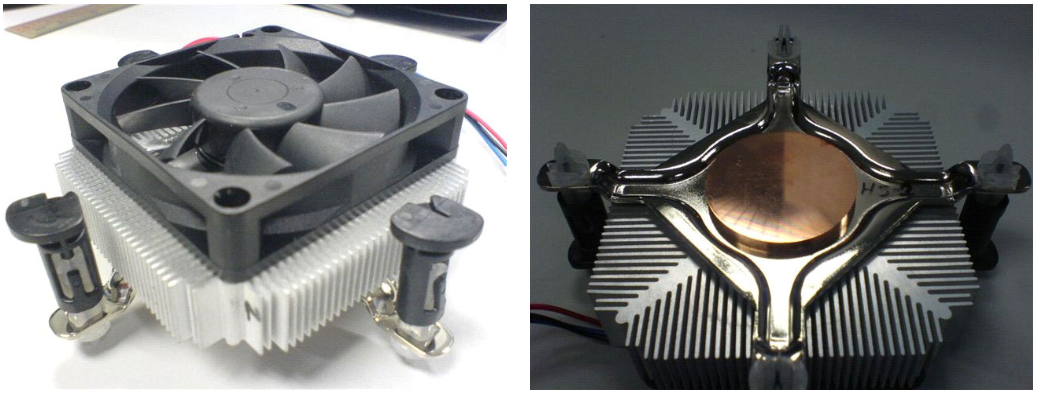

5.4.1 Active Fan Wire Configuration

Below is the wire configuration of the embedded ATX reference solution.

Table 5. Fan Wire Configuration Embedded ATX Solution

Pin	Function
Pin1 (Black Wire)	-
Pin 2 (Red Wire)	+
Pin 3 (Blue Wire	F00

5.4.2 Active Fan Curve

Below is curve of the fan pressure and volumetric air volume.

Figure 14. Fan P-Q Curve for Embedded ATX Solution

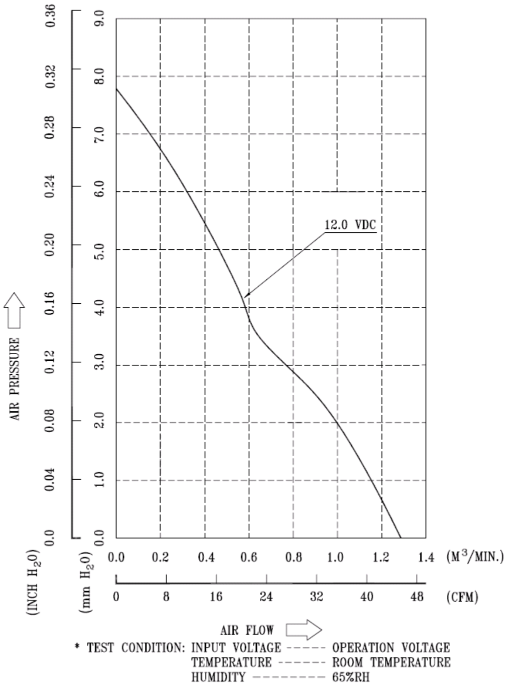

5.5 Altitude

The reference heatsink solutions will be evaluated at sea level. However, many companies design products that must function reliably at high altitude, typically 1,500 m [5,000 ft.] or more. Air-cooled temperature calculations and measurements at sea level must be adjusted to take into account altitude effects like variation in air density and overall heat capacity. This often leads to some degradation in thermal solution performance compared to what is obtained at sea level, with lower fan performance and higher surface temperatures. The system designer needs to account for altitude effects in the overall system thermal design to make sure that the T_C requirement for the processor is met at the targeted altitude.

5.6 Geometric Envelope for Intel Reference PICMG 1.3 Thermal Mechanical Design

Figure 48, “1U or PICMG 1.3 Motherboard Keep-out, Secondary Side” on page 67in Appendix E gives detailed reference PICMG 1.3 motherboard keep-out information for the reference thermal/mechanical enabling design. These drawings include height restrictions in the enabling component region.

6.0 Intel® Quiet System Technology (Intel® QST)

In the Intel® 965 Express family chipset, a new control algorithm for fan speed control is being introduced. It is composed of a Manageability Engine (ME) in the Graphics Memory Controller Hub (GMCH) which executes the Intel QST algorithm and the ICH8 containing the sensor bus and fan control circuits.

The ME provides integrated fan speed control in lieu of the mechanisms available in a SIO or a stand-alone ASIC. The Intel QST is time based as compared to the linear or state control used by the current generation of FSC devices.

A short discussion of Intel QST will follow along with thermal solution design recommendations. For a complete discussion of programming the Intel QST in the ME, consult the *Intel® Quiet System Technology (Intel® QST) Configuration and Tuning Manual*.

Note: Fan speed control algorithms and Intel QST in particular rely on a thermal solution being compliant to the processor thermal profile. It is unlikely that any fan speed control algorithm can compensate for a non-compliant thermal solution. See Chapter 5 and Chapter 6 for thermal solution requirements that should be met before evaluating or configuring a system with Intel QST.

6.1 Intel® Quiet System Technology (Intel® QST)

The objective of Intel QST is to minimize the system acoustics by more closely controlling the thermal sensors to the corresponding processor or chipset device $T_{CONTROL}$ value. This is achieved by the use of a Proportional-Integral-Derivative (PID) control algorithm and a Fan Output Weighting Matrix. The PID algorithm takes into account the difference between the current temperature and the target ($T_{CONTROL}$), the rate of change and direction of change to minimize the required fan speed change. The Fan Output Weighting Matrix uses the effects of each fan on a thermal sensor to minimize the required fan speed changes.

Figure 15 on page 37 shows in a very simple manner how Intel QST works. See the *Intel® Quiet System Technology (Intel® QST) Configuration and Tuning Manual* for a detail discussion of the inputs and response.

Figure 15. Intel® QST Overview

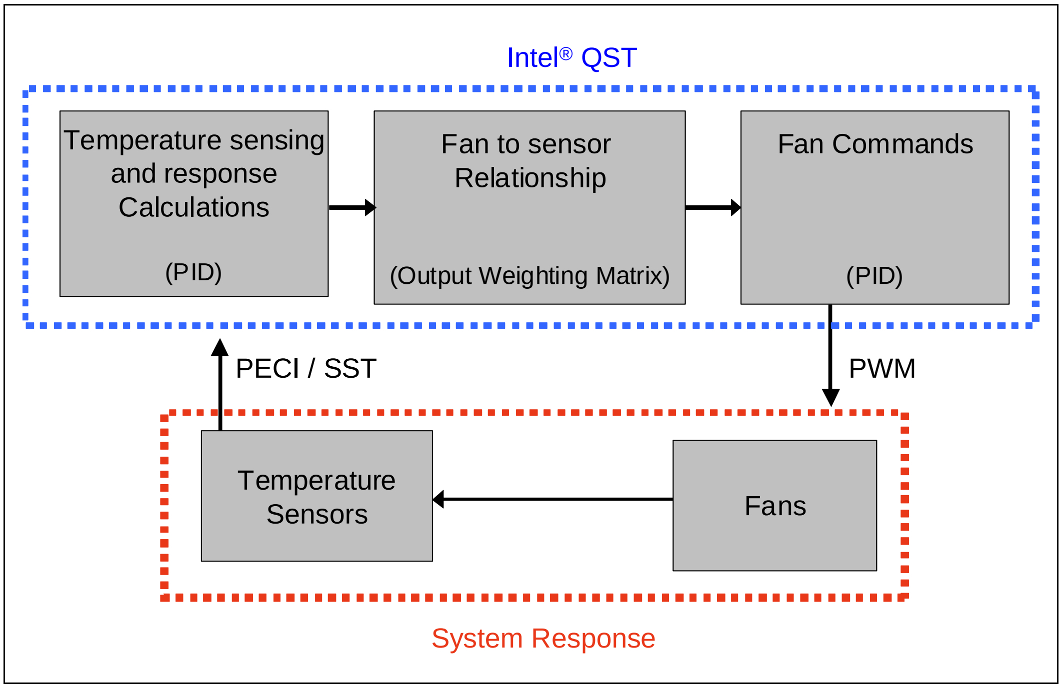

6.1.1 Output Weighting Matrix

Intel QST provides an Output Weighting Matrix that provides a means for a single thermal sensor to affect the speed of multiple fans. An example of how the matrix could be used is if a sensor placed next to the memory is sensitive to changes in both the processor heatsink fan and a 2nd fan in the system. By placing a sensor in this matrix, the Intel QST could command both the processor thermal solution fan and this 2nd fan to accelerate by a small amount. At the system level, these two small changes would result in a smaller change in acoustics rather than having a single fan respond to this sensor.

6.1.2 Proportional-Integral-Derivative (PID)

The use of Proportional-Integral-Derivative (PID) control algorithms allow the magnitude of fan response to be determined based upon the difference between current temperature readings and specific temperature targets. A major advantage of a PID Algorithm is the ability to control the fans to achieve sensor temperatures much closer to the $T_{CONTROL}$.

Figure 16 on page 38 is an illustration of the PID fan control algorithm. As illustrated in the figure, when the actual temperature is below the target temperature, the fan will slow down. The current FSC devices have a fixed temperature vs. PWM output relationship and miss this opportunity to achieve additional acoustic benefits. As the actual temperature starts ramping up and approaches the target temperature, the algorithm will instruct the fan to speed up gradually, but will not abruptly increase the fan speed to respond to the condition. It can allow an overshoot over the target temperature for a short period of time while ramping up the fan to bring the actual temperature to the target temperature. As a result of its operation, the PID control algorithm can enable an acoustic-friendly platform.

Figure 16. PID Controller Fundamentals

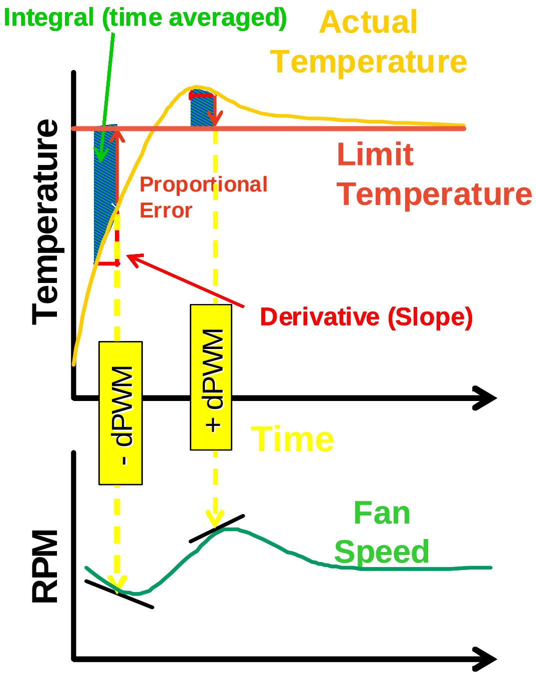

For a PID algorithm to work, limit temperatures are assigned for each temperature sensor. For Intel QST, the $T_{CONTROL}$ for the processor and chipset are to be used as the limit temperature. The ME will measure the error, slope and rate of change using the equations below:

- Proportional Error (P) = T_{LIMIT} - T_{ACTUAL}
- Integral (I) = Time averaged error
- Derivative (D) = ΔTemp / ΔTime

Three gain values are used to control response of algorithm.

- Kp = proportional gain
- Ki = Integral gain
- Kd = derivative gain

The *Intel® Quiet System Technology (Intel® QST) Configuration and Tuning Manual* provides initial values for the each of the gain constants. In addition, it provides a methodology to tune these gain values based on system response. Finally, the fan speed change will be calculated using the following formula:

ΔPWM = -Px(Kp) - Ix(Ki) + Dx(Kd)

6.2 Board and System Implementation of Intel® QST

To implement, the board must be configured as shown in Figure 17 with its components listed below:

- ME system (S0-S1) with Controller Link connected and powered
- DRAM with Channel A DIMM 0 installed and 2MB reserved for Intel QST FW

execution

- SPI Flash with sufficient space for the Intel QST Firmware
- SST-based thermal sensors to provide board thermal data for Intel QST algorithms
- Intel QST firmware

Figure 17. Intel® QST Platform Requirements

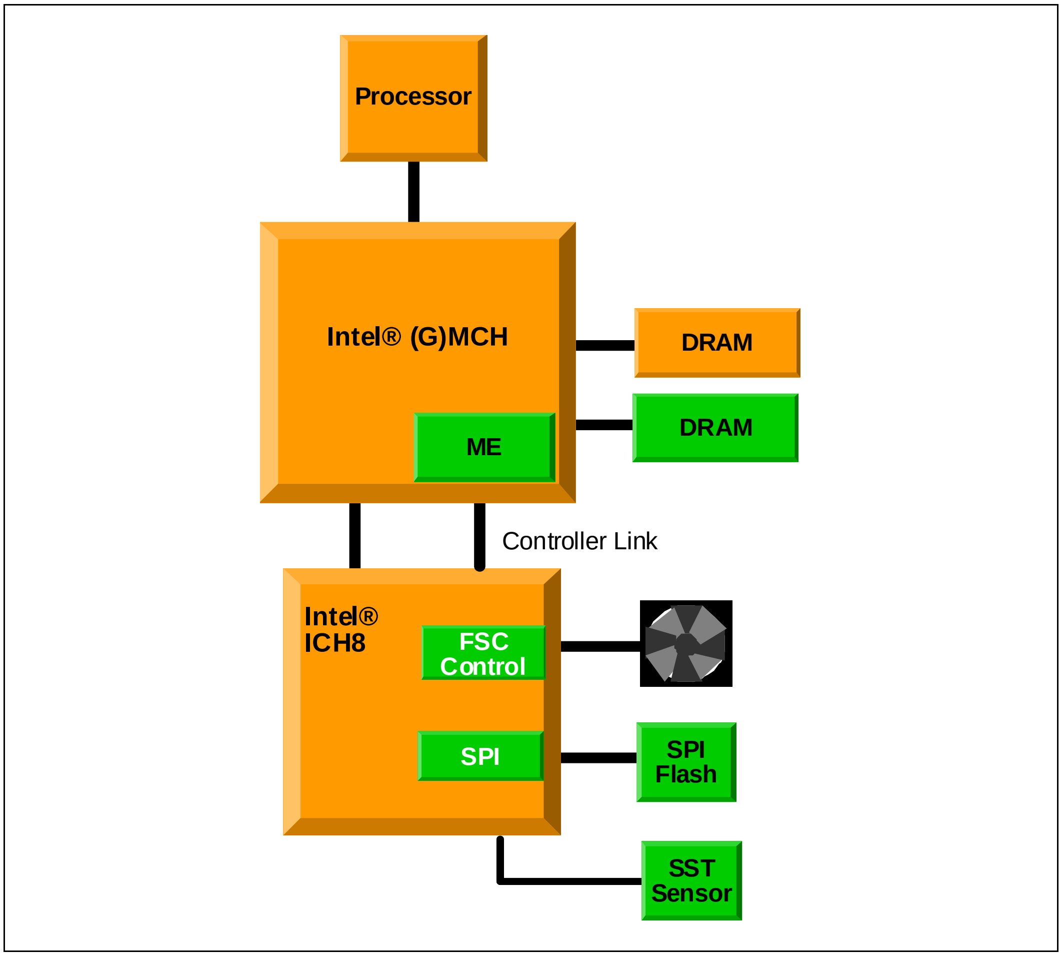

Note: Simple Serial Transport (SST) is a single wire bus that is included in the ICH8 to provide additional thermal and voltage sensing capability to the Manageability Engine (ME)

Figure 18 on page 41 shows the major connections for a typical implementation that can support processors with a Digital thermal sensor or a thermal diode. In this configuration, a SST Thermal Sensor has been added to read the on-die thermal diode that is in all of the processors in the 775-land LGA packages shipped before the Intel Celeron processor 440. With the proper configuration information, the ME can accommodate inputs from PECI or SST for the processor socket. Additional SST sensors can be added to monitor the system temperature (see Appendix F for BTX recommendations for placement).

Figure 18. Example Acoustic Fan Speed Control Implementation

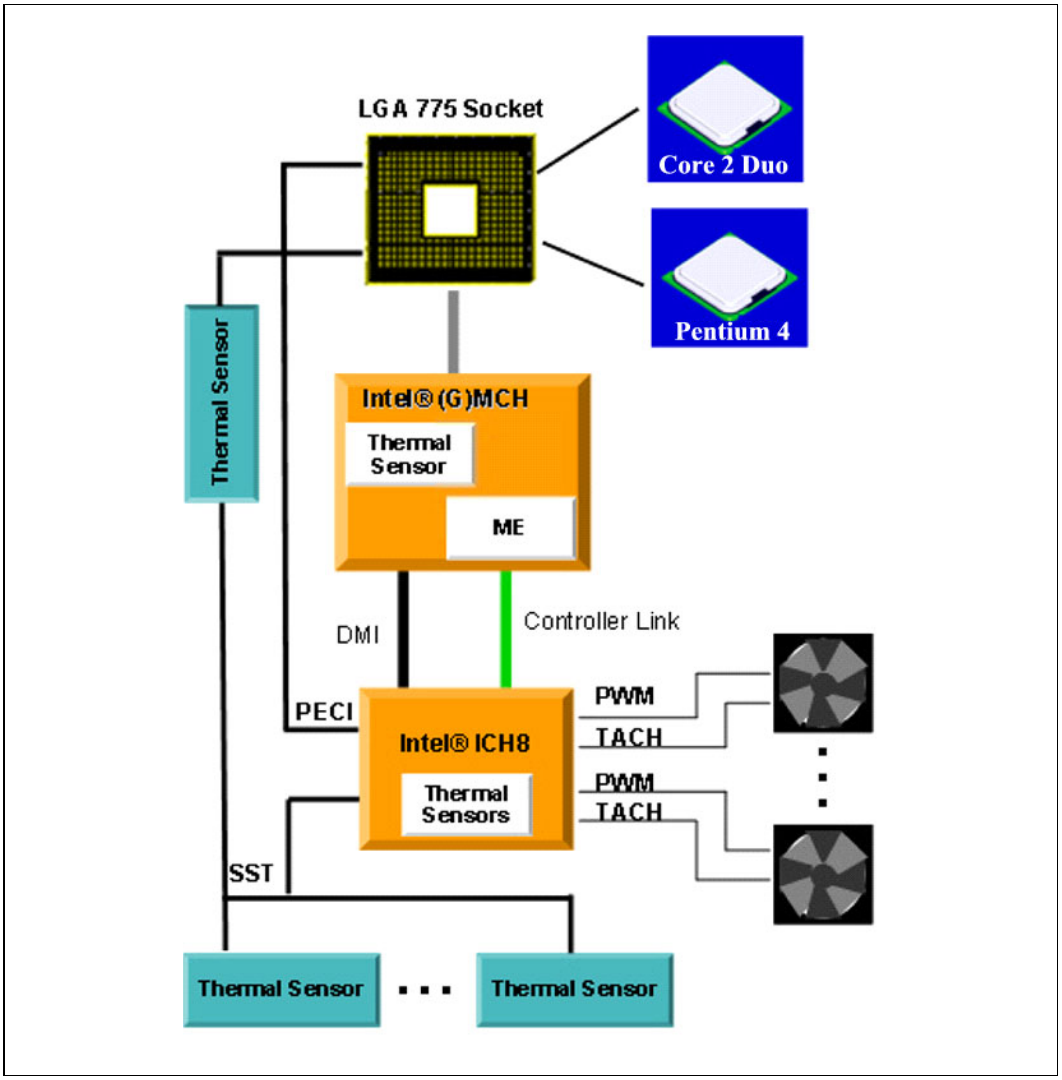

Intel has engaged with a number of major manufacturers of thermal / voltage sensors to provide devices for the SST bus. Contact your Intel Field Sales representative for the current list of manufacturers and visit their web sites or local sales representatives for a part suitable for your design.

6.3 Intel® QST Configuration and Tuning

The board manufacturer is responsible for the initial configuration of the Intel QST. The SPI flash should be programmed with the hardware configuration of the motherboard and initial settings for fan control, fan monitoring and voltage and thermal monitoring. This initial data is generated using the Intel provided Configuration Tool.

At the system integrator level, the Configuration Tool can be used to tune the Intel QST subsystem to reflect the shipping system configuration. In the tuning process, the Intel QST can be modified to have a proper association between the installed fans and sensors in the shipping system. A Weighting Matrix Utility and Intel QST Log program assist in optimizing fan management and to achieve an acoustic goal.

See your Intel field sales representative for availability of these tools.

6.4 Fan Hub Thermistor and Intel® QST

There is no closed loop control between Intel QST and the thermistor, but they can work in tandem to provide the maximum fan speed reduction. The BTX reference design includes a thermistor on the fan hub. This Variable Speed Fan curve as shown in Figure 19 will determine the maximum fan speed as a function of the inlet ambient temperature, and by design, provides a Ψ_{CA} sufficient to meet the thermal profile of the processor. Intel QST, by measuring the processor Digital thermal sensor, will command the fan to reduce speed below the VSF curve in response to processor workload. Conversely, if the processor workload increases, the FSC will command the fan via the PWM duty cycle to accelerate the fan up to the limit imposed by the VSF curve. Care needs to be taken in BTX designs to ensure the fan speed at the minimum operating speed provides sufficient air flow to support the other system components.

Figure 19. Digital Thermal Sensor and Thermistor

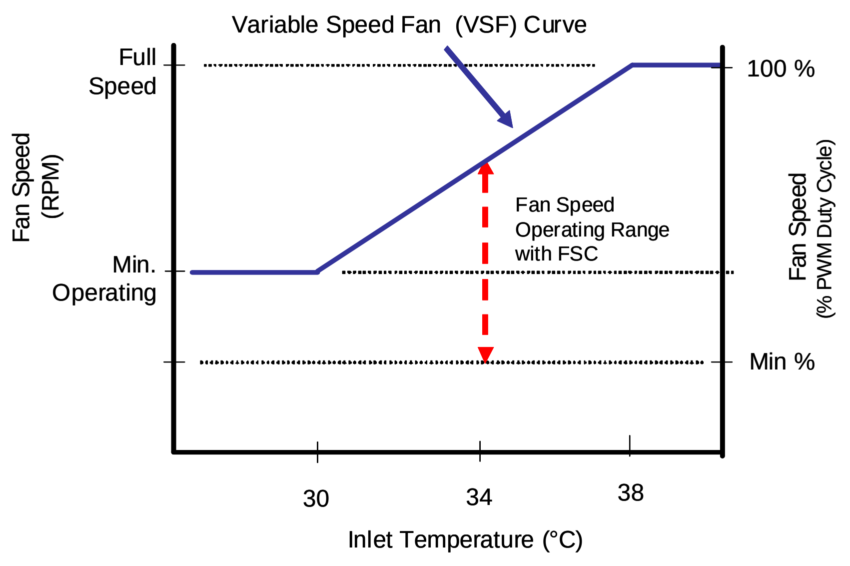

Appendix A LGA775 Socket Heatsink Loading

A.1 LGA775 Socket Heatsink Considerations

Heatsink clip load is traditionally used for:

- Mechanical performance in mechanical shock and vibration
- Thermal interface performance
  - Required preload depends on TIM

— Preload can be low for thermal grease

In addition to mechanical performance in shock and vibration and TIM performance, LGA775 socket requires a minimum heatsink preload to protect against fatigue failure of socket solder joints.

Solder ball tensile stress is originally created when, after inserting a processor into the socket, the LGA775 socket load plate is actuated. In addition, solder joint shear stress is caused by a coefficient of the thermal expansion (CTE) mismatch induced by shear loading. The solder joint compressive axial force (Faxial) induced by the heatsink preload helps to reduce the combined joint tensile and shear stress.

Overall, the heatsink required preload is the minimum preload needed to meet all of the above requirements: Mechanical shock and vibration and TIM performance and LGA775 socket protection against fatigue failure.

A.2 Metric for Heatsink Preload for Designs Non-Compliant with Intel Reference Design

A.2.1 Heatsink Preload Requirement Limitations

Heatsink preload by itself is not an appropriate metric for solder joint force across various mechanical designs and does not take into account, for example (not an exhaustive list):

- Heatsink mounting hole span
- Heatsink clip/fastener assembly stiffness and creep
- Board stiffness and creep
- Board stiffness is modified by fixtures like backing plate, chassis attach, etc.

A simulation shows that the solder joint force (F_{axial}) is proportional to the board deflection measured along the socket diagonally. The matching of Faxial required to protect the LGA775 socket solder joint in temperature cycling is equivalent to matching a target MB deflection.

Therefore, the heatsink preload for LGA775 socket solder joint protection against fatigue failure can be more generally defined as the load required to create a target board downward deflection throughout the life of the product.

This board deflection metric provides guidance for mechanical designs that differ from the reference design for the ATX//µATX form factor.

A.2.2 Motherboard Deflection Metric Definition

Motherboard deflection is measured along either diagonal (refer to Figure 20 below):

d = dmax – (d1 + d2)/2

d’ = dmax – (d’1 + d’2)/2

Configurations in which the deflection is measured are defined in the Table 6, “Board Deflection Configuration Definitions” below.

To measure board deflection, follow industry standard procedures (such as IPC) for board deflection measurement. Height gauges and possibly dial gauges may also be used.

Table 6. Board Deflection Configuration Definitions

Configuration Parameter	Processor + Socket load plate	Heatsink	Parameter Name
d_ref	yes	no	BOL deflection, no preload
d_BOL	yes	yes	BOL deflection with preload
d_EOL	yes	yes	EOL deflection

BOL: Beginning of Life

EOL: End of Life

Figure 20. Board Deflection Definition

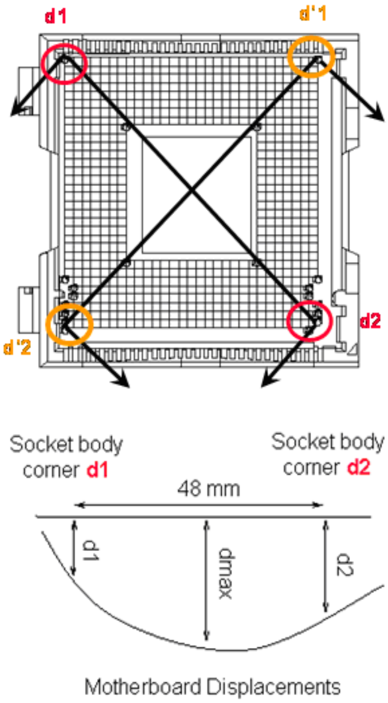

A.2.3 Board Deflection Limits

Deflection limits for the ATX/µATX form factor are:

d_BOL - d_ref 0.09 mm and d_EOL - d_ref 0.15 mm

And

d'_BOL – d'_ref 0.09 mm and d_EOL' – d_ref' 0.15 mm

Notes:

1. The heatsink preload must remain within the static load limits defined in the processor EMTS at all times.
2. Board deflection should not exceed motherboard manufacturer specifications.

This section is for illustration only, and relies on the following assumptions:

- 72 mm x 72 mm hole pattern of the reference design
- Board stiffness = 900 lb./in. at BOL, with degradation that simulates board creep over time
    - — Though these values are representative, they may change with selected material and board manufacturing process. Check with your motherboard vendor.
- Clip stiffness assumed constant – No creep.

Using Figure 21 on page 45, the heatsink preload at beginning of life is defined to comply with d_EOL – d_ref = 0.15mm depending on the clip stiffness assumption.

Note that the BOL and EOL preload and board deflection differ. This is a result of the creep phenomenon. The example accounts for the creep expected to occur in the motherboard. It assumes no creep to occur in the clip. However, there is a small amount of creep accounted for in the plastic fasteners. This situation is somewhat similar to the Intel Reference Design.

The impact of the creep to the board deflection is a function of the clip stiffness:

- The relatively compliant clips store strain energy in the clip under the BOL preload condition and tend to generate increasing amounts of board deflection as the motherboard creeps under exposure to time and temperature.
- In contrast, the stiffer clips stores very little strain energy, and therefore, do not generate substantial additional board deflection through life.

Notes:

1. Board and clip creep modify board deflection over time and depends on board stiffness, clip stiffness and the selected materials.
2. Designers must define the BOL board deflection that will lead to the correct end of life board deflection.

Figure 21. Example: Defining Heatsink Preload Meeting Board Deflection Limit

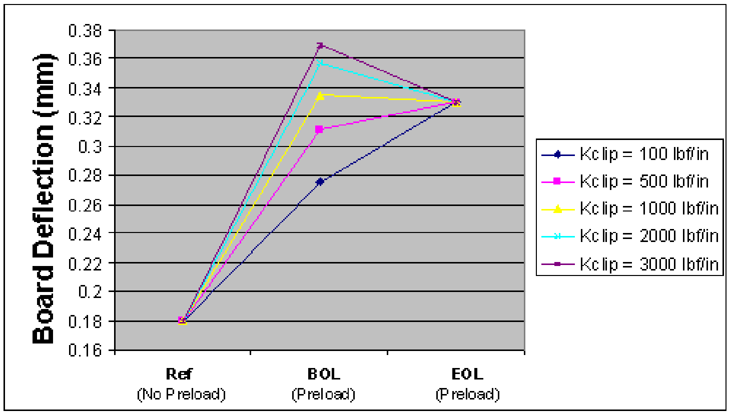

Kclip (lbf/in)	BOL Preload (lb)	(d_BOL - d_ref) (mm)	EOL Preload (lb)	(d_EOL - d_ref) (mm)
100	21.1	0.09	18.0	0.15
500	29.2	0.13	18.1	0.15
1000	34.4	0.15	18.0	0.15
2000	39.4	0.18	18.0	0.15
3000	42	0.19	18.0	0.15

A.2.4 Additional Considerations

Intel recommends to design to {d_BOL - d_ref = 0.15mm} at BOL when EOL conditions are not known or difficult to assess.

The following information is given for illustration only. It is based on the reference keep-out, assuming there is no fixture that changes board stiffness:

d_ref is expected to be 0.18 mm on average, and be as high as 0.22 mm

As a result, the board should be able to deflect 0.37 mm minimum at BOL.

Additional deflection as high as 0.09 mm may be necessary to account for additional creep effects impacting the board/clip/fastener assembly. As a result, designs could see as much as 0.50mm total downward board deflection under the socket.

In addition to board deflection, other elements need to be considered to define the space needed for the downward board total displacement under load, like the potential interference of through-hole mount component pin tails of the board with a mechanical fixture on the back of the board.

Notes:

1. The heatsink preload must remain below the maximum load limit of the package at all times (Refer to processor EMTS).
2. Board deflection should not exceed motherboard manufacturer specifications.

A.2.4.1 Motherboard Stiffening Considerations

To protect LGA775 socket solder joint, designers need to drive their mechanical design to:

- Allow downward board deflection to put the socket balls in a desirable force state to protect against fatigue failure of socket solder joint (refer to Chapter A.2.1, “Heatsink Preload Requirement Limitations,” Appendix A.2.2, “Motherboard Deflection Metric Definition,” and Appendix A.2.3, “Board Deflection Limits”).
- Prevent board upward bending during mechanical shock event.
- Define load paths that keep the dynamic load applied to the package within specifications published in the processor EMTS.

Limiting board deflection may be appropriate in some situations like:

- Board bending during shock
- Board creep with high heatsink preload

However, the load required to meet the board deflection recommendation (refer to Appendix A.2.3, “Board Deflection Limits”) with a very stiff board may lead to heatsink preloads exceeding package maximum load specification. For example, such a situation may occur when using a backing plate that is flush with the board in the socket area, and prevents the board to bend underneath the socket.

A.3 Heatsink Selection Guidelines

Carefully evaluate heatsinks coming with motherboard stiffening devices (like backing plates), and conduct board deflection assessments based on the board deflection metric.

Solutions derived from the reference design comply with the reference heatsink preload, for example:

- The Boxed Processor
- The Intel® RCFH-4 reference design available from licensed suppliers (refer to Appendix F, "Intel Enabled Reference Solution Information" for contact information)

Intel will collaborate with vendors participating in its third party test house program to evaluate third party solutions. Vendor information will be available after product launch.

Appendix B Thermal Interface Management

To optimize a heatsink design, it is important to understand the impact of factors related to the interface between the processor and the heatsink base. Specifically, the bond line thickness, interface material area and interface material thermal conductivity should be managed to realize the most effective thermal solution.

B.1 Bond Line Management

Any gap between the processor integrated heat spreader (IHS) and the heatsink base degrades thermal solution performance. The larger the gap between the two surfaces, the greater the thermal resistance. The thickness of the gap is determined by the flatness and roughness of both the heatsink base and the integrated heat spreader, plus the thickness of the thermal interface material (for example thermal grease) used between these two surfaces and the clamping force applied by the heatsink attach clip(s).

B.2 Interface Material Area

The size of the contact area between the processor and the heatsink base will impact the thermal resistance. However, there is a point of diminishing returns. Unrestrained incremental increases in thermal interface material area do not translate to a measurable improvement in thermal performance.

B.3 Interface Material Performance

Two factors impact the performance of the interface material between the processor and the heatsink base:

- Thermal resistance of the material
- Wetting/filling characteristics of the material

Thermal resistance is a description of the ability of the thermal interface material to transfer heat from one surface to another. The higher the thermal resistance, the less efficient the interface material is at transferring heat. The thermal resistance of the interface material has a significant impact on the thermal performance of the overall thermal solution. The higher the thermal resistance, the larger the temperature drop is across the interface and the more efficient the thermal solution (heatsink, fan) must be to achieve the desired cooling.

The wetting or filling characteristic of the thermal interface material is its ability, under the load applied by the heatsink retention mechanism, to spread and fill the gap between the processor and the heatsink. Since air is an extremely poor thermal conductor, the more completely the interface material fills the gaps, the lower the temperature drops across the interface. In this case, thermal interface material area also becomes significant; the larger the desired thermal interface material area, the higher the force required to spread the thermal interface material.

Appendix C Case Temperature Reference Metrology

C.1 Objective and Scope

This appendix defines a reference procedure for attaching a thermocouple to the IHS of a 775-land LGA package for T_C measurement. This procedure takes into account the specific features of the 775-land LGA package and of the LGA775 socket for which it is intended. This procedure is applicable for both Thermal Test Vehicles (TTV) and functional processors. The recommended equipment for the reference thermocouple installation, including tools and part numbers are also provided. In addition a video *Thermocouple Attach Using Solder – Video CD-ROM* is available and shows the process in real time.

Note: This procedure is applicable only to processors that comply with the 2005 Performance or Mainstream / Value FMB. The T_C measurements from this metrology are not compatible with processor specifications based on the 2004 FMB targets.

C.2 Supporting Test Equipment

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

Table 7. Recommended Equipment

Item	Description	Part Number
Measurement and Output		
Microscope	Olympus* Light microscope or equivalent	SZ-40
DMM	Digital Multi Meter for resistance measurement	Fluke 79 Series
Thermal Meter	Hand held thermocouple meter	Multiple Vendors
Solder Station (see note 1 for ordering information)		
Heater Block	Heater assembly to reflow solder on IHS	30330
Heater	WATLOW120V 150W Firerod	0212G G1A38-L12
Transformer	Superior Powerstat transformer	05F857
Miscellaneous Hardware		
Solder	Indium Corp. of America Alloy 57BI / 42SN / 1AG 0.010 Diameter	52124
Flux	Indium Corp. of America	5RMA
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	Not Available
Thermocouple	Omega*, 36 gauge, "T" Type (see note 2 for ordering information)	OSK2K1280/5SRTC-TT-T-36-72
Calibration and Control		
Ice Point Cell	Omega*, stable 0 ºC temperature source for calibration and offset	TRCIII
Hot Point Cell	Omega*, temperature source to control and understand meter slope gain	CL950-A-110

Notes:

1. A Solder Station consisting of the Heater Block, Heater, Press and Transformer is available from Jemelco Engineering (480-804-9514).

2. This part number is a custom part with the specified insulation trimming and packaging requirements necessary for a quality thermocouple attachment; see Figure 22 on page 49. Order from Omega Anthony Alvarez, direct phone: (203) 359-7671, direct fax: (203) 968-7142, email: aalvarez@omega.com

Figure 22. Omega Thermocouple

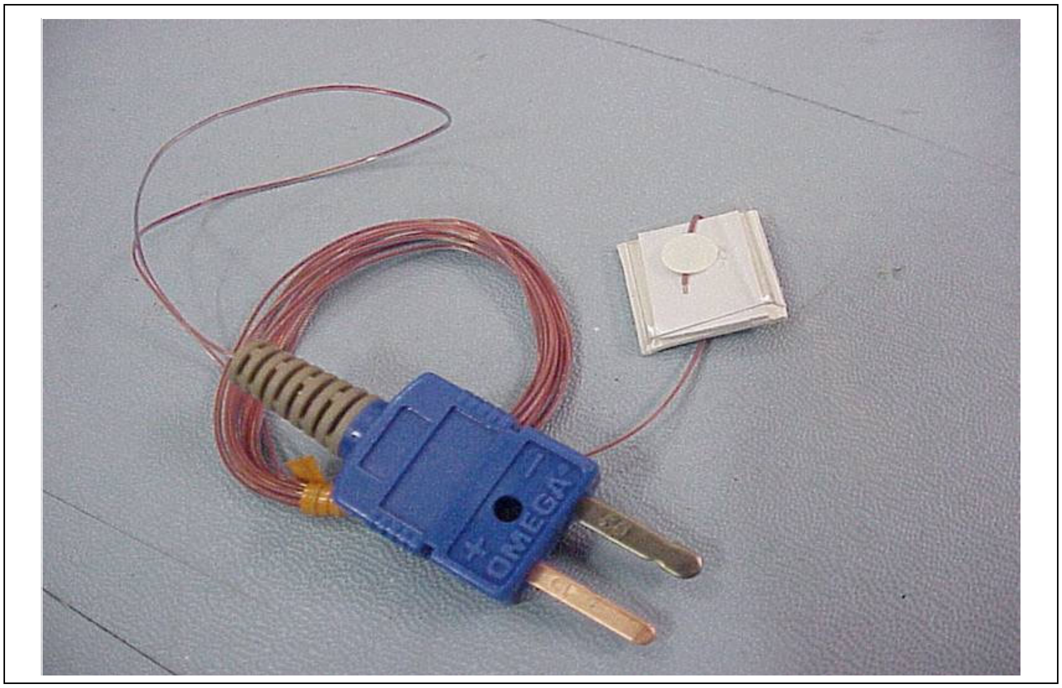

C.3 Thermal Calibration and Controls

It is recommended that full and routine calibration of temperature measurement equipment be performed before attempting to perform temperature case measurement of TTVs and live products. Intel recommends checking the meter probe set against known standards. This should be done at 0ºC (using ice bath or other stable temperature source) and at an elevated temperature, around 80ºC (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 the implementation of proper controls for thermal measurements.

Notes:

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. Please ask your Intel field sales representative if you need assistance to groove and/or install a thermocouple according to the reference process.

C.4 IHS Groove

Cut a groove in the package IHS according to the drawing given in Figure 23 on page 50.

Figure 23. 775-LAND LGA Package Reference Groove Drawing

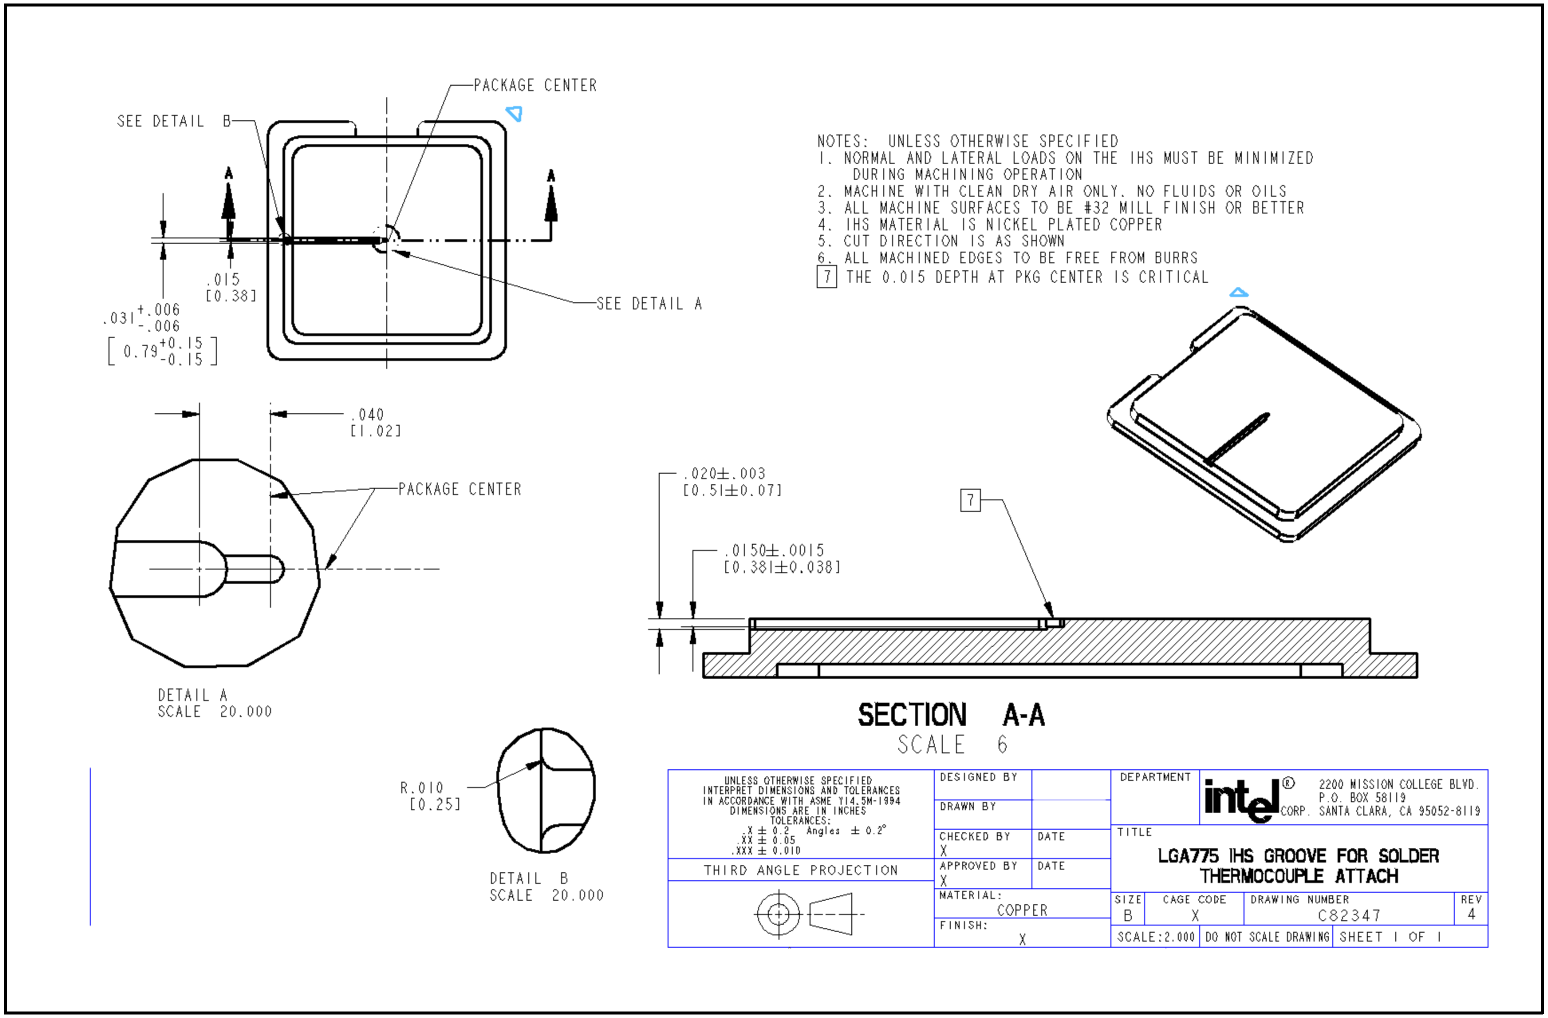

The orientation of the groove relative to the package pin 1 indicator (gold triangle in one corner of the package) is shown below. Figure 24 shows the 775-Land LGA package IHS.

Figure 24. IHS Groove on the 775-LAND LGA Package

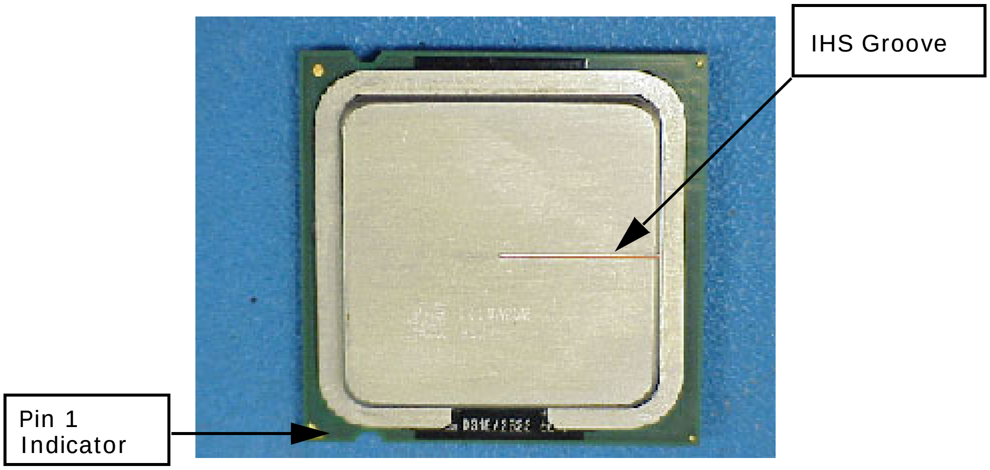

When the processor is installed in the LGA775 socket, the groove is perpendicular to the socket load lever, and on the opposite side of the lever, as shown Figure 25.

Figure 25. IHS Groove Orientation Relative to the LGA775 Socket

Select a machine shop that is capable of holding drawing specified tolerances. IHS groove 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.

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

C.5 Thermocouple Attach Procedure

The procedure to attach a thermocouple with solder takes about 15 minutes to complete. Before proceeding, turn on the solder block heater, as it can take up to 30 minutes to reach the target temperature of 153 - 155°C.

Note: To avoid damage to the TTV or processor, ensure the heater temperature does not exceed 155°.

As a complement to the written procedure, a video, *Thermocouple Attach Using Solder – Video CD-ROM*, is available.

C.5.1 Thermocouple Conditioning and Preparation

1. Utilize a calibrated thermocouple as specified in Appendix C.2, "Supporting Test Equipment" and Appendix C.3, "Thermal Calibration and Controls".
2. Under a microscope, verify the thermocouple insulation meets the quality requirements. The insulation should be about 1/16 inch (0.62 ± 0.030) from the end of the bead (Figure 26).

Figure 26. Inspection of Insulation on Thermocouple

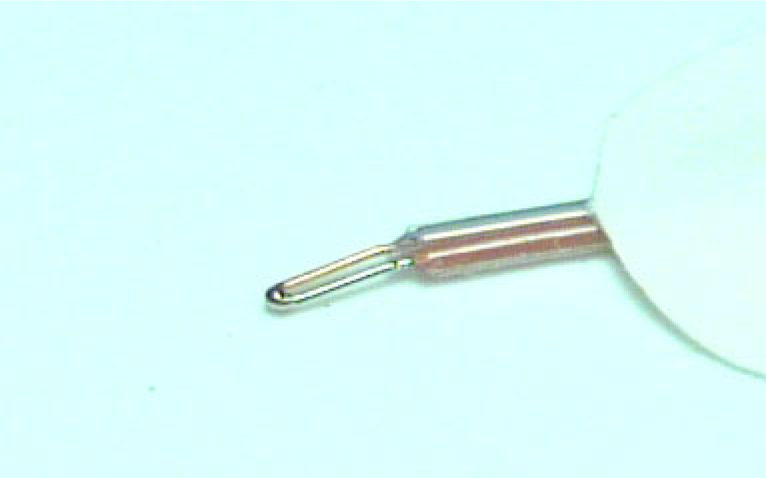

3. Measure the thermocouple resistance by holding both contacts on the connector on one probe and the tip of thermocouple to the other probe of the DMM (measurement should be about~3.0 ohms for 36-gauge type T thermocouple).
4. Straighten the wire for about 38 mm [1 ½ inch] from the bead.
5. Using the microscope and tweezers, bend the tip of the thermocouple at approximately 10 degree angle by about 0.8 mm [.030 inch] from the tip (Figure 28 on page 53).

C.5.2 Thermocouple Attachment to the IHS

6. Clean groove and IHS with Isopropyl Alcohol (IPA) and a lint free cloth removing all residues prior to thermocouple attachment.
7. Place the thermocouple wire inside the groove, letting the exposed wire and bead extend about 1.5 mm [0.030 inch] past the end of groove. Secure it with Kapton* tape (Figure 27). Clean the IHS with a swab and IPA.

8. Verify under the microscope that the thermocouple wires are straight and parallel in the groove and that the bead is still bent.

Figure 27. Securing Thermocouple Wires with Kapton* Tape Prior to Attach

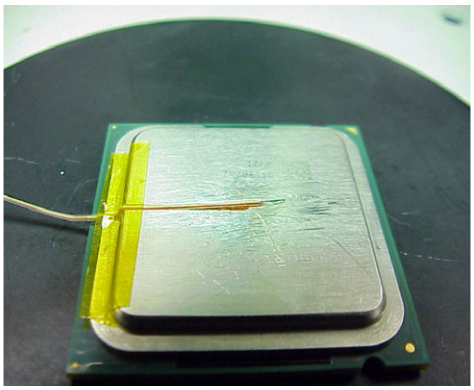

9. Lift the wire at the middle of groove with tweezers and bend the front of the wire to place the thermocouple in the groove ensuring the tip is in contact with the end and bottom of the groove in the IHS (Figure 28- A and B on page 53).

Figure 28. Thermocouple Bead Placement

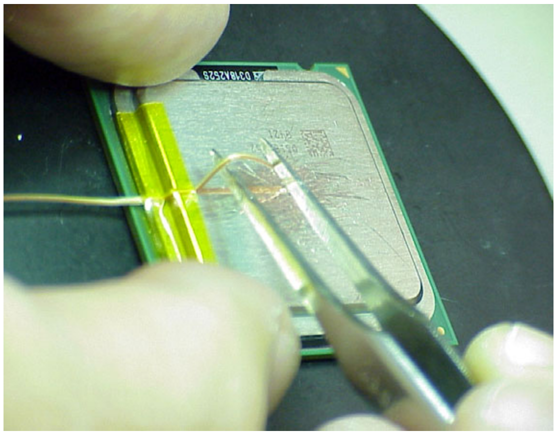

10. Place the package under the microscope to continue with the process. It is also recommended to use a fixture (like processor tray or a plate) to help hold the unit in place for the rest of the attach process.
11. While still at the microscope, press the wire down about 6mm [0.125”] from the thermocouple bead using the tweezers or your finger. Place a piece of Kapton* tape to hold the wire inside the groove (Figure 29 on page 54). Refer to Figure 30 on page 54 for detailed bead placement.

Figure 29. Position Bead on the Groove Step

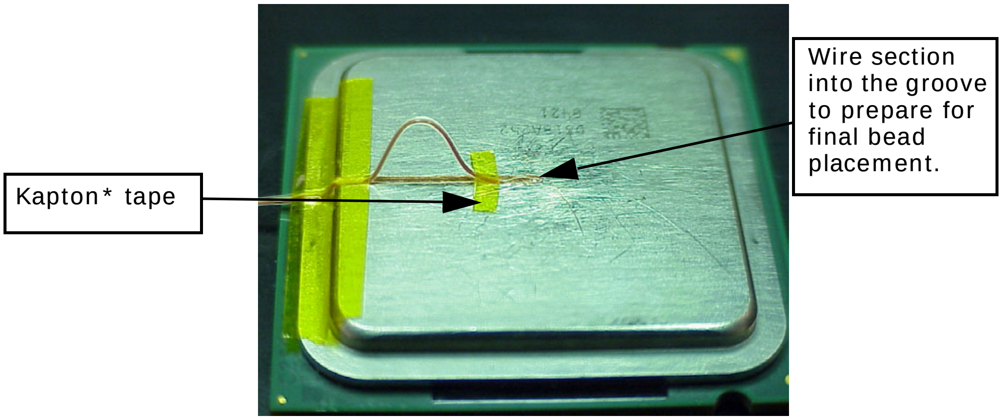

Figure 30. Detailed Thermocouple Bead Placement

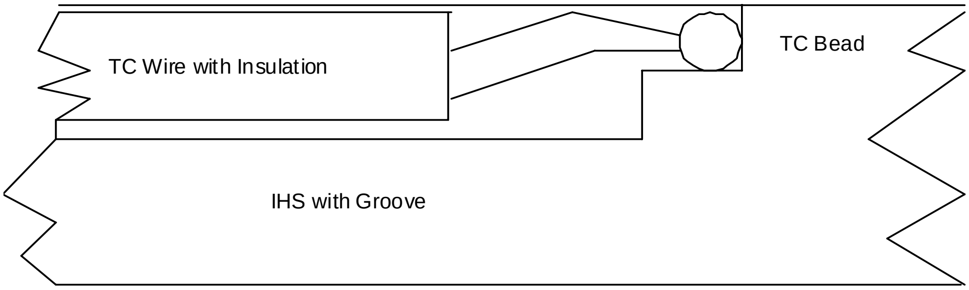

Figure 31. Third Tape Installation

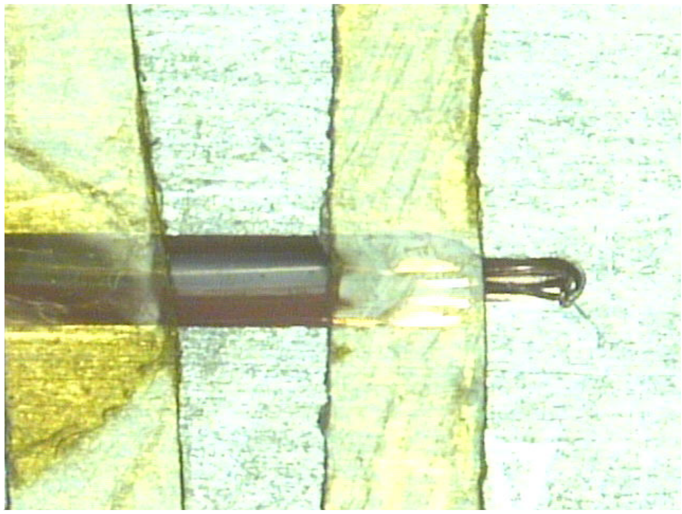

12. Place a 3rd piece of tape at the end of the step in the groove as shown in Figure 31 on page 54. This tape will create a solder dam to prevent solder from flowing into the larger IHS groove section during the melting process.
13. Measure resistance from thermocouple end wires (hold both wires to a DMM probe) to the IHS surface. This should be the same value (Figure 32) as measured during the thermocouple conditioning in Appendix C.5.1, "Thermocouple Conditioning and Preparation", step 3.

Figure 32. Measuring Resistance between Thermocouple and IHS

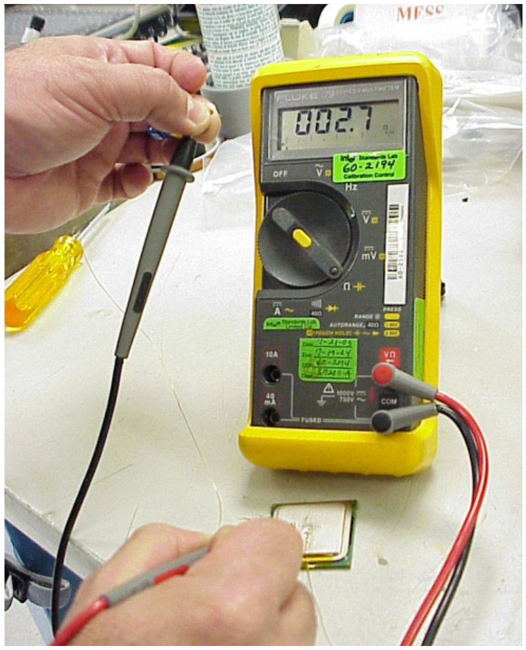

14. Using a fine point device, place a small amount of flux on the thermocouple bead. Be careful not to move the thermocouple bead during this step (Figure 33). Ensure the flux remains in the bead area only.

Figure 33. Applying Flux to the Thermocouple Bead

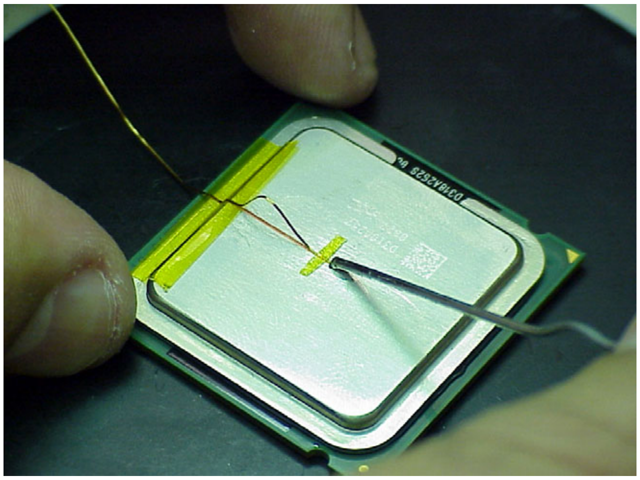

15. Cut two small pieces of solder 1/16 inch (0.065 inch / 1.5 mm) from the roll using tweezers to hold the solder while cutting with a fine blade (Figure 34).

Figure 34. Cutting Solder

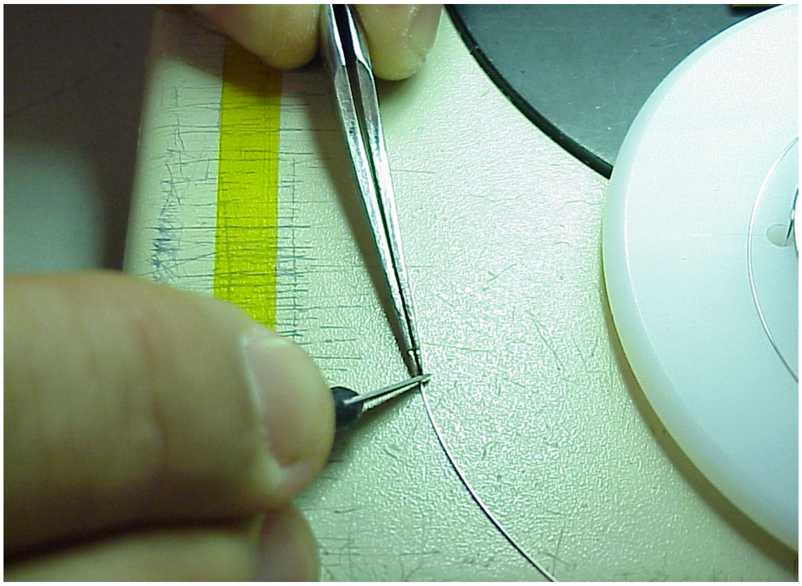

16. Place the two pieces of solder in parallel, directly over the thermocouple bead (Figure 35).

Figure 35. Positioning Solder on IHS

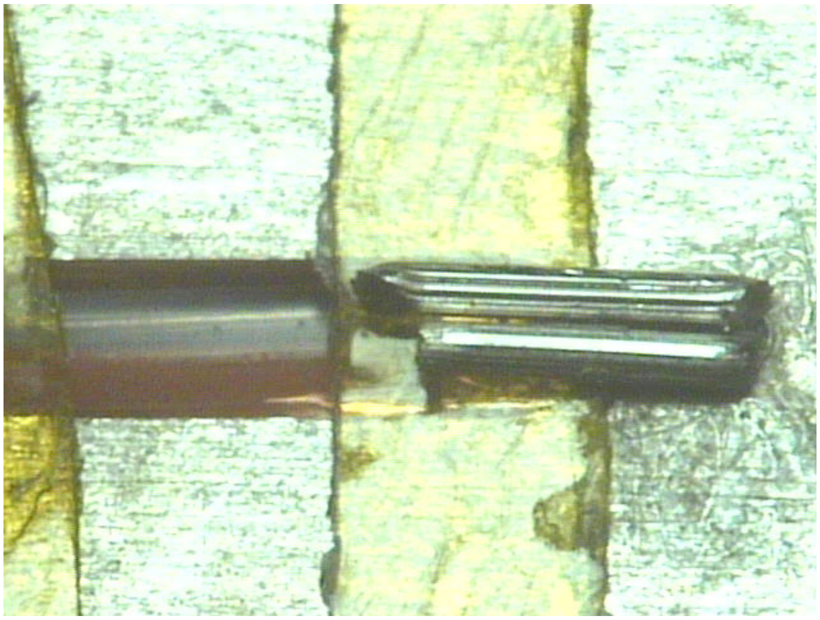

17. Measure the resistance from the thermocouple end wires again using the DMM (refer to Appendix C.5.1, "Thermocouple Conditioning and Preparation," step 2) to ensure the bead is still properly contacting the IHS.

C.5.3 Solder Process

18. Make sure the thermocouple that monitors the Solder Block temperature is positioned on the Heater block.
19. Connect the thermocouple to a handheld meter to monitor the heater block temperature.
20. Verify the temperature of the Heater block station has reached 155°C ±5°C before you proceed.
21. Connect the thermocouple for the device being soldered to a second hand held meter to monitor IHS temperature during the solder process.

Figure 36. Solder Station Setup

22. Remove the land side protective cover and place the device to be soldered in the solder station. Make sure the thermocouple wire for the device being soldered is exiting the heater toward you.

Note: Don’t touch the copper heater block at any time as this is very hot.

23. Move a magnified lens light close to the device in the solder status to get a better view when the solder begins to melt.
24. Lower the Heater block onto the IHS. Monitor the device IHS temperature during this step to ensure the maximum IHS temperature is not exceeded.

Note: The target IHS temperature during reflow is 150°C ±3°C. At no time should the IHS temperature exceed 155 °C during the solder process as damage to the device may occur.

25. You may need to move the solder back toward the groove as the IHS begins to heat. Use a fine tip tweezers to push the solder into the end of the groove until a solder ball is built up (Figure 37 and Figure 38).

Figure 37. View Through Lens at Solder Station

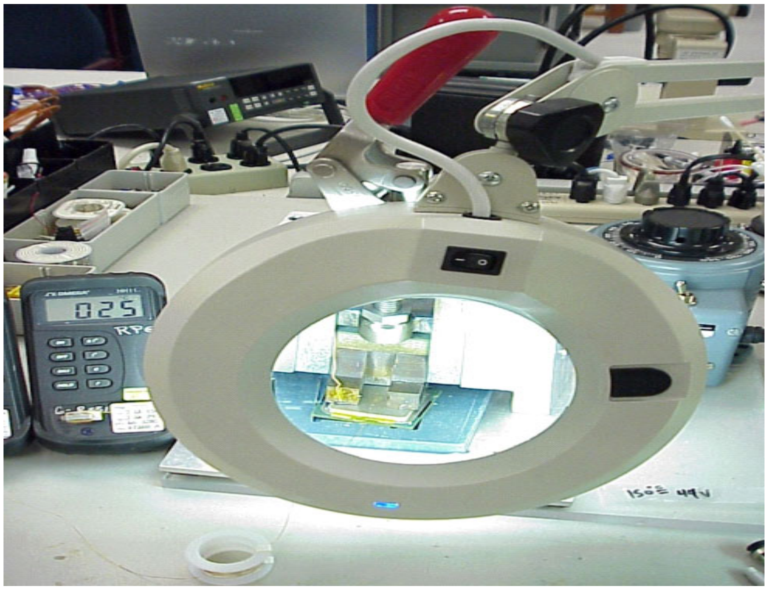

Figure 38. Moving Solder back onto Thermocouple Bead

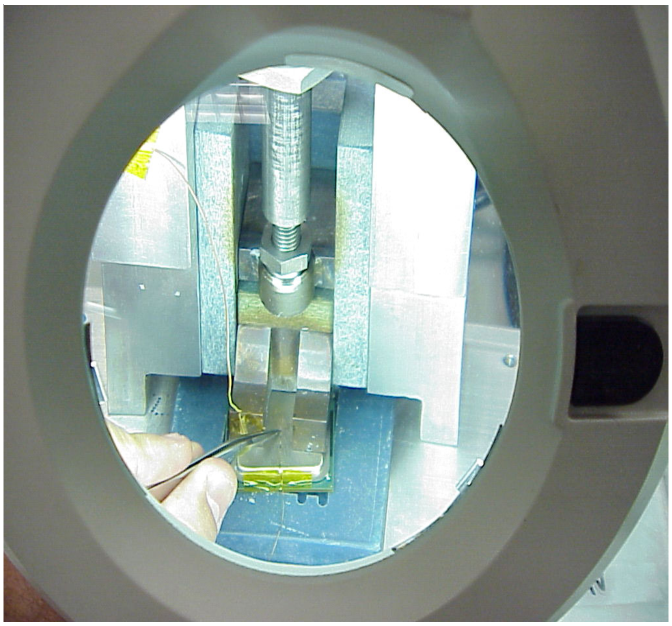

26. Lift the heater block and magnified lens, and using tweezers, quickly rotate the device 90 degrees clockwise. Using the back of the tweezers, press down on the solder. This will force out the excess solder (Figure 39).

Figure 39. Removing Excess Solder

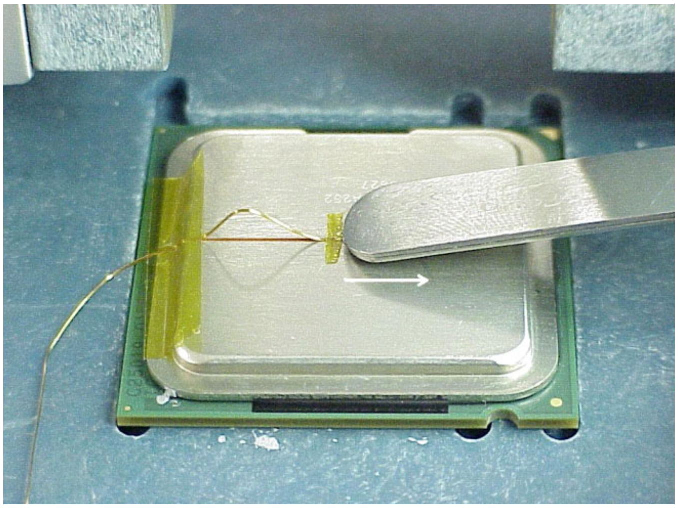

27. Allow the device to cool down. Blowing compressed air on the device can accelerate the cooling time. Monitor the device IHS temperature with a handheld meter until it drops below 50° C before moving it to the microscope for the final steps.

C.5.4 Cleaning & Completion of Thermocouple Installation

1. Remove the device from the solder station and continue to monitor the TTV IHS temperature with a handheld meter. Place the device under the microscope and remove the three pieces of Kapton* tape with Tweezers, keeping the longest for re-use.
2. Straighten the wire and work the wire into the groove. Bend the thermocouple over the IHS. Replace the long piece of Kapton* tape at the edge of the IHS.

Note: The wire needs to be straight so it doesn't sit above the IHS surface at anytime (Figure 40 on page 59).

Figure 40. Thermocouple placed into groove

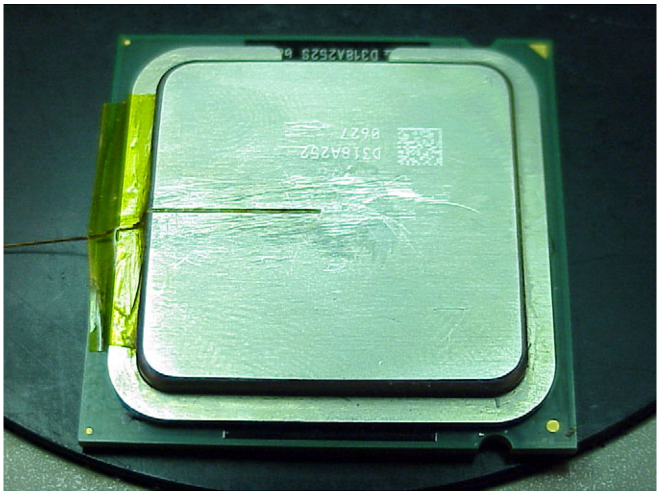

3. Using a blade, carefully shave the excess solder above the IHS surface. Only shave in one direction until solder is flush with the groove surface (Figure 41).

Figure 41. Removing Excess Solder

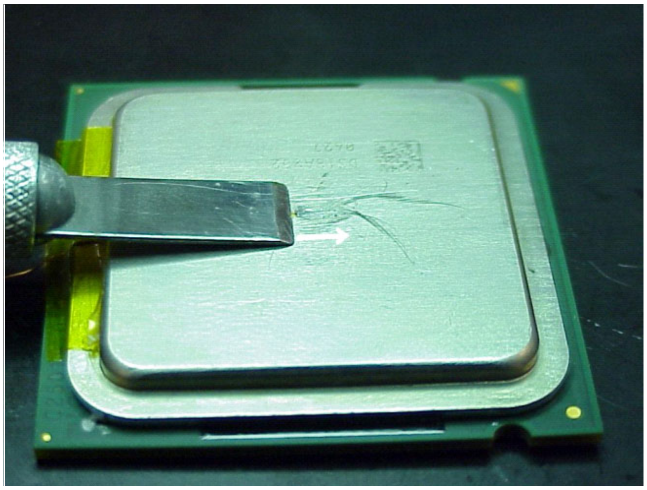

Note: Take usual precautions when using open blades

4. Clean the surface of the IHS with alcohol and use compressed air to remove any remaining contaminants.
5. Fill the rest of the groove with Loctite* 498 adhesive. Verify under the microscope that the thermocouple wire is below the surface along the entire length of the IHS groove (Figure 42 on page 60).

Figure 42. Filling Groove with Adhesive

6. To speed up the curing process, apply Loctite* accelerator on top of the adhesive and let it set for a couple of minutes (Figure 43).

Figure 43. Application of Accelerant

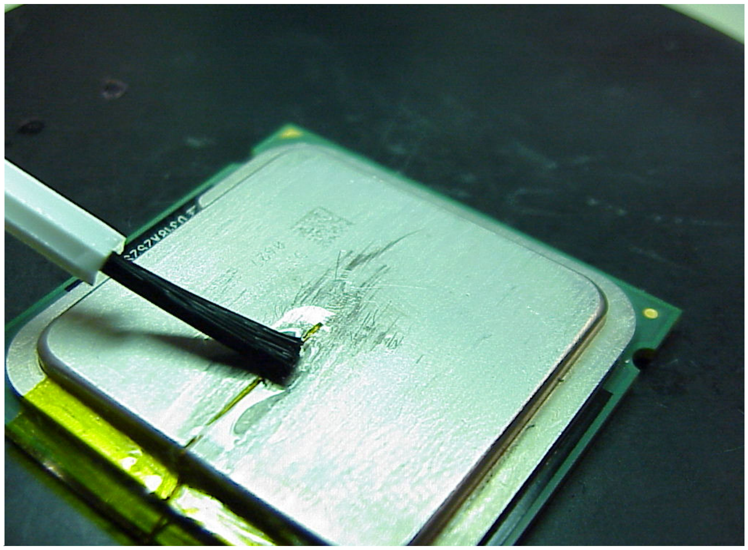

Figure 44. Removing Excess Adhesive from IHS

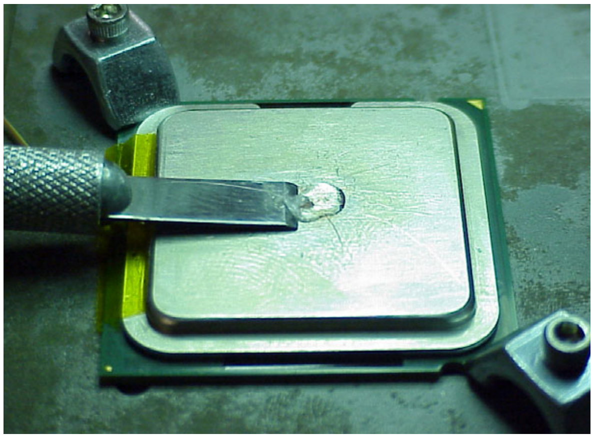

7. Using a blade, carefully shave any adhesive that is above the IHS surface (Figure 44). The preferred method is to shave from the edge to the center of the IHS.

Note: The adhesive shaving step should be performed while the adhesive is partially cured, but still soft. This will help to keep the adhesive surface flat and smooth with no pits or voids. If there are voids in the adhesive, refill the voids with adhesive and shave a second time.

8. Clean IHS surface with IPA and a wipe.
9. Clean the LGA pads with IPA and a wipe.
10. Replace the land side cover on the device.
11. Perform a final continuity test.
12. Wind the thermocouple wire into loops and secure, or if provided by the vendor, back onto the plastic roll (Figure 45).

Figure 45. Finished Thermocouple Installation

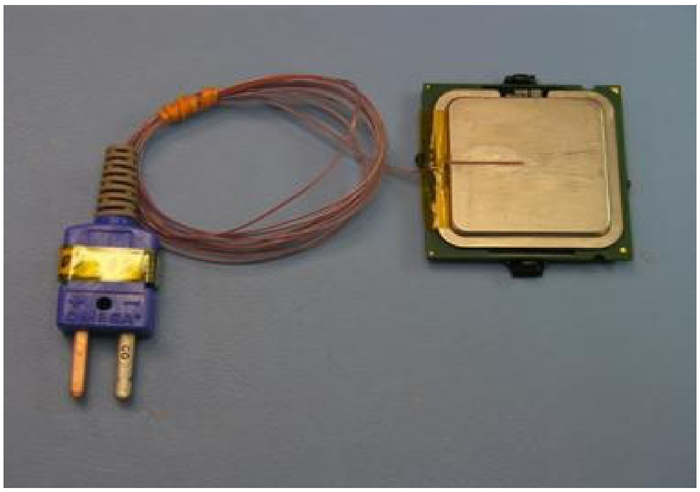

13. Place the device in a tray or bag until it's ready to be used for thermal testing use.

C.6 Thermocouple Wire Management

When installing the processor into the socket, make sure that the thermocouple wires exit above the load plate as shown in Figure 46. 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 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 on the edge of the IHS/socket load plate area. This temperature is usually 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 46. Thermocouple Wire Management

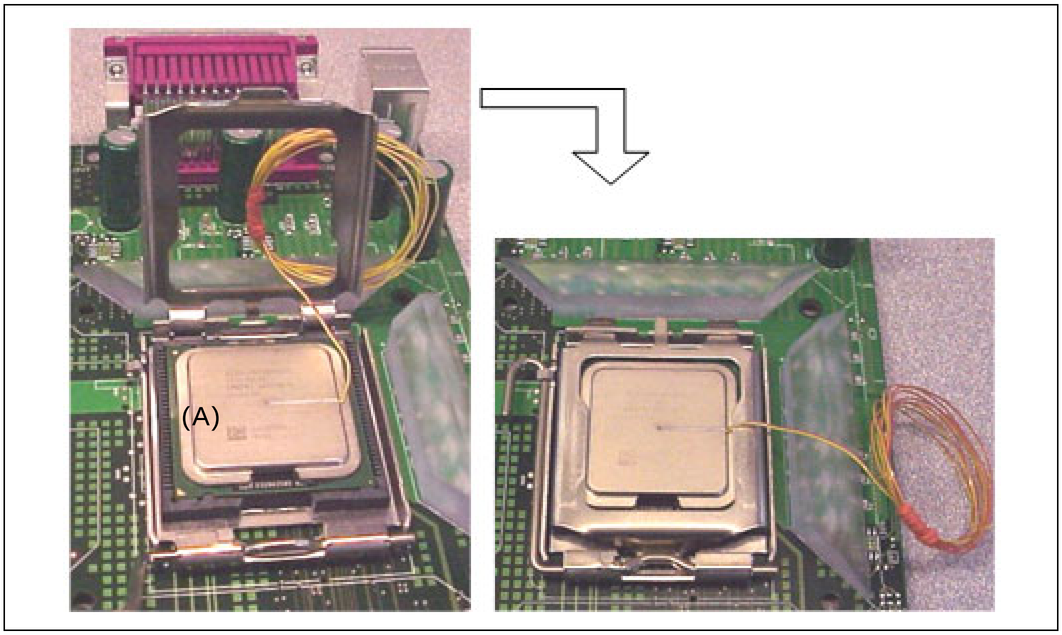

Appendix D Validation of System Thermal Solution

This section provides system level validation information for an active fan thermal solution that monitors the on-die thermal diode. It is assumed the system integrator has already validated the thermal solution and meets the thermal profile using the TTV as described in Chapter 3.0, “Thermal Metrology.”

The functional validation needs to account for the entire operating environment of the system. In addition, the fan speed control implementation must be verified. The system engineer will need to account for the following factors:

- System ambient temperature range
- Interaction of internal heat loads such as hard drives, optical drives, PCI, PCI Express*, graphics, etc.
- The Thermal Profile specification is for components at end of life not beginning of life
- Variations in heatsink and thermal interface material
- Equipment calibration

D.1 Processor Power Dissipation

Each processor is calibrated at the factory with a $T_{CONTROL}$ value. That value is determined in large part by the leakage power. A processor with a higher $T_{CONTROL}$ will inherently dissipate more power than a part with a low $T_{CONTROL}$ when running the same application workload.

Each processor will have characteristic maximum power dissipation. The characteristic maximum power dissipation is defined at the TDP workload percentage. To determine the maximum characteristic power for a specific processor, run the max power program at the TDP workload percentage and near $T_{CASE-MAX}$. At the time of publication, the recommended TDP workload percentage for the Intel Celeron processor 440 is 80%. Intel continues to evaluate this value and may make changes in the future.

The characteristic power dissipation of any randomly selected processor is unlikely to be equal to the TDP power as listed in the EMTS. The TDP power may be approached by running the max power program at greater than TDP workload. There is a risk when operating greater than TDP workload that the PROCHOT# signal or even THERMTRIP# signal may go active and invalidate the test.

D.2 Preparation

Follow the electrical load line characterization procedure for the board(s) that will be used in the system thermal test. An accurate load line is necessary to determine the power dissipated during the test. The power dissipated can then be plotted on the thermal profile to determine compliance to the maximum case temperature.

Install thermocouples on selected functional processors as described in Appendix C, “Case Temperature Reference Metrology”. It is beneficial, but not required to test with processors having different $T_{CONTROL}$ values. Selecting a pair of processors with a high and low $T_{CONTROL}$ value can help confirm compliance to thermal and acoustic requirements over a range of $T_{CONTROL}$ values.

D.3 System Setup

Instrument the system to capture the following data:

- Processor case temperature
- Processor on-die thermal diode temperature
- Heatsink local ambient temperature
- Processor power
- PROCHOT# signal
- Fan speed to estimate acoustic signature with dBA vs. RPM curves (optional)
- Other points of interest for overall system validation

The data collection process should sample all parameters at least 1/sec. Reading the on-die thermal diode more frequently is encouraged as the rate of change for the thermal diode can be up to 50°C/sec.

D.4 System Test Conditions

The system integrator should select system thermal loads and ambient temperatures that are consistent with end use conditions of the system under test. The choice of active or passive system loads is up to the system integrator. The processor can be exercised with the max power program or other suitable high power applications. Example load scenarios for thermal / acoustic validation are in Table 8. The conditions outlined here may differ from your own system requirements.

Table 8. Example System Test Conditions

Condition	Ambient (°C) (External)	Power
Minimum Loading for Acoustic Verification	25	Windows XP* Idle
Typical Operation Load and Environment	25	User Defined
Maximum Operating Load & Environment	35	80%

D.5 Pass / Fail Criteria

The data logs should be checked to identify the test conditions where T_{DIODE} exceeds $T_{CONTROL}$. To determine pass/fail, the power dissipation must be plotted on the thermal profile to get the maximum allowed case temperature. If the measured case temperature is less than the maximum case temperature from the thermal profile at that power dissipation, the test is successful. By design, the thermal solution fan should be at 100% of the fan operating speed for the ambient temperature being examined when T_{DIODE} exceeds $T_{CONTROL}$.

If PROCHOT# goes active during any validation run, that test should be considered a failure. When PROCHOT# is active, Thermal Monitor will take steps to reduce the processor power, see Chapter 4.0, "Thermal Management Logic and Thermal Monitor Feature" for further details. The system integrator should carefully review the test data to determine the conditions that led to the activation of Thermal Monitor.

For all other conditions where T_{DIODE} is less than $T_{CONTROL}$, the thermal solution is providing sufficient cooling to meet the processor thermal specification.

Appendix E Mechanical Drawings

Table 9 lists the mechanical drawings included in this appendix. These drawings refer to the reference thermal mechanical enabling components for the processor.

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

Table 9. Mechanical Drawings

Drawing Description
Figure 47, "1U or PICMG 1.3 Motherboard Keep-out Footprint Definition and Height Restrictions for Enabling Components, Primary Side" on page 66
Figure 48, "1U or PICMG 1.3 Motherboard Keep-out, Secondary Side" on page 67
Embedded ATX Motherboard Keep-out Footprint Definition and Height Restriction for Enabling Components (Please refer to doc #657586 Processor Thermal Mechanical Design Guide appendix section.)
Figure 49, "Embedded ATX Reference Solution Assembly Stack-up on Motherboard" on page 68

Figure 47. 1U or PICMG 1.3 Motherboard Keep-out Footprint Definition and Height Restrictions for Enabling Components, Primary Side

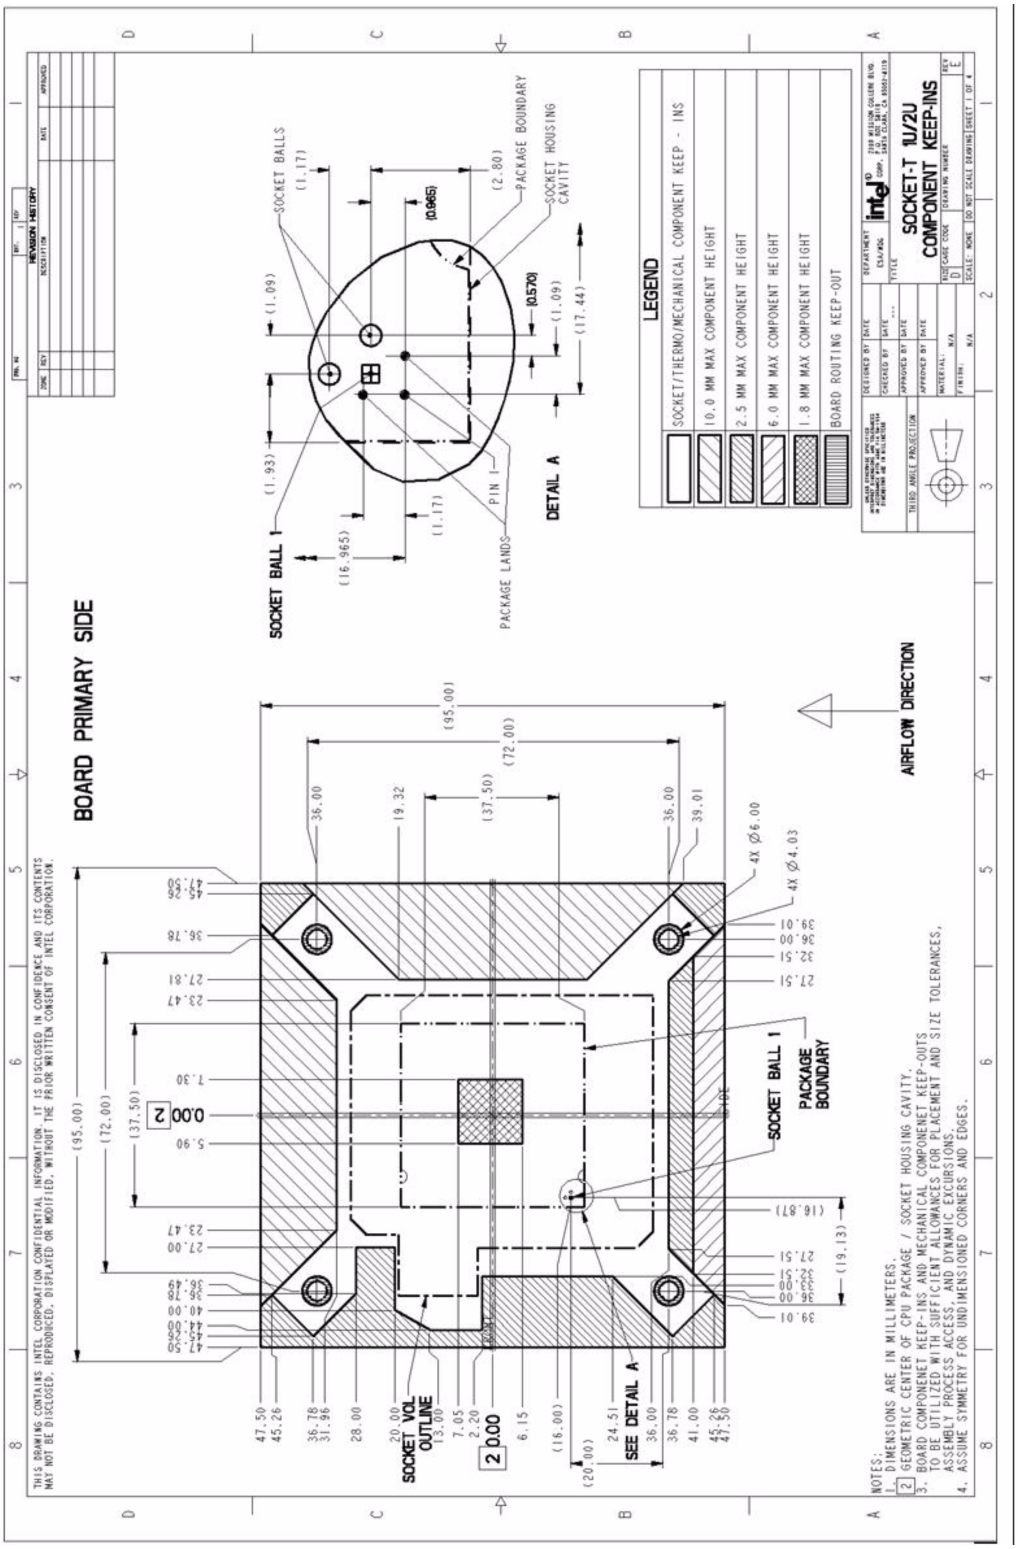

Figure 48. 1U or PICMG 1.3 Motherboard Keep-out, Secondary Side

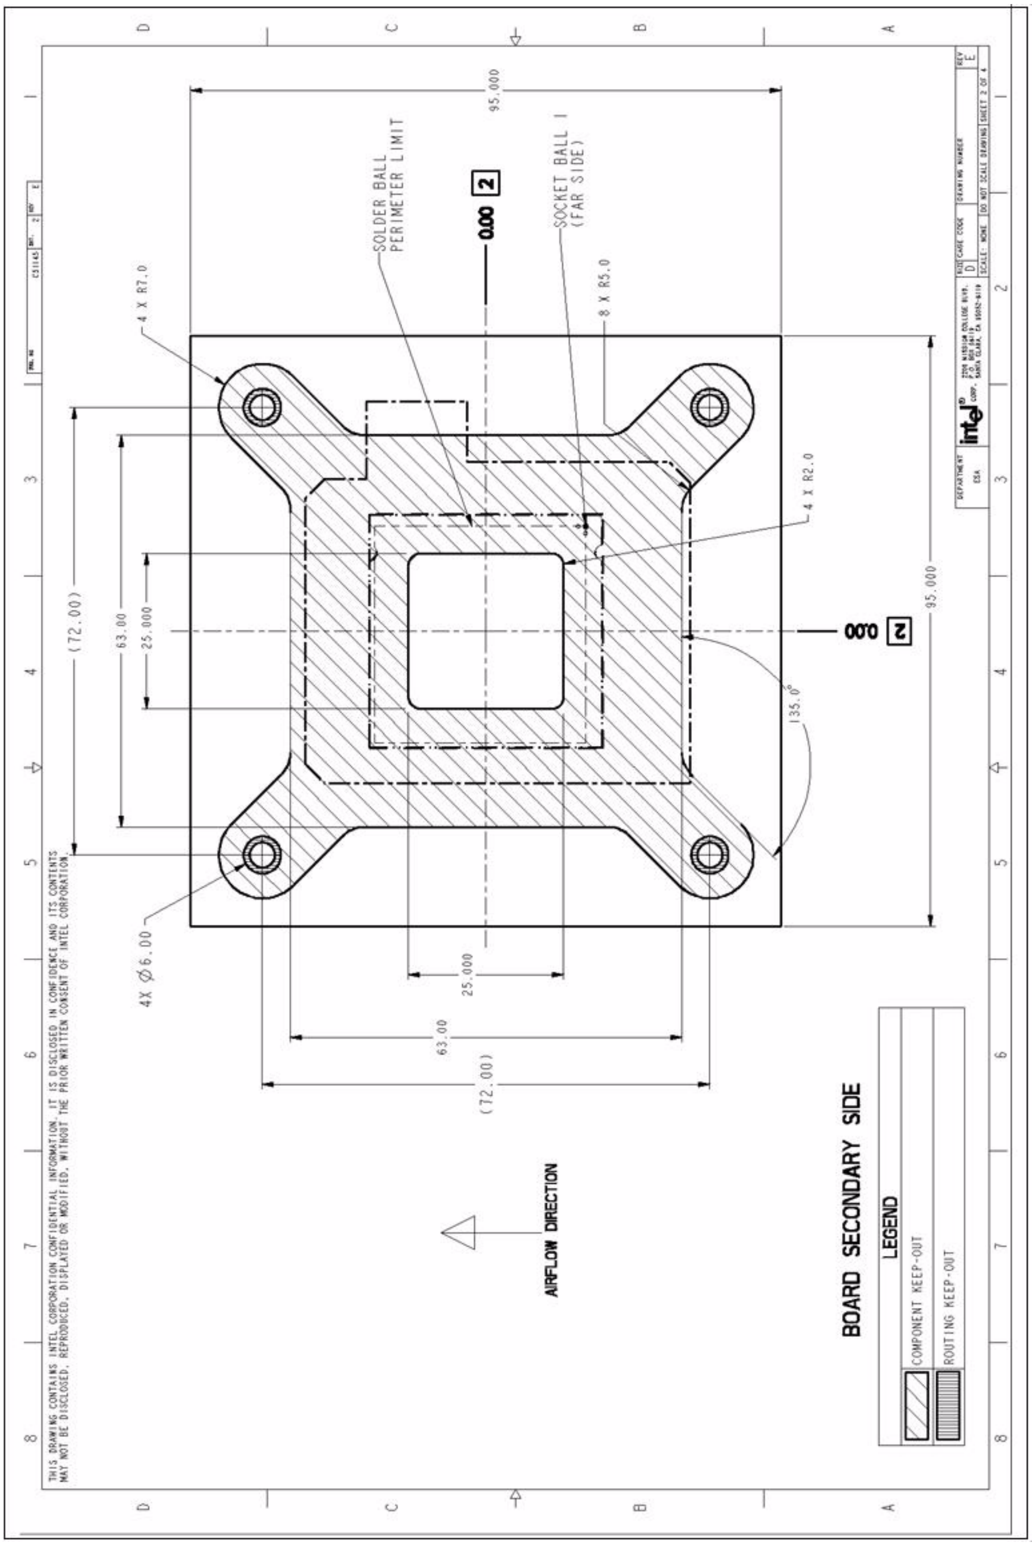

Figure 49. Embedded ATX Reference Solution Assembly Stack-up on Motherboard

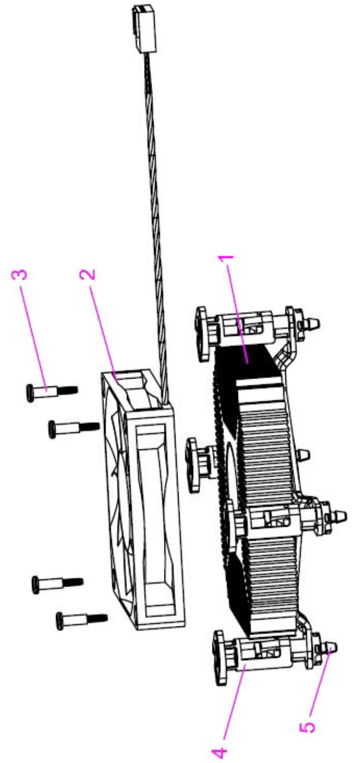

	DESCRIPTION / COMPONENTS	QTY
1	TIM	1
2	PUSH PIN 01	4
3	PUSH PIN 01	4
4	MOUNTING SCREW	4
5	DC FAN	1
6	HEATSINK	1

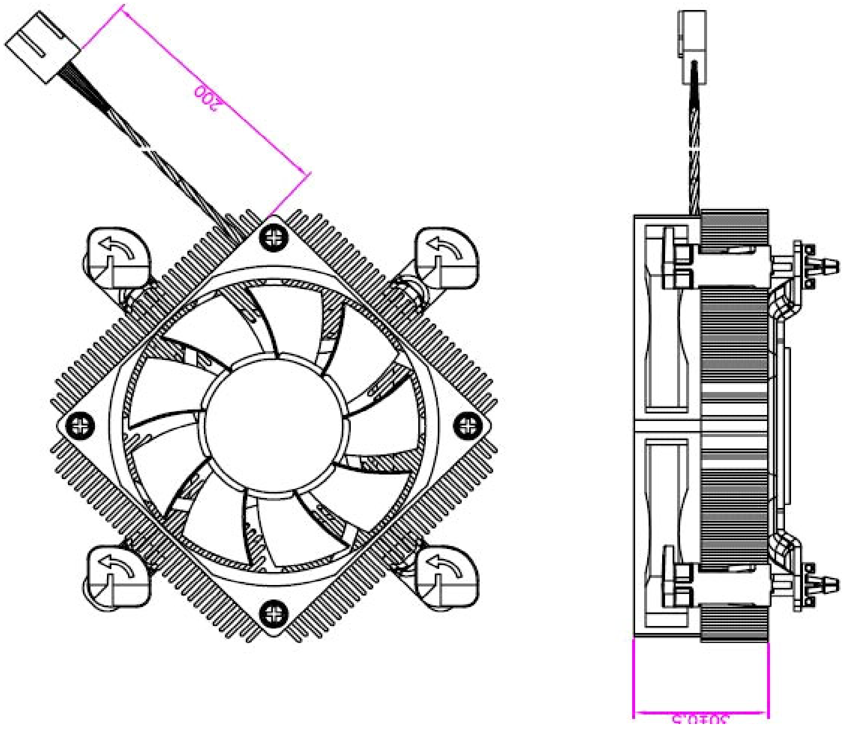

Appendix F Intel Enabled Reference Solution Information

This appendix includes supplier information for Intel enabled vendors for the 1U or PICMG 1.3, and Embedded ATX thermal solution.

Table 10 lists suppliers that produce Intel enabled reference components. The part numbers listed below identifies these reference components. End-users are responsible for the verification of the Intel enabled component offerings with the supplier. OEMs and System Integrators are responsible for thermal, mechanical and environmental validation of these solutions.

Table 10. Intel Reference Component PICMG 1.3 Thermal Solution Providers

Supplier	Part Number	Contact	Phone	Email
AVC* (ASIA Vital Components Co., Ltd.)	C69175	David Chao	+886-2-22996930 Extension: 619	david_chao@avc.com.tw
Cooler Master*	CONROE-L-HS-T1	Ivan Yang	+886-2-3234-0050 ext: 255	Ivan_yang@coolermaster.com.tw

Note: These vendors and devices are listed by Intel as a convenience to Intel's general customer base, but Intel does not make any representations or warranties whatsoever regarding quality, reliability, functionality or compatibility of these devices. This list and/or these devices may be subject to change without notice.