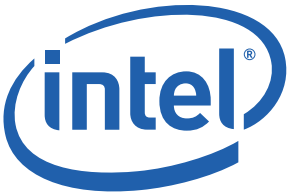

Intel® Atom™ Processor S1200 Product Family for Microserver

Thermal and Mechanical Design Guidelines (TMDG)

November 2012

Document Number: 328196-001

Revision History

Date	Revision	Description
November 2012	1.0	Initial Release

Contents

Figures

Tables

§ §

1 Introduction

The Intel® Atom™ Processor S1200 Product Family for Microserver is the new generation of System-On-Chip (SoC) 64-bit processor; dual-core processors built on 32-nanometer process technology. Throughout this document, the Intel® Atom™ Processor S1200 Product Family for Microserver is also referred to as the processor or S12x0.

This document relates to the following Intel® Atom™ Processor S1200 Product Family for Microserver:

Table 1-1. Intel® Atom™ Processor S1200 Product Family for Microserver

SKU Name	Thermal Design Power (TDP)
Intel® Atom™ Processor S1260	8.5W
Intel® Atom™ Processor S1240	6.3W
Intel® Atom™ Processor S1220	8.1W

1.1 Document Goals and Scope

This document provides thermal and mechanical specifications and design guidelines for the S12x0 processor.

Key content in this document include:

- Thermal and mechanical specifications for the S12x0 dual-core processor.
- Reference thermal solutions for the processor.
- Thermal management features.

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, airflow over the processor, and 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 (with more transistors and smaller size). As operating frequencies increase and packaging size decreases, the resulting power density increases. At the same time, the thermal solution space and airflow are becoming more constrained as system sizes trend smaller. The result is an increased importance on system design to ensure that thermal design requirements are met for each component including the processor.

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 multi-processor (multi-node) systems using S12x0 dual-core processors.

The concepts given in this document are applicable to any system form factor. Specific examples used are the Intel enabled reference solution for a microserver platform system.

1.1.3 Document Scope

This specifications and design guide supports the following processors:

- S12x0 dual-core processors

In this document when a reference is made to the processor, the reference includes all the processors supported by this document. If needed for clarity, the specific processor is listed.

In this document, when a reference is made to the reference design, the reference includes all reference designs supported by this document. If needed for clarify, the specific reference design is listed.

1.2 Reference Documents

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

Table 1-2. Reference Documents

Document	Document No./Location
Intel® Atom™ Processor S1200 Product Family for Microserver Datasheet Vol. 1 of 2	328194
Intel® Atom™ Processor S1200 Product Family for Microserver Datasheet Vol. 2 of 2	328195
System Mechanical Design Guidance for Dynamic Events – Application Notes Briefs	383578
Board Flexure Initiative (BFI) - Manufacturing Advantage Service (MAS)	http://www.intel.com/design/quality/cme.htm

Notes:
1. Contact the local Intel field sales representative for the latest revision and order number of these documents.

1.3 Definition of Terms

Table 1-3. Definition of Terms

Term	Description
CTF	Critical To Function.
DTS	Digital Thermal Sensor.
FCBGA	Flip Chip Ball Grid Array.
MD	Metal Defined (Pad).
NCTF or nCTF	Non-Critical To Function.
Ψ_{JA} (Psi-ja)	Junction-to-ambient thermal characterization parameter (Psi). A measure of thermal solution performance using total package power. Defined as $(T_J - T_A) / TDP$. ***Note:*** Heat source must be specified for Ψ_{JA} measurements.
SMD	Solder Mask Defined (Pad).
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_J	Processor junction temperature and defined as the maximum temperature of the die active surface.
$T_{J\text{-}MAX}$	Maximum processor junction temperature specification.
$T_{J\text{-}MIN}$	Minimum processor junction temperature specification (defined as the minimum temperature of the die active surface).
THS	Thin Heat Spreader.
TIM	Thermal Interface Material: The thermally conductive compound between the heatsink and the processor die surface. This material fills the air gaps and voids, and enhances the transfer of the heat from the processor die surface to the heatsink.
TDP	Thermal Design Power: A power dissipation target based on worst-case applications. Thermal solutions should be designed to dissipate the thermal design power.
WTMD	Wide Trace Metal Defined (Pad).

§ §

2 Package Mechanical Specifications

2.1 Package Mechanical Description

The processor is packaged in a 1283 ball Flip-Chip Ball Grid Array (FCBGA) package that is soldered directly to a PCB surface without the use of a socket. The package size is nominally 34 mm x 28 mm. The ball pitch is variable with a minimum 0.7 mm pitch. The nominal die size is 9.86 mm x 10.0 mm.

2.2 Package Mechanical Drawing

The detailed package mechanical drawing is included in Appendix A, "Mechanical Drawings." Refer to the *Intel® Atom™ Processor S1200 Product Family for Microserver Solder Ball Location Coordinates* spreadsheet for solder ball location coordinates.

2.2.1 Processor Component Keep-Out Zones

The FCBGA package may have capacitors placed in the area surrounding the processor die. The die-side capacitors, which are only slightly shorter than the die height, are electrically conductive and contact with electrically conductive materials should be avoided. The use of an insulating material between the capacitors and any thermal and mechanical solution should be considered to prevent capacitors shorting. A thermal and mechanical solution design must not intrude into the required keep-out zones as specified in Appendix A, "Mechanical Drawings."

2.3 PCB Pad Recommendations

Intel recommends implementing board pad sizes and types shown in Table 2-1 on page 13 and Figure 2-1 on page 12 for processor-printed circuit-board designs. These recommendations allow for trace breakout in a typical 12-layer board design. In addition, the corner pad design has been optimized for mechanical strength and protection to inner critical-to-function interconnects. The processor PCB CAD symbol has the recommended pad sizes and types implemented. See Section 1.2, "Reference Documents" for the document number for the processor PCB symbol.

2.3.1 Pad Type Recommendations

In the processor footprint, Intel defines two pad types based on how the pad is constructed. A Metal Defined (MD) pad is where a copper pad is individually etched into the PCB with a minimum width trace exiting it. The Solder Mask Defined (SMD) pad is typically a pad in a flood plane where the solder mask opening defines the pad size for soldering to the component.

In thermal cycling a MD pad is more robust than a SMD pad type. The solder mask that defines the SMD pad can create a sharp edge on the solder joint as the solder ball/paste conforms to the window created by the solder mask.

The solder joints under the die (die shadow) can experience increased stress due to a Coefficient of Thermal Expansion (CTE) mismatch between the package and board. The size of the die tends to influence the localized CTE of the package substrate driving higher stresses on the solder joints under the die shadow. For this reason the recommended pad type in this region is MD.

For certain failure modes, the MD pad may not be as robust in shock and vibration (S&V). During S&V, the predominant failure mode for an MD pad in the corner of the BGA layout is a pad crater and solder joint crack. A corner MD pad can be designed to absorb more dynamic energy by having a wide trace exiting the pad (if other escape routing constraints allow). This trace should be 10-mils minimum (but should not exceed the pad diameter) and should exit the pad at a 45-degree angle parallel to the diagonal of the package. During board flexure that results from shock and vibration, a SMD pad, due to the larger surface area, is less susceptible to a crack initiating.

Figure 2-1. Pad Size and Shape Recommendations (looking at PCB land)

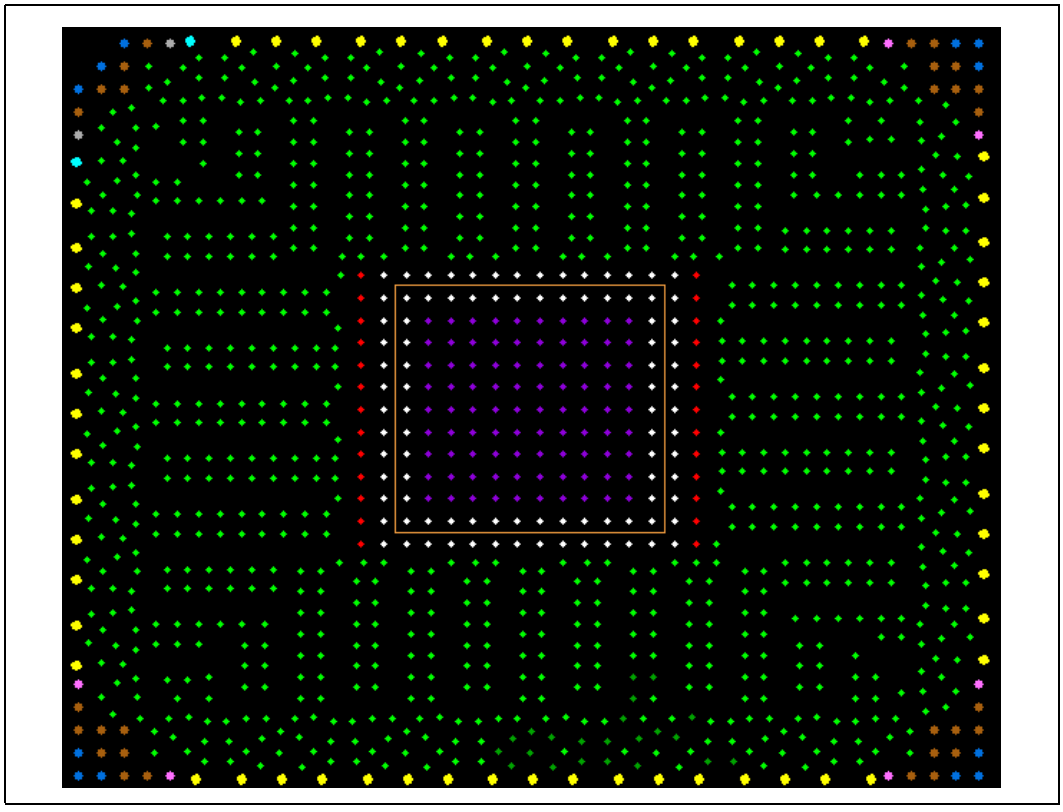

Table 2-1. Pad Size and Shape Recommendations

Pad Color	Pad Description	Quantity
Blue	Corner nCTF 16 mil/SMD	12
Orange	Corner nCTF 16 mil/MD	29
Gray	11 x 13.5 mil/MD-Oval/nCTF	2
Cyan	10 x 13.5 mil/MD-Oval/nCTF	2
Pink	11 x 13.5 mil/MD-Oval/CTF	6
Yellow	10 x 13.5 mil/MD/WTMD-Oval/CTF	58
Green	12 mil/MD/WTMD/SMD/CTF	966
Red	15 mil/MD/WTMD/CTF	26
White	Die Shadow nCTF 15 mil/MD/WTMD	92
Purple	Die Shadow CTF 15 mil/MD/WTMD	90
	Total	1283

Notes:
1. nCTF = Non-Critical to Function.
2. CTF = Critical to Function.
3. MD = Metal Defined Pad.
4. SMD = Solder Mask Defined Pad.
5. WTMD = Wide-Trace Metal Defined Pad.
6. 2WTMD/SMD pads are NOT allowed under die shadow.
7. I/O trace = 4 mil, WTMD = 10 mil.
8. No uVia or Via in pad is allowed.
9. The S12x0 processor is planned for a 62-mil PCB ONLY.

2.3.2 Solder Pad Recommendation

Intel has defined selected solder joints as Non-Critical-to-Function (NCTF) when evaluating package solder joints after environmental testing. The signals at NCTF locations are typically redundant ground or non-critical reserved. The loss of the solder joint continuity at end-of-life conditions does not affect the overall product functionality. Table 2-1 and Figure 2-1 on page 12 identify NCTF solder joints.

Additional protection from pad cratering on the motherboard has been demonstrated through the usage of thick traces at the corner NCTF ball locations. NCTF trace thicknesses of 60-80% of the pad diameter were tested in board-level shock tests with metal defined pads which reduced the occurrence of pad cratering failures. Pad cratering is the failure mode in which solder pads on the motherboard separate from the PCB.

The thick traces shown in Figure 2-2 on page 14 are an example of how thick traces may be used at NCTF pads.

Note: The NCTF locations shown in Figure 2-2 on page 14 are not the NCTF locations of the processor package and are shown to illustrate the application of thick traces. Designers are encouraged to use thick traces in designs where pad cratering has occurred along the corners of the package. The thick traces effectively increase the strength of the pad to motherboard interface and may cause a crack to initiate in a different failure mode in the NCTF solder ball while increasing the shock margin.

Figure 2-2. Example of Thick Traces Used in a Desktop BGA

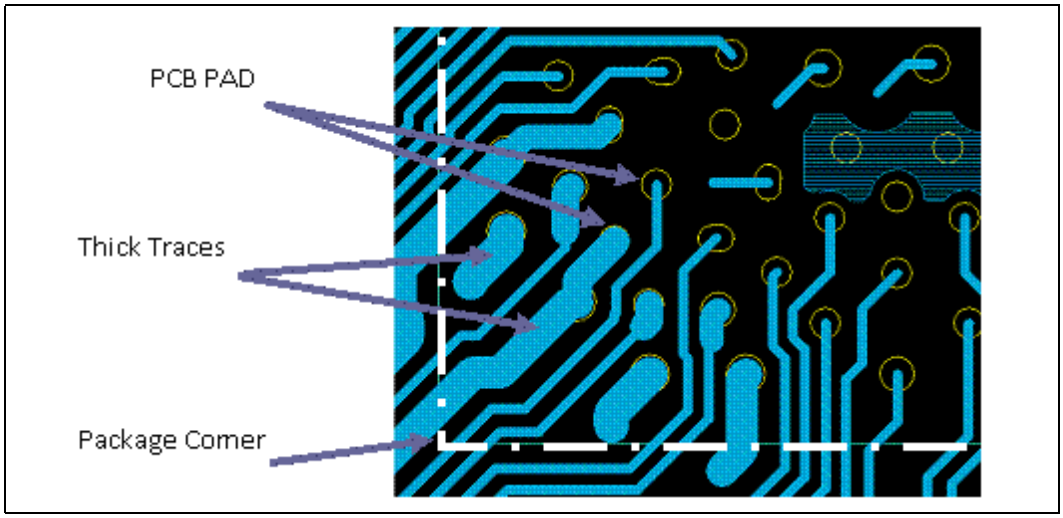

2.4 Package Mechanical Loading Specifications

The processor thermal solution applies a mechanical load on the processor die for both thermal and mechanical functionality and must comply with the maximum compressive loading specification.

- When a compressive static load is necessary to ensure thermal performance of the thermal interface material between the heatsink base and the processor die, the compressive static load should not exceed the corresponding specification.
- When a compressive static load is necessary to ensure mechanical performance, it should remain below the maximum limit specified.

No portion of the substrate should be used as a mechanical reference or load-bearing surface for the thermal or mechanical solution.

Table 2-2. Processor Loading Specifications

Parameter	Maximum
Static Compressive Load	67N [15 lbf]

Notes:
1. These specifications apply to uniform compressive loading in a direction normal to the processor die.
2. This is the maximum static force that can be applied by the heatsink and retention solution to maintain the heatsink and processor interface.
3. These specifications are based on limited testing for design characterization. Loading limits are for the package only.

2.5 Processor Markings

Figure 2-3 shows the topside markings on the processor. This diagram is to aid in the identification of the processor.

Figure 2-3. Processor Top-Side Markings

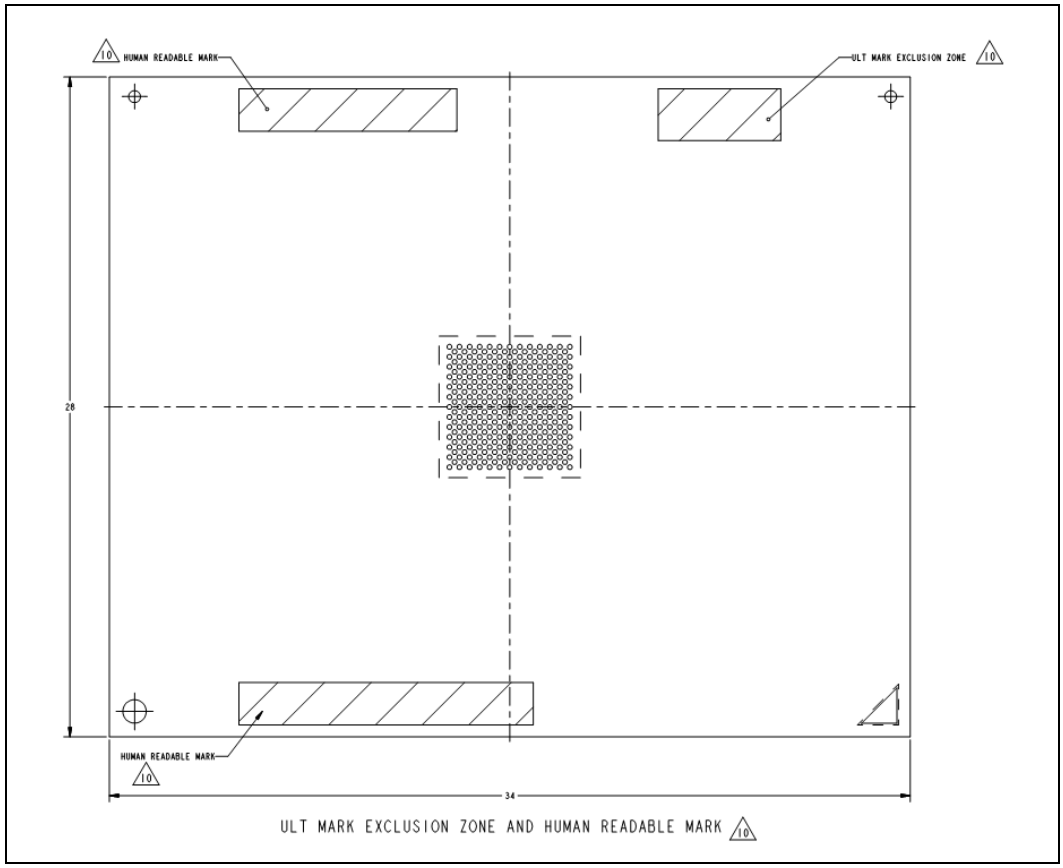

2.5.1 Shock Strain Guidance

Solder Joint Reliability (SJR) remains a major topic of concern in designing systems especially for surface mounted components such as the S12x0 processor. Solder ball cracking and fracture is a failure mode associated with overstressing the surface mounted component on the motherboard. The overstressing typically occurs when the motherboard is subjected to bending deflection. The deflection of the motherboard applies loads to these surface-mounted components that attempt to peel the component from the board. These loads stress the solder balls of the component and either initiate cracks, which grow through the solder during thermal and power cycling, or cause fracture, which results in an electrical open.

Loading conditions such as shock typically stress the motherboard and generate the stresses at the solder joint that lead to either crack initiation or complete fracture of the ball. This section addresses guidance specific to this SoC product. Refer to the *Desktop System Mechanical Design Guidance for Dynamic Events* for more information on system design guidance and best practices.

Section 2.3.2, "Solder Pad Recommendation" describes the function of the Non-Critical-to-Function (NCTF) solder balls. These balls are located in the corners of the ball grid array where they are most susceptible to stressing from motherboard flexure. These NCTF balls do well to mitigate degradation to component performance once damage has occurred at the solder balls. General design guidance is available in the *Desktop System Mechanical Design Guidance for Dynamic Events for Solder Joint Reliability*.

A useful metric to compare the impact of design modifications to Solder Joint Reliability (SJR) and assess SJR risk during shock events is strain measurement. This strain measurement, also referred to as shock strain, utilizes strain gauges to measure the surface strain of a motherboard.

Note: Intel also publishes strain guidance specifically for manufacturing. This manufacturing guidance is part of the Board Flexure Initiative (BFI) and those strain limits are commonly referred to as BFI strain. More information is available in the BFI Manufacturing Advantage Service (MAS). DO NOT use BFI strain values for shock strain testing and DO NOT use shock strain guidance for BFI. These two strain metrics are significantly different and are not interchangeable. Using the BFI strain values for a design metric will likely result in a poor system design.

Given parameters unique to the board of interest such as board thickness, the board surface strain directly correlates to the amount of board curvature. The amount of motherboard curvature in the critical locations directly beneath the solder balls is indicative of the reliability of the component solder joints. This measurement is typically made at the corners of the BGA components. The shock strain results are sensitive to the application of the strain gauges. Guidance for strain-gauge application is available in the *Shock Strain Monitoring Customer Reference Document (CRD)*. An Intel Corporate Quality Engineer is also available for help with strain-gauge attach training. This Shock Strain Monitoring CRD outlines the proper selection, application, and usage of the strain gauges and strain instrumentation to attain repeatable and valid results. The Shock Strain Monitoring CRD also discusses proper reduction of the data in order to use the data to compare to the Intel strain guidance.

The strain guidelines have been developed from simulations. The resulting strain guidance in Table 2-3 is provided for a board thickness commonly used in microserver systems. Solder joint failure in critical-to-function solder joints is unlikely when shock strain levels are kept at or below these guidelines. The non-critical-to-function solder balls may have some cracking and fractures when the strain measurements are at these levels.

Table 2-3. Shock Strain Guidelines for the Processor

Shock Strain (Micro Strain, με)	Board Thickness	Associated Risk	Recommendation/Comments
Emin < 1800	62 mils (1.57 mm)	Low	Solder joint failure is unlikely.
1800 < Emin < 2500	62 mils (1.57 mm)	Medium	Larger sample size and failure analysis is suggested for design validation.
2500 < Emin	62 mils (1.57 mm)	High	Solder joint failure is likely; consider design changes to improve reliability.

Notes:
1. Emin is the minimum principal strain as defined in the *Shock Strain Monitoring Customer Reference Document (CRD)*.
2. The shock strain guidelines are applicable to all S12x0 processors which exist across microserver platforms.
3. The strain value guidance is different for different board thicknesses. Contact the local Intel Field Sales Representative if the design uses a different board thickness.
4. These values apply to both FR4 and Halogenated Flame Retardant Free (HFR-Free) boards.

2.5.2 Halogenated Flame Retardant Free PCB

A limited number of FR4 and Halogenated Flame Retardant Free (HFR-Free) boards available from PCB vendors were tested by Intel for shock performance and bending stiffness. Both board materials exhibited similar shock performance as measured by solder crack assessment from the die and pull-failure analysis when the optimized pad definitions were used. Also, the mean flexural modulus of the boards was similar. The flexural modulus variations within the HFR-free samples were significantly larger than the modulus variation observed within the FR4 samples indicating there may be more variation in HFR-Free shock performance.

A high sensitivity to pad definition was found from the Intel testing. The usage of wide or thick traces as incorporated into the package pad definition was found to improve shock performance. The pad definitions defined in the customer reference board are **recommended** for use in both HFR-Free and FR4 boards to optimize shock performance.

§ §

3 Thermal Specifications

The S12x0 processor requires a thermal solution to maintain temperatures within operating limits. Any attempt to operate the processor outside these operating limits may result in permanent damage to the processor and potentially other components in the system. As processor technology changes, thermal management becomes increasingly crucial when building computer systems. Maintaining the proper thermal environment is crucial to reliable, long-term system operation.

3.1 Thermal Specifications

To allow for the optimal operation and long-term reliability of Intel processor-based systems, the system/processor thermal solution should keep the processor within the minimum and maximum junction temperature ($T_{J\text{-}MAX}$) specifications at any processor power level as listed in Table 3-1 on page 19. Designing to this specification allows optimization of thermal designs for processor performance.

The thermal limit for the processor is the maximum junction temperature ($T_{J\text{-}MAX}$). The maximum junction temperature is defined as the maximum temperature on the processor-die-active surface. Analysis indicates that real applications are unlikely to cause the processor to consume the theoretical maximum power dissipation for sustained time periods. Intel recommends that complete thermal solution designs target the TDP indicated in Table 3-1 on page 19 instead of the maximum processor power consumption. The Intel® Thermal Monitor feature is designed to help protect the processor in the unlikely event that an application exceeds the TDP recommendation for a sustained period of time. Refer to Section 3.2.2, "Intel® Thermal Monitor Implementation" for more details on the usage of this feature.

In additional to the maximum junction temperature and TDP specifications, guidance values for a lower TDP are provided at lower junction temperatures in Table 3-1 on page 19. This information is intended to allow power-sensitive microserver-system designs to reduce total power consumption using increased thermal solution capability. By maintaining junction temperatures below the maximum allowable specification, the system designer can decrease total power consumption for the total system.

3.1.1 Processor Thermal Specifications

Table 3-1. Thermal Specifications for S12x0 Dual-Core Processors

Processor	Processor Number	Core Frequency	Thermal Design Power [1, 2]	$T_{J\text{-}MIN}$ [3, 4]	$T_{J\text{-}MAX}$ [3, 4]
Intel® Atom™ Processor S1240	0x0C72	1.6 GHz	6.0 W[5]	0 °C	90 °C[5]
			6.3 W	0 °C	102 °C
Intel® Atom™ Processor S1220	0x0C73	1.6 GHz	7.9 W[5]	0 °C	90 °C[5]
			8.1 W	0 °C	95 °C
Intel® Atom™ Processor S1260	0x0C75	2.0 GHz	8.1 W[5]	0 °C	90 °C[5]
			8.5 W	0 °C	95 °C

Notes:
1. Thermal Design Power (TDP) should be used for general thermal solution design targets. The TDP is not the maximum theoretical power the processor can generate.
2. Not 100% tested. These power specifications are determined by characterization of the processor currents at higher temperatures and extrapolating the values for the temperature indicated.
3. As measured by the activation of the on-die Intel® Thermal Monitor. Intel Thermal Monitor's automatic mode is used to indicate that the maximum processor operating temperature has been reached.
4. Intel Thermal Monitor automatic mode must be enabled for the processor to operate within specifications.
5. These TDP and $T_{j\text{-}max}$ values are provided as guidance only. The lower TDP value in this table for a given processor can be used if a thermal solution with increased capability is implemented to maintain the corresponding lower Tj listed in this table.

3.2 Processor Thermal Features

3.2.1 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, power increases linearly with frequency and with the square of voltage. In the absence of power saving technologies, ever increasing frequencies 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 pursuing low-power design techniques. For example, decreasing the operating voltage, reducing unnecessary transistor activity, and using more power efficient circuits can reduce processor power consumption.

An on-die thermal management feature called the Thermal Monitor is available on the processor. This feature 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. Intel® Thermal Monitor can reduce cooling solution costs by allowing thermal designs to target the TDP.

3.2.2 Intel® Thermal Monitor Implementation

The Intel Thermal Monitor helps control the processor temperature by activating the TCC (Thermal Control Circuit) when the processor silicon reaches its maximum operating temperature. The temperature at which the Intel Thermal Monitor activates the TCC is not user configurable. Bus traffic is snooped in the normal manner and interrupt requests are latched (and serviced during the time that the clocks are on) while the TCC is active. With a properly designed and characterized thermal solution, the TCC would only be activated for very short periods of time when running the most power intensive applications. The processor performance impact due to these brief periods of TCC activation is expected to be minor and hence not detectable. An under-designed thermal solution that is not able to prevent excessive activation of the TCC in the anticipated ambient environment may cause a noticeable performance loss and may affect the long-term reliability of the processor.

In addition, a thermal solution that is significantly under designed may not be capable of cooling the processor even when the TCC is active continuously. The Intel Thermal Monitor controls the processor temperature by modulating (starting and stopping) the processor core clocks when the processor silicon reaches its maximum operating temperature. The Intel Thermal Monitor uses two modes to activate the TCC: automatic mode and on-demand mode. If both modes are activated, automatic mode takes precedence.

This processor features one automatic thermal mode of operation called Intel Thermal Monitor 1 (TM1). This mode is selected by writing values to the Model Specific Registers (MSRs) of the processor. After automatic mode is enabled, the TCC activates only when the internal die temperature reaches the maximum allowed value for operation. Intel Thermal Monitor automatic mode must be enabled through BIOS for the processor to be operating within specifications. Intel recommends TM1 be enabled on the S12x0 processor.

When TM1 is enabled and a high-temperature situation exists, the clocks are modulated by alternately turning the clocks off and on at a 50% duty cycle. Cycle times are processor speed dependent and decreases linearly as processor core frequencies increase. Once the temperature has returned to a non-critical level, modulation ceases and TCC goes inactive. A small amount of hysteresis has been included to prevent rapid active/inactive transitions of the TCC when the processor temperature is near the trip point. The duty cycle is factory configured and cannot be modified. Also, the automatic mode does not require any additional hardware, software drivers, or interrupt handling routines. Processor performance decreases by the same amount as the duty cycle when the TCC is active.

The TCC may also be activated via on-demand mode. If bit 4 of the ACPI Intel Thermal Monitor Control register is written to a 1, the TCC is activated immediately independent of the processor temperature. When using on-demand mode to activate the TCC, the duty cycle of the clock modulation is programmable via bits 3:1 of the same ACPI Intel Thermal Monitor Control register. In automatic mode, the duty cycle is fixed at 50% on, 50% off; however, in on-demand mode, the duty cycle can be programmed from 12.5% on/87.5% off, to 87.5% on/12.5% off in 12.5% increments. On-demand mode may be used at the same time automatic mode is enabled; however, if the system tries to enable the TCC via on-demand mode at the same time automatic mode is enabled and a high-temperature condition exists, automatic mode takes precedence. An external signal, PROCHOT# (processor hot), is asserted when the processor detects that its temperature is above the thermal trip point. Bus snooping and interrupt latching are also active while the TCC is active.

Besides the thermal sensor and thermal control circuit, the Intel Thermal Monitor also includes one ACPI register, one performance counter register, three MSR, and one I/O pin (PROCHOT#). All are available to monitor and control the state of the Intel Thermal Monitor feature. The Intel Thermal Monitor can be configured to generate an interrupt upon the assertion or de-assertion of PROCHOT#. PROCHOT# is not asserted when the processor is in the Stop Grant power states; hence, the thermal diode reading must be used as a safeguard to maintain the processor junction temperature within the maximum specification. If the platform thermal solution is not able to maintain the processor junction temperature within the maximum specification, the system must initiate an orderly shutdown to prevent damage. If the processor enters one of the above power states with PROCHOT# already asserted, PROCHOT# remains asserted until the processor exits the Stop Grant power state and the processor junction temperature drops below the thermal trip point. If Intel Thermal Monitor automatic mode is disabled, the processor operates out of specification. Regardless of enabling the automatic or on-demand modes, in the event of a catastrophic cooling failure, the processor automatically shut downs when the silicon has reached a temperature of approximately 125 °C. At this point, the THERMTRIP# signal goes active. THERMTRIP# activation is independent of processor activity and does not generate any bus cycles. When THERMTRIP# is asserted, the processor core voltage must be shut down within the time specified in the *Intel® Atom™ Processor S1200 Product Family for Microserver Datasheet.*

The Intel 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 TCC.
- A TCC that attempts 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.

3.2.3 PROCHOT# Signal

The primary function of the PROCHOT# signal is to provide an external indication that the processor has exceeded its maximum operating temperature. While PROCHOT# is asserted, the TCC is active. Assertion of the PROCHOT# signal is independent of any register settings within the processor. This signal 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# goes 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# activates the TCC for both cores. The TCC remains active until the system de-asserts PROCHOT#.

As an output, 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 and once configured in manufacturing process, the temperature at which the PROCHOT# signal is asserted is not re-configurable.

One possible application for PROCHOT# as an input 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), 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 specifications should rarely experience activation of the TCC as indicated by the PROCHOT# signal going active.

3.2.4 Thermal Control Circuit (TCC)

The Thermal Control Circuit portion of the Thermal Monitor must be enabled for the processor to operate within specifications. The Intel Thermal Monitor's TCC, when active, attempts to lower the processor temperature by reducing the processor power consumption. In the original implementation of Thermal Monitor, power is reduced by modulating 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 disable the internal clocks for a shorter time period. Figure 3-1 illustrates the relationship between the internal processor clocks and PROCHOT#.

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

Figure 3-1. Concept for Clocks under Thermal Monitor Control

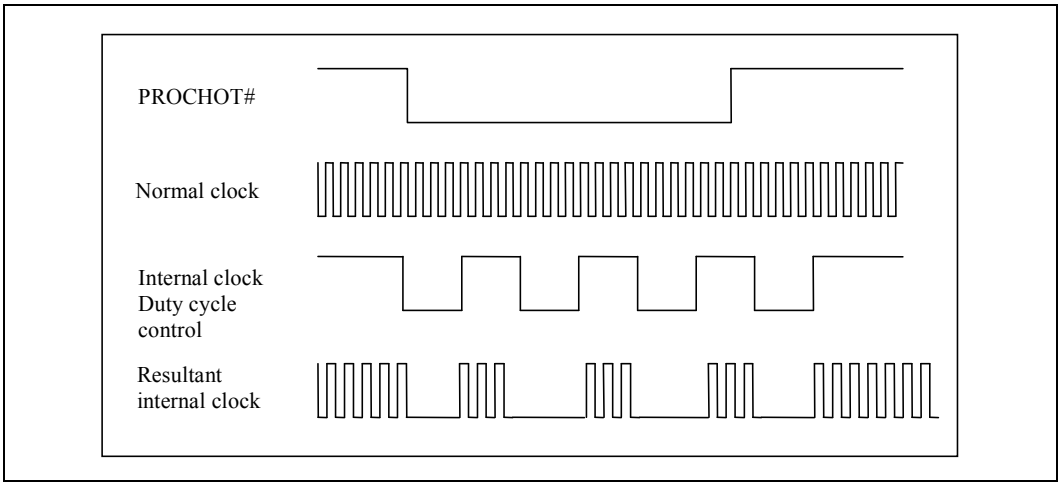

3.2.5 Operation and Configuration

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. The 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. 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 is enabled, processor power consumption is 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# always indicates the thermal status of the processor.

The power reduction mechanism of thermal monitor is also activated manually using an "on-demand" mode. Refer to Section 3.2.6, "On-Demand Mode" for details on this feature.

3.2.6 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.

3.2.7 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, compute 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 datasheet.

A system designed to meet the thermal specification of $T_{J\text{-}MAX}$ values published in the processor datasheet 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 is not capable of maintaining a safe operating temperature and the processor could shut down and signal THERMTRIP#.

3.2.8 Operating System and Application Software Considerations

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

3.2.9 THERMTRIP# Signal

In the event of a catastrophic cooling failure, the processor automatically shuts 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. The THERMTRIP# activation is independent of processor activity and does not generate any bus cycles. Refer to the processor datasheet 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.

3.2.10 Cooling System Failure Warning

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 allows the system to continue functioning or allows 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.

3.2.11 Digital Thermal Sensor

The S12x0 processor uses a Digital Thermal Sensor (DTS) as the on-die sensor for processor temperature monitoring. The DTS monitors the same sensor that activates the TCC. See Section 3.2.4, "Thermal Control Circuit (TCC)." Only the DTS is available for temperature monitoring since a thermal diode is not supported.

The S12x0 processor DTS can only be accessed via a MSR. The value read via the MSR is an unsigned number of degrees Celsius away from TCC activation temperature. The readings from the DTS are relative to the activation of the TCC which occurs at a DTS reading of zero (0). Multiple digital thermal sensors can be implemented within the package without adding a pair of signal pins per sensor as required with a thermal diode. The digital thermal sensor is easier to place in thermally sensitive locations of the processor than thermal diodes. This is achieved due to a smaller footprint 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 for each processor family.

§ §

4 Design Considerations

This section describes heatsink and system design considerations that should be understood for developing the system thermal solution.

4.1 Heatsink Design Considerations

To dissipate 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 die. One method used to improve thermal performance is by attaching a heatsink to the die. A heatsink can increase the effective heat transfer surface area by conducting heat out of the die 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 die 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 die and the bottom surface of the heatsink and improve the overall performance of the stack-up (die-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 and the pressure applied to it. Refer to Section 4.1.3, “Thermal Interface Material” for further information.
- 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. Heat transfer is characterized by the local ambient temperature of the air 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 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 are cooled by air with lower air speed. These heatsinks are typically larger (and heavier) than active heatsinks due to the increase in fin surface required to meet a required performance. As the heatsink fin density (the number of fins in a given cross-section) increases, the resistance to the airflow increases. The air likely 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.

4.1.1 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 motherboard recommendations may preclude using it in a system adhering strictly to the form factor requirements in compliance with the form factor documentation.

The following server form factor standard may be useful for determining maximum heatsink size (using footprint keep-out and height restriction specifications):

- Server System Infrastructure – ssiforum.org

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 (fan performance, for example) and other design considerations (air duct, etc.).

4.1.2 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-thermal, conductive materials like copper to increase heatsink thermal conduction performance results in physically heavier solutions. As mentioned in Section 2.4, "Package Mechanical Loading Specifications" the heatsink mass must take into consideration the package 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 ensures the system integrity under the mechanical shock and vibration profile targets may become prohibitive.

The mass limit for a reference design heatsink is based on the capabilities of reference design components that retain the heatsink to the board and apply the necessary preload. Any reuse of the clip and fastener in alternate or derivative designs should not exceed the recommended mass limit. Designs that have a mass of greater than recommended mass should analyze the preload force, retention limits of the fastener and the shock and vibration impact on the package soldered on motherboard.

4.1.3 Thermal Interface Material

Thermal interface material application between the processor die and the heatsink base is generally required to improve thermal conduction from the die 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.

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

When pre-applied material is used, Intel recommends having a protective film applied. This film must be removed prior to heatsink installation.

4.1.4 Heatsink Attach Mechanism Design Considerations

4.1.4.1 General Guidelines

There are no features on the FCBGA package for a direct heatsink attachment; 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 processor die, this mechanism plays a significant role in the thermal and mechanical robustness of the system by:

- Maintaining thermal performance of the Thermal Interface Material (TIM) applied between the processor die and the heatsink. TIMs based on phase-change materials are very sensitive to applied pressure: the higher the pressure, the better the initial performance. Designs should incorporate a possible decrease in applied pressure over time due to potential structural relaxation in retention components (creep effect causing the clip to lose its preload and causing anchor pull-out). Intel does not recommend utilizing TIMs such as thermal greases onto small-bare die package, due to the TIM pump-out concern after heatsink is assembled.
- 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 should be considered in designing the heatsink attach mechanism. The design should provide a means for protecting the solder joints.

4.1.4.2 Attach Mechanism Design Considerations

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

- Solder joint reliability compliant with Intel quality specification before and after a reliability tests such as shock and vibration. The Critical-to-Function (CTF) corner solder joints of the processor package might experience high-stress concentration during shock and vibration testing. The retention design is critical to prevent solder joints failures.
- The heatsink should be held in place under mechanical shock and vibration events. The applied force to the heatsink base should maintain the desired pressure on the thermal interface material.

Note: The load applied by the heatsink attach mechanism must comply with the package specifications described in this document.

The height of the package (the package seating plane to the top of the die) needs to account for nominal variation and tolerances of the manufacturing process. It must:

- Engage 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.
- Minimize contact with the motherboard surface during installation and actuation to avoid scratching/damaging the motherboard.
- Account for the associated heatsink keep-out zone.
- Not use a board stiffening device (backing plate, chassis attach, etc.).

The load generated by the attach mechanism is designed to help protect the package against solder joint risk in dynamic loading (e.g., shock and vibration) and fatigue loading experienced in temperature cycling.

Note: The loading specifications in Section 2.4, "Package Mechanical Loading Specifications" are required to be followed while simultaneously ensuring a minimum load is applied throughout the product life. The tolerance and nominal load is based on the reference design and will differ for an alternate thermal solution provided by a third party.

Taking into account potential load degradation from creep over time is important when designing the clip or fastener to the required minimum load. The initial preload, depending on the clip stiffness, at the beginning of the life of the product may be significantly higher than the minimum preload that must be met throughout the life of the product. Refer to Appendix C, "Heatsink Retention Load Metrology" for clip load metrology guidelines.

4.2 System Thermal Solution Considerations

4.2.1 Chassis Thermal Capabilities

The microserver platform reference thermal solution for the S12x0 processor is a passive heatsink that requires an adequate system thermal design to maintain thermal specifications. The system must provide an adequate airflow rate corresponding to the local ambient temperature at the heatsink. Refer to Section 4.2.2, "Improving Chassis Thermal Performance" for more chassis design details.

To evaluate the system thermal capability of a given chassis, conduct an in-chassis system thermal test.

In a system using customized thermal solution on the processor, the thermal pass requirement for a given chassis can be met, if

$$T_J \leq T_{J\text{-}MAX}$$

4.2.2 Improving Chassis Thermal Performance

The heat generated by components within the chassis must be removed to provide an adequate operating environment for the processor and all other components in the system. Moving airflow through the chassis brings in fresh cool 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/or vents determine the chassis thermal performance, and the resulting ambient temperature around the processor.

4.2.3 Summary

In summary, heatsink design considerations for the S12x0 processor include:

- When implementing any customized thermal solution using either a thermal interface material, or clip, or heatsink design that is different from the Intel reference thermal solution described in this document, Intel recommends using the junction temperature ($T_{J\text{-}MAX}$) for the processor as the criteria to judge the thermal performance of the processor when operated at TDP for an external ambient temperature of 35 °C. Compliance relative to $T_{J\text{-}MAX}$ at all times is needed to ensure processor reliability.
- Heatsink interface-to-die surface characteristics such as flatness and roughness influence the contact resistance of the heatsink to the die.
- The resistance of the thermal interface material used between the heatsink and the die is a large portion of the heatsink performance. Intel recommends use of Honeywell* PCM45F with sufficiently high preload applied by the retention mechanism.
- The surface area of the heatsink should be adequate to provide sufficient convective cooling in a forced convection environment. The reference design maximizes the fin area while minimizing the weight of the heatsink.
- While the material used to construct the heatsink is generally important, the expected operating condition of the reference design heatsink is such that aluminum is more than adequate for proper conduction in the heatsink. At low flow rates, the resistance of the heatsink is largely limited by its convective resistance rather than conduction.
- Proper venting is required on the chassis to ensure sufficient air flow for the passive heatsink and other thermal critical components within the system.

§ §

5 Reference Thermal Solutions

This section describes the overall requirements for the platform heatsink reference thermal solutions supporting the S12x0 processor. A heat spreader concept is also included which may offer a low-cost approach for cooling the S12x0 power-optimized processor.

Intel has developed reference thermal solutions to meet the cooling needs of the S12x0 processor under operating environments and specifications defined in this document. This section describes the overall requirements for the heatsink reference thermal solution including critical-to-function dimensions, operating environment, and validation criteria. Other chipset components may or may not need attached thermal solutions depending on the specific system local-ambient operating conditions.

Table 5-1. Recommended Reference Thermal Solutions for S12x0 Dual-Core Processors

Processor	Reference Thermal Solution[1]	Description
Intel® Atom™ Processor S1240	Low-profile torsional clip heatsink	40 mm x 40 mm, 7 mm tall Aluminium 6063 T5, PCM45F TIM
	Thin Heat Spreader (THS) concept[2]	25 mm x 25 mm, 0.33 mm thick Copper
Intel® Atom™ Processor S1220 and S1260	Low-profile torsional clip heatsink	40 mm x 40 mm, 7 mm tall Aluminium 6063 T5, PCM45F TIM

Notes:
1. The reference thermal mechanical solution information shown in this document represents the current state of the data and may be subject to changes. The information represents design targets and not commitments by Intel.
2. The THS is presented here as a reference concept only for system-design consideration. The system integrator is encouraged to work with a heatsink vendor to fully develop a similar thermal solution.

5.1 Low-Profile Torsional Clip Heatsink Reference Thermal Solution

This section describes a reference thermal solution for the S12x0 processor using an aluminum heatsink.

5.1.1 Operating Environment

The reference thermal solution was designed assuming a high-fan speed, a maximum local-ambient temperature of 50 °C, and the minimum recommended airflow velocity (1.2 m/s).

The approaching airflow temperature is assumed to be equal to the local-ambient temperature. The thermal designer must select the location to measure airflow to obtain an accurate estimate. The local-ambient conditions (Table 5-2) are based on a 35 °C (high-fan speed condition) external-ambient temperature at an altitude of 1500 meters. (External-ambient refers to the environment external to the system.)

Table 5-2. S12x0 Processor Operating Conditions

	High Speed
Velocity (m/sec)	1.2
Altitude (m)	1500
Tsystem_ambient (^{o}C)	35
Trise (^{o}C)	15
T_A (^{o}C)	50

5.1.2 Heatsink Performance

Figure 5-1 depicts the simulated thermal performance (Psi-ja) of the reference thermal solution versus average air velocity through the fins (in m/s). Since the simulation was performed for sea level, a correction would be required to estimate thermal performance at 1500 m altitude.

The following curve fit equation can be used to determine Psi-ja for sea level based on expected system airflow:

$$\textbf{Psi-ja} = 0.9346 + 3.559 \times \textbf{Velocity}^{-0.4332}$$

The following equation with altitude correction can be used for 1500 m altitude:

$$\textbf{Psi-ja} = 0.9346 + 3.932 \times \textbf{Velocity}^{-0.4332}$$

The system integrator then determines whether the $T_{J\text{-}MAX}$ specification can be met with the expected local ambient temperature:

$$T_J = \textbf{Psi-ja} \times \textbf{TDP} + T_A$$

Figure 5-1. Low-Profile Torsional Clip Heatsink Measured Thermal Performance vs. Average Velocity at Sea Level

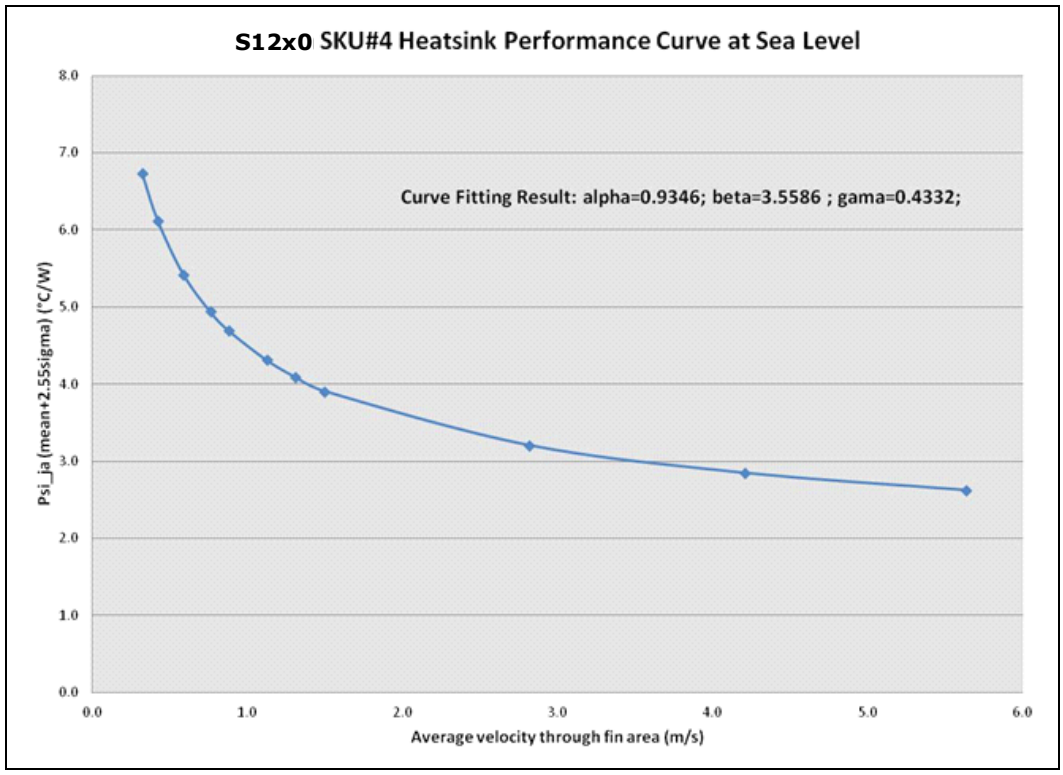

5.1.3 Mechanical Design Keep-Out Zones

The location of hole patterns and keep-out zones for the reference thermal solution are available in Appendix A, "Mechanical Drawings" for two possible orientations of the package relative to the heatsink assembly. The heatsink assembly volumetric envelope is also included in Appendix A, "Mechanical Drawings."

5.1.4 Torsional Clip Heatsink Thermal Solution Assembly

The reference thermal solution is a passive extruded heatsink with thermal interface material. It is attached using a clip with each end hooked through an anchor soldered to the board. Figure 5-2 and Figure 5-3 on page 37 shows the reference thermal solution assembly and associated components.

Full-mechanical drawings of the thermal solution assembly and the heatsink clip are provided in Appendix A, "Mechanical Drawings." Appendix B, "Thermal Solution Supplier Information" contains vendor information for thermal solution assembly.

Figure 5-2. Low-Profile Torsional Clip Heatsink Assembly (View 1)

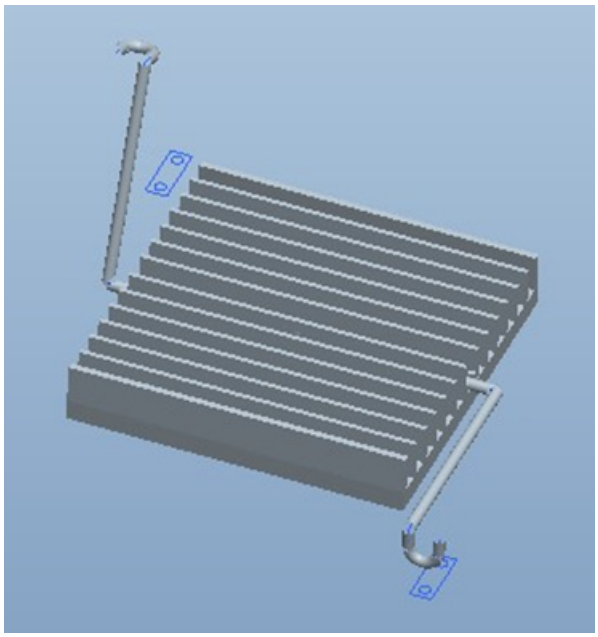

Figure 5-3. Low-Profile Torsional Clip Heatsink Assembly (View 2)

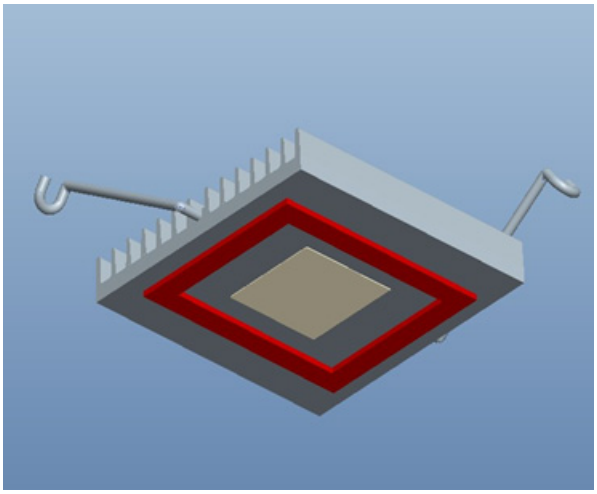

5.1.5 Extruded Heatsink Profiles

The reference thermal solution uses an extruded heatsink for cooling the processor. A full-mechanical drawing of this heatsink is provided in Appendix A, "Mechanical Drawings."

5.1.6 Mechanical Interface Material

The low-profile torsional clip heatsink uses a 34 mm x 28 mm foam pad to provide a protective mechanical interface between the package and heatsink. This pad helps reduce heatsink tilt during installation and protect package features and components from damage. The interior cutout of the pad (27.5 mm x 21.5 mm) is sized to avoid any interference with package top-surface components.

5.1.7 Thermal Interface Material

A Thermal Interface Material (TIM) provides improved conductivity between the die surface and heatsink. The reference thermal solution uses Honeywell* PCM45F, 0.25 mm (0.010 in.) thick, 15 mm x 15 mm (0.6 in. x 0.6 in.) square.

Note: Unflowed or dry Honeywell PCM45F has a material thickness of 0.010 inches. The flowed or wet Honeywell PCM45F has a material thickness of ~0.003 inches after it reaches its phase change temperature.

5.1.7.1 Effect of Pressure on TIM Performance

As mechanical pressure increases on the TIM, the thermal resistance of the TIM decreases. This phenomenon is due to the decrease of the Bond-Line Thickness (BLT). BLT is the final settled thickness of the thermal interface material after installation of heatsink. The effect of pressure on the thermal resistance of the Honeywell PCM45F TIM is shown in Table 5-3.

Intel provides both end-of-line and end-of-life TIM thermal resistance values of the Honeywell PCM45F. The end-of-line value represents the TIM performance post heatsink assembly while the end-of-life value is the predicted TIM performance when the product and TIM reaches the end of its life. The heatsink clip provides enough pressure for the TIM to achieve end-of-line thermal impedance of 0.19 °C-cm^2/W and end-of-life thermal impedance of 0.39 °C-cm^2/W.

Table 5-3. Honeywell PCM45F TIM Performance as a Function of Attach Pressure

Pressure on Thermal Solution and Package Interface (PSI)	Thermal Impedance (°C-cm^2)/W	
	End of Line	End of Life
40	0.19	0.39

5.1.8 Heatsink Clip

The reference solution uses a wire clip with hooked ends. The hooks attach to wire anchors to fasten the clip to the board. See Appendix A, "Mechanical Drawings" for a mechanical drawing of the clip.

5.1.9 Clip Retention Anchors

For S12x0 processor-based platforms that have very limited board space, a clip retention anchor has been developed to minimize the impact of clip retention on the board. Based on a standard three-pin jumper, this anchor is soldered to the board like any common through-hole header. A new anchor design is available with 45° bent leads to increase the anchor attach reliability over time. See Appendix B, "Thermal Solution Supplier Information" for part number and supplier information.

5.1.10 Heatsink Clip Load Requirement

The platform-reference-design-attach mechanism for the heatsink creates a nominal static compressive preload on the package of 11.3 lbf ± 3.6 lbf throughout the life of the product. The reference-design-attach mechanism also assumes the following important mechanical characteristics:

- Utilizing TIM Honeywell PCM45F (pad version) or alternate TIM with similar thickness.
- Attach pattern per keep-out zone drawing in Appendix A, "Mechanical Drawings."

5.1.11 Reliability Guidelines

Each motherboard, heatsink, and attach combination may vary the mechanical loading of the component. Based on the end-user environment, define the appropriate reliability test criteria and carefully evaluate the completed assembly prior to use in high volume.

The test profiles for the S12x0 processor-reference solutions are unpackaged system-level limits. The reference solution is to be mounted to a fully-configured system. The environmental reliability requirements for the reference-thermal solution are shown in Table 5-4. These could be considered as general guidelines.

Table 5-4. Reliability Guidelines

Test (1)	Example of Test Description	Pass/Fail Criteria (2)
Mechanical Shock	System-level unpackaged test: • Two drops for + and - directions in each of three perpendicular axes (i.e., total 12 drops). • Profile: 25g, Trapezoidal waveform, velocity change depending on system weight.	Visual Check and Electrical Functional Test
Random Vibration	System-level unpackaged test: • Duration: 10 min/axis, 3 axes. • Frequency Range: 0.002 g2/Hz @ 5 Hz, ramping to 0.01 g2/Hz @ 20 Hz, 0.01 g2/Hz @ 20 Hz to 500 Hz. • Power Spectral Density (PSD) Profile: 2.20g RMS.	Visual Check and Electrical Functional Test
Temperature Cycle	–40 °C to +85 °C	Visual Check and Electrical Functional Test

Notes:
1. Intel recommends that the above tests be performed on a sample size of at least 12 assemblies from three lots of material.
2. Additional pass/fail criteria may be added at the discretion of the user.

5.2 Thin-Heat Spreader Thermal Solution Concept

A low-profile Thin-Heat Spreader (THS) may provide sufficient processor cooling for low-power applications. This section provides design data for an S12x0 power-optimized processor thermal solution using a THS concept. The THS solution provides the advantage of a low profile and elimination of the retention clip and solder-in anchors. Elimination of the clip and solder-in anchors provides additional board space for components and routing. The concept presented here is based on thermal simulation results and is intended to guide system integrators in developing a low-power, cost-optimized solution in conjunction with a heatsink vendor. Off the shelf heat spreader solutions may exist that can meet product thermal and reliability requirements.

5.2.1 Description

The THS concept consists of a flat piece of copper measuring 25 mm x 25 mm x 0.33 mm. For simulation purposes, it is assumed to be attached to the processor die top surface using an adhesive TIM with thermal conductivity of ~2 W/m-K and a bond-line thickness of 0.1 mm.

5.2.2 THS Thermal Performance for the S12x0 Power-Optimized Processor

Thermal simulation using a 25 mm x 25 mm x 0.33 mm thick copper heat spreader concept yielded the following thermal performance is shown in Figure 5-4.

The following curve fit equation can be used to determine Psi-ja for sea level based on expected system airflow:

$$\text{Psi-ja} = 9.9938 \times \text{Velocity}^{-0.48}$$

The system integrator then determines whether the $T_{J\text{-}MAX}$ specification can be met with the expected local ambient temperature:

$$T_J = \text{Psi-ja} \times \text{TDP} + T_A$$

Figure 5-4. Intel® Atom™ Processor S1220 THS Thermal Performance

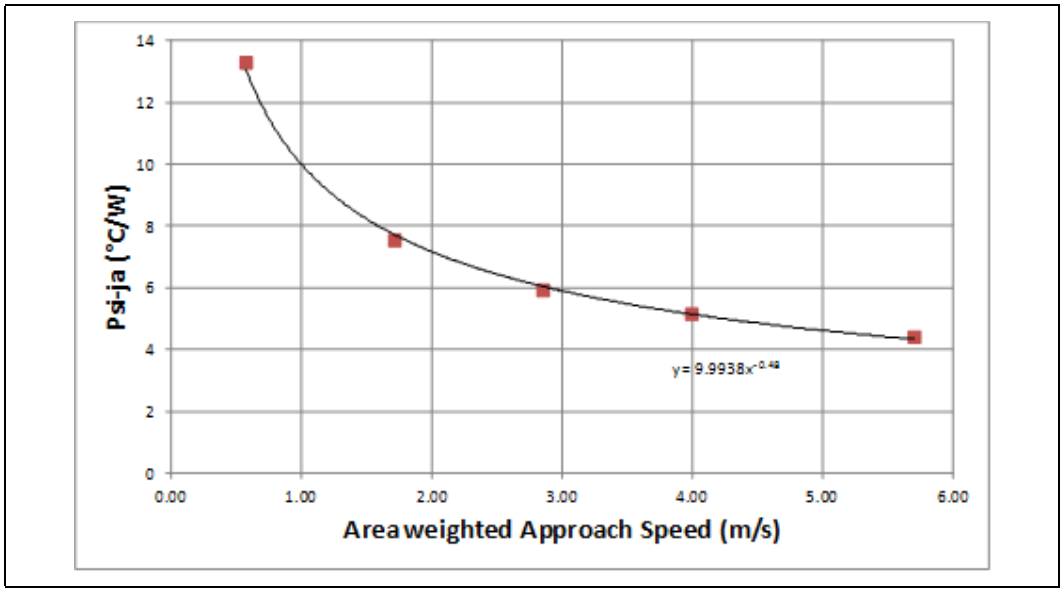

5.2.3 THS Thermal and Mechanical Reliability

The use of adhesive TIM attach may require additional testing to assess thermal and mechanical performance at end-of-life conditions, i.e., lifetime environmental stress. The thermal interface material must provide adequate thermal performance and THS mechanical attach over the operating life of the processor. System integrators should work with a heatsink vendor to develop an adequate adhesive TIM if desired.

§ §

6 System Thermal/Mechanical Design Information

6.1 Overview of the Design Requirements

This chapter documents the requirements for designing a passive heatsink that meets the maximum usage power consumption mentioned in Section 3.1, "Thermal Specifications." The platform reference thermal solution satisfies the specified thermal requirements for the S12x0 processor.

6.1.1 Altitude

Products may be required to function reliably at high altitude, typically 1,500 m [5,000 ft.] or more. Air-cooled temperature calculations and measurements at the test site elevation 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 $T_{j\text{-}max}$ requirement for the processor is met at the targeted altitude for the reference thermal solution.

6.1.2 Heatsink Thermal Validation

Intel recommends evaluation of the heatsink within the product specific boundary conditions by measuring T_J using the DTS and ensuring $T_{J\text{-}MAX}$ is not exceeded. For more information on the DTS, refer to Section 3.2.11, "Digital Thermal Sensor" and Appendix D, "Thermal Metrology."

6.2 Environmental Reliability Testing

6.2.1 Structural Reliability Testing

Structural reliability tests consist of unpackaged, board-level vibration and shock tests of a given thermal solution in the assembled state. The thermal solution should meet the specified thermal performance targets after these tests are conducted; however, the test conditions outlined here may differ from the designer's system requirements.

6.2.1.1 Random Vibration Test Procedure

- Duration: 10 min/axis, 3 axes
- Frequency Range: 5 Hz to 500 Hz
- Power Spectral Density (PSD) Profile: 3.13 GRMS

Figure 6-1. Random Vibration PSD

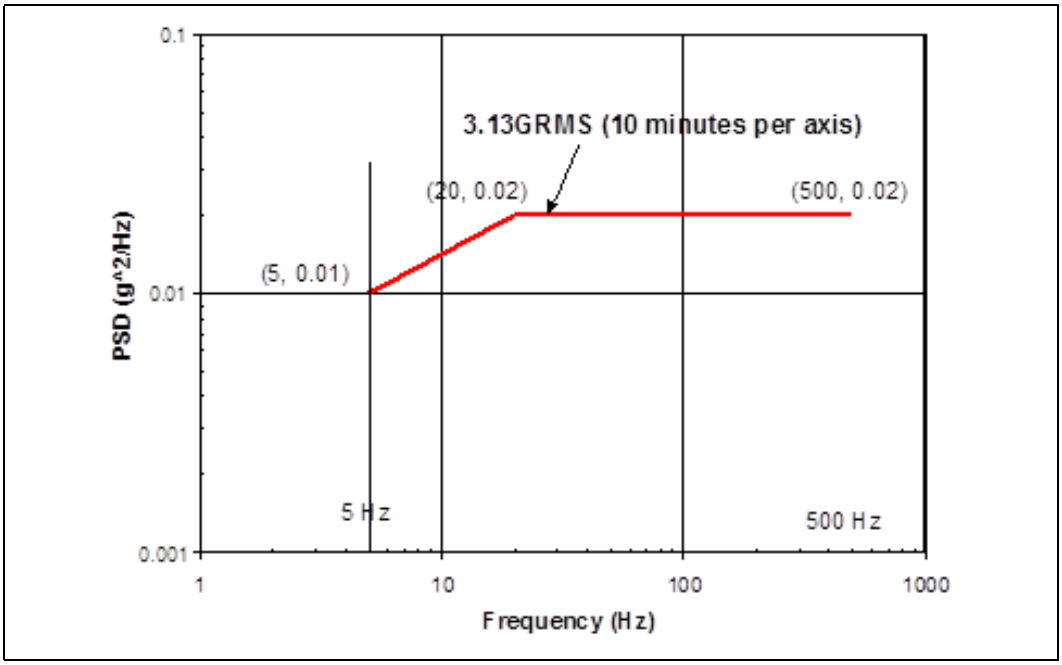

6.2.1.2 Shock Test Procedure

Recommended performance requirements for a motherboard:

- Quantity: Three drops for + and - directions in each of three perpendicular axes (i.e., total 18 drops).
- Profile: 50G trapezoidal waveform, 170 in./sec. minimum velocity change.
- Setup: Mount sample board on test fixture.

Figure 6-2. Shock Acceleration Curve

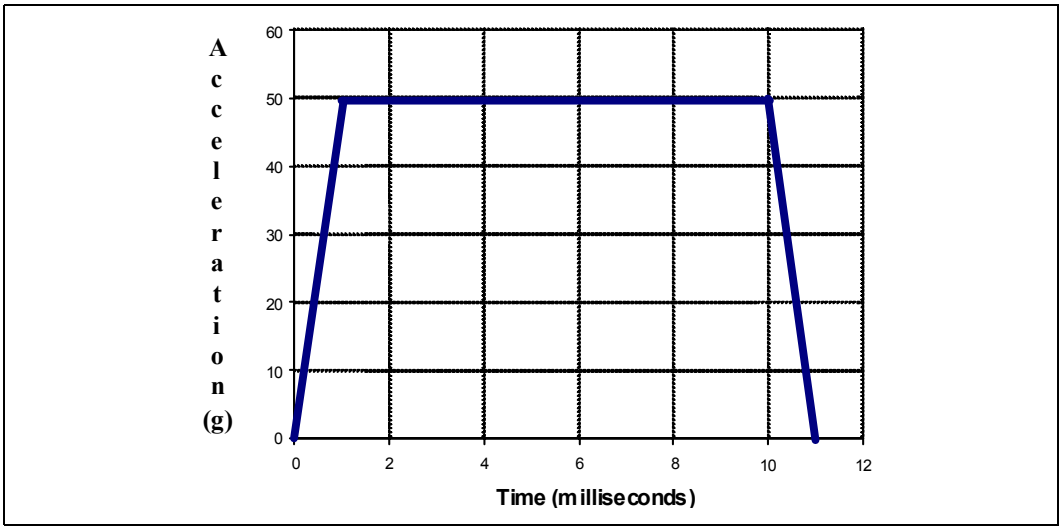

6.2.1.2.1 Recommended Test Sequence

Each test sequence should start with components (i.e., motherboard, heatsink assembly, etc.) that have never been previously submitted to any reliability testing.

The test sequence should always start with a visual inspection after assembly and BIOS/CPU/Memory test. Refer to Section 6.2.1.2.2, "Post-Test Pass Criteria."

Prior to the mechanical shock and vibration test, the units under test should be pre-conditioned for 72 hours at 45 ºC. The purpose is to account for load relaxation during burn-in stage.

Stress test should be followed by a visual inspection and then BIOS/CPU/Memory test.

6.2.1.2.2 Post-Test Pass Criteria

The post-test pass criteria include:

1. No significant physical damage to the heatsink attach mechanism (including such items as clip and motherboard fasteners).
2. The heatsink must remain attached to the motherboard.
3. The heatsink remains seated and its bottom remains mated flat against the die surface: no visible gap between the heatsink base and processor die and no visible tilt of the heatsink with respect to the attach mechanism.
4. No signs of physical damage on the motherboard surface due to impact of the heatsink or the heatsink attach mechanism.
5. No visible physical damage to the processor package.
6. Successful BIOS/CPU/memory test of post-test samples.
7. Thermal compliance testing to demonstrate that the case temperature specification can be met.

6.2.2 Recommended BIOS/CPU/Memory Test Procedures

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

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

- Appropriate system motherboard
- Processor
- All enabling components including thermal solution parts
- Power supply
- Disk drive
- Add-in card
- DIMM
- Keyboard
- Monitor

The pass, the system under test successfully completes the checking of BIOS, basic processor functions, and memory without any errors.

6.3 Material and Recycling Requirements

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

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

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

6.4 Safety Requirements

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

- UL Recognition approved for flammability at the system level. All mechanical and thermal enabling components must be a minimum UL94V-2 approved.
- CSA Certification. All mechanical and thermal enabling components must have CSA certification.
- All components (the heatsink fins) must meet the test requirements of UL1439 for sharp edges.
- If the International Accessibility Probe specified in IEC 950 can access the moving parts of the fan, consider adding a safety feature so that there is no risk of personal injury.

6.5 Reference Attach Mechanism

6.5.1 Structural Design Strategy

Structural design strategy for the Intel reference thermal solution is to minimize upward and downward board deflection during shock test.

The design uses a high-clip stiffness that resists local-board curvature under the heatsink and minimizes upward board deflection.

6.5.2 Mechanical Interface to the Reference Attach Mechanism

The attach mechanism component from the reference thermal solution can be used by other third-party cooling solutions. The attach mechanism consists of:

- 4 fasteners (for example: ITW P/N:83FT-02-37-9909)

§ §

A Mechanical Drawings

Table A-1 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 A-1. Mechanical Drawings

Figure A-1. S12x0 Processor Package Drawing

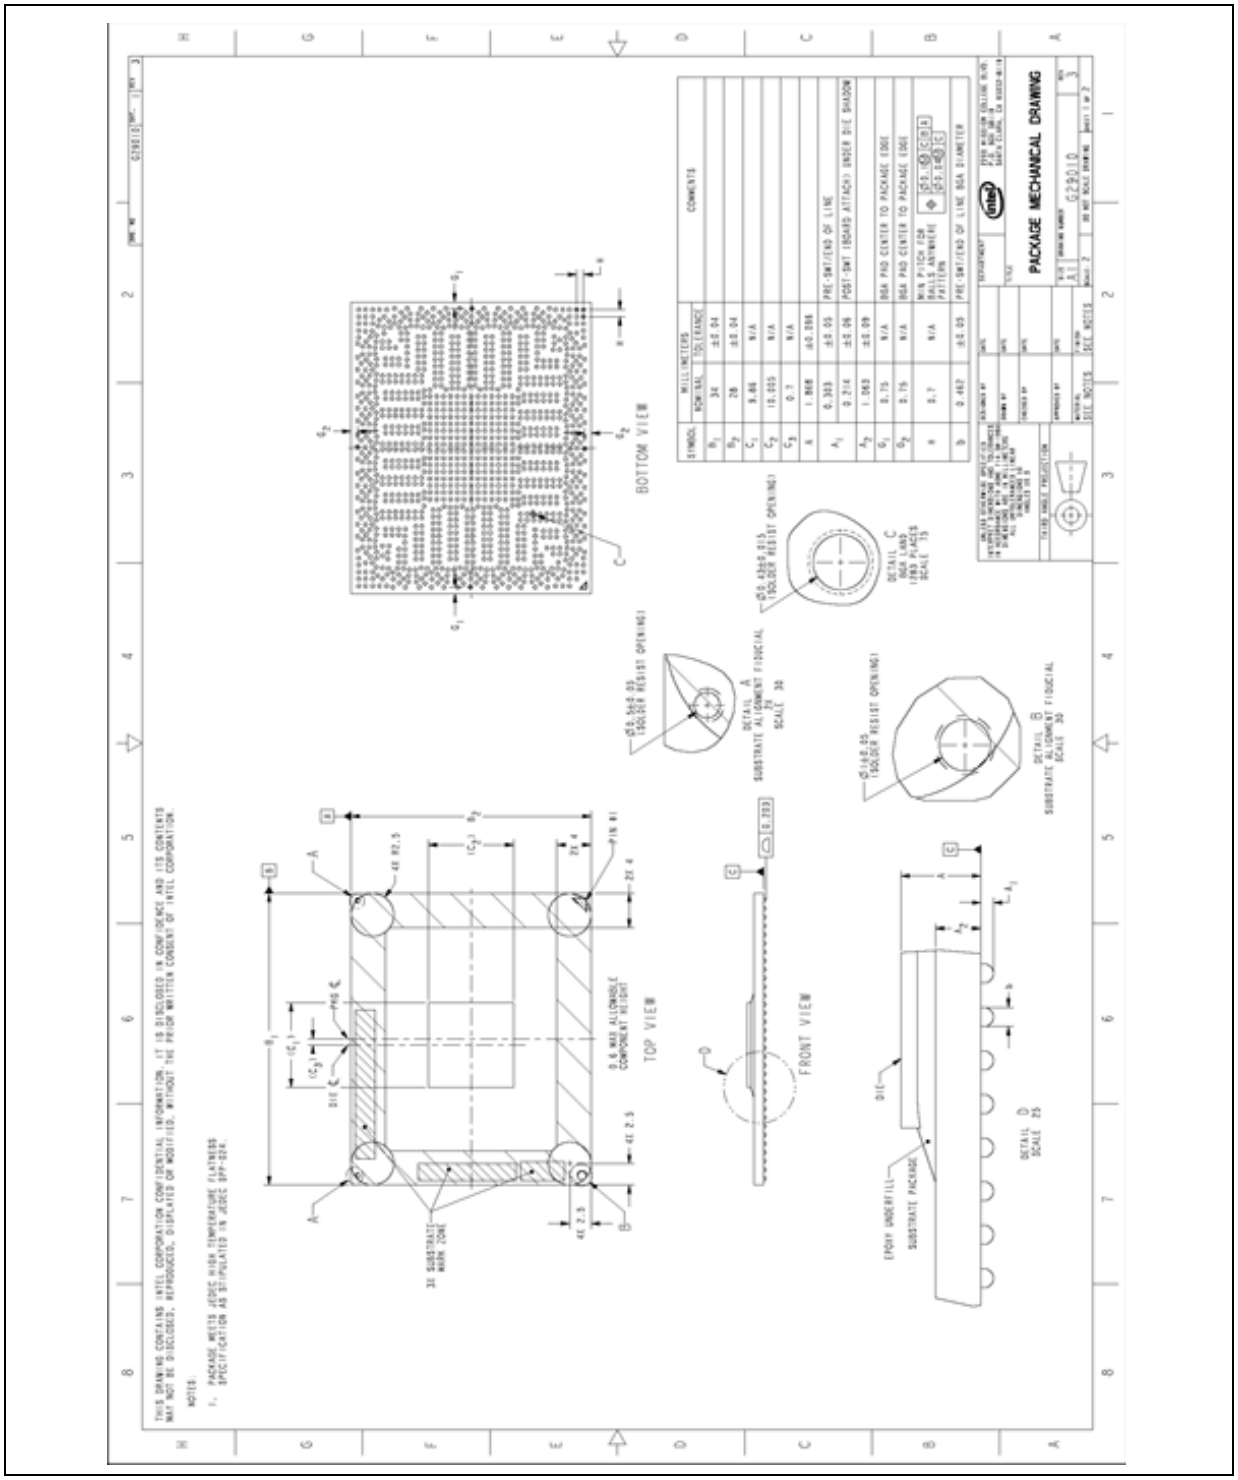

Note: Refer to the *Intel® Atom™ Processor S1200 Product Family for Microserver Processor Solder Ball Location Coordinates* spreadsheet for solder ball location coordinates.

Figure A-2. Low-Profile Torsional Clip Heatsink Assembly Drawing

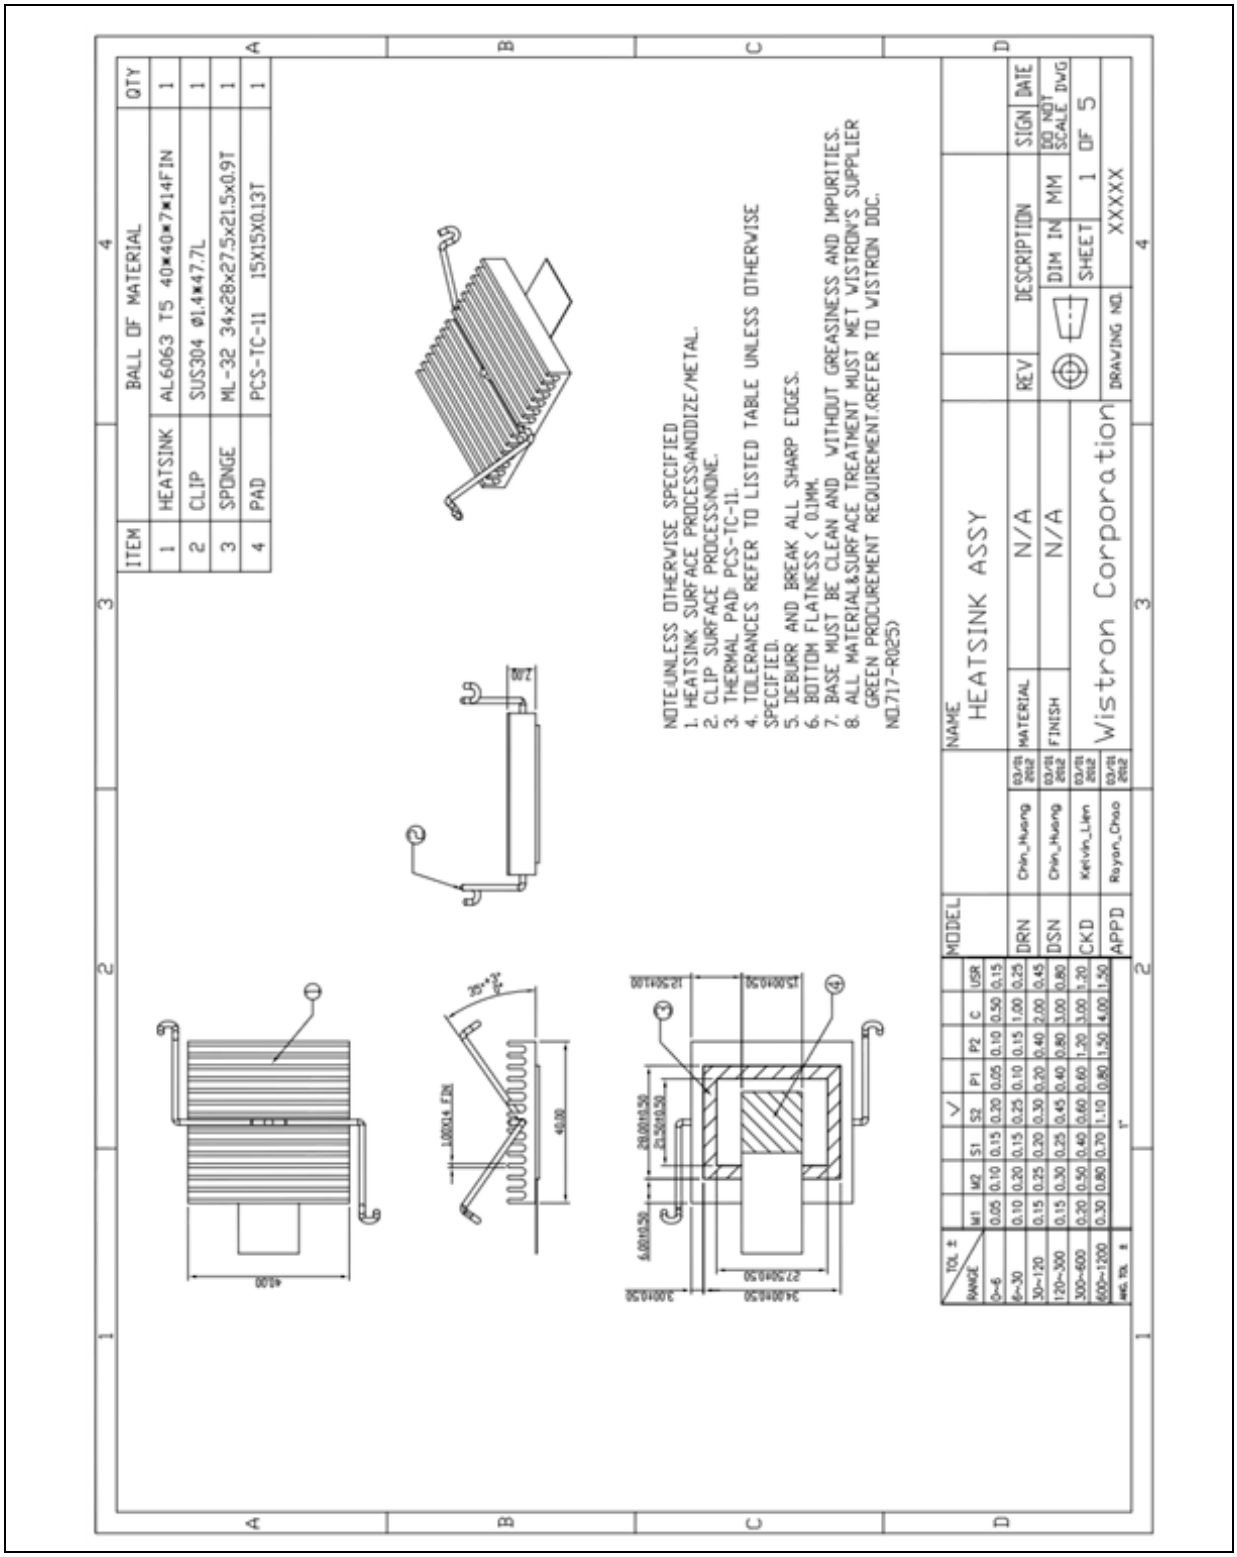

Figure A-3. Low-Profile Torsional Clip Heatsink Drawing

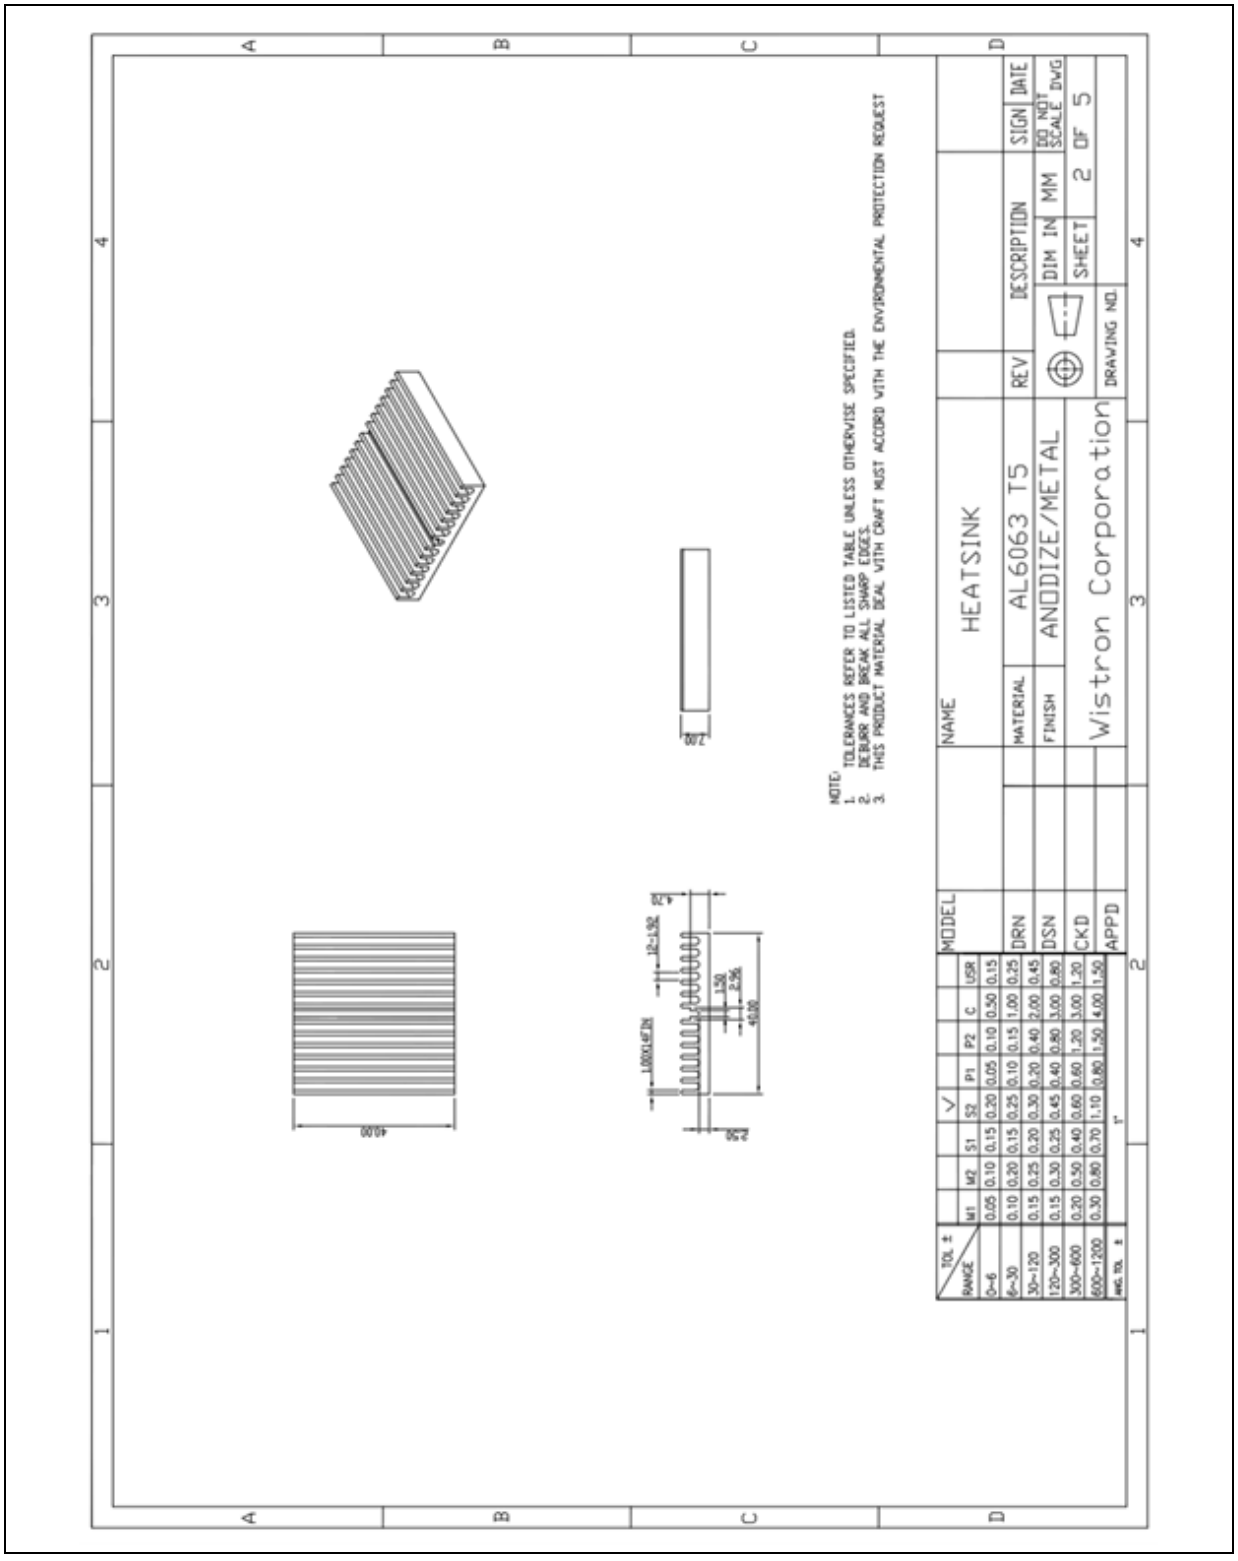

Figure A-4. Low-Profile Torsional Clip Drawing

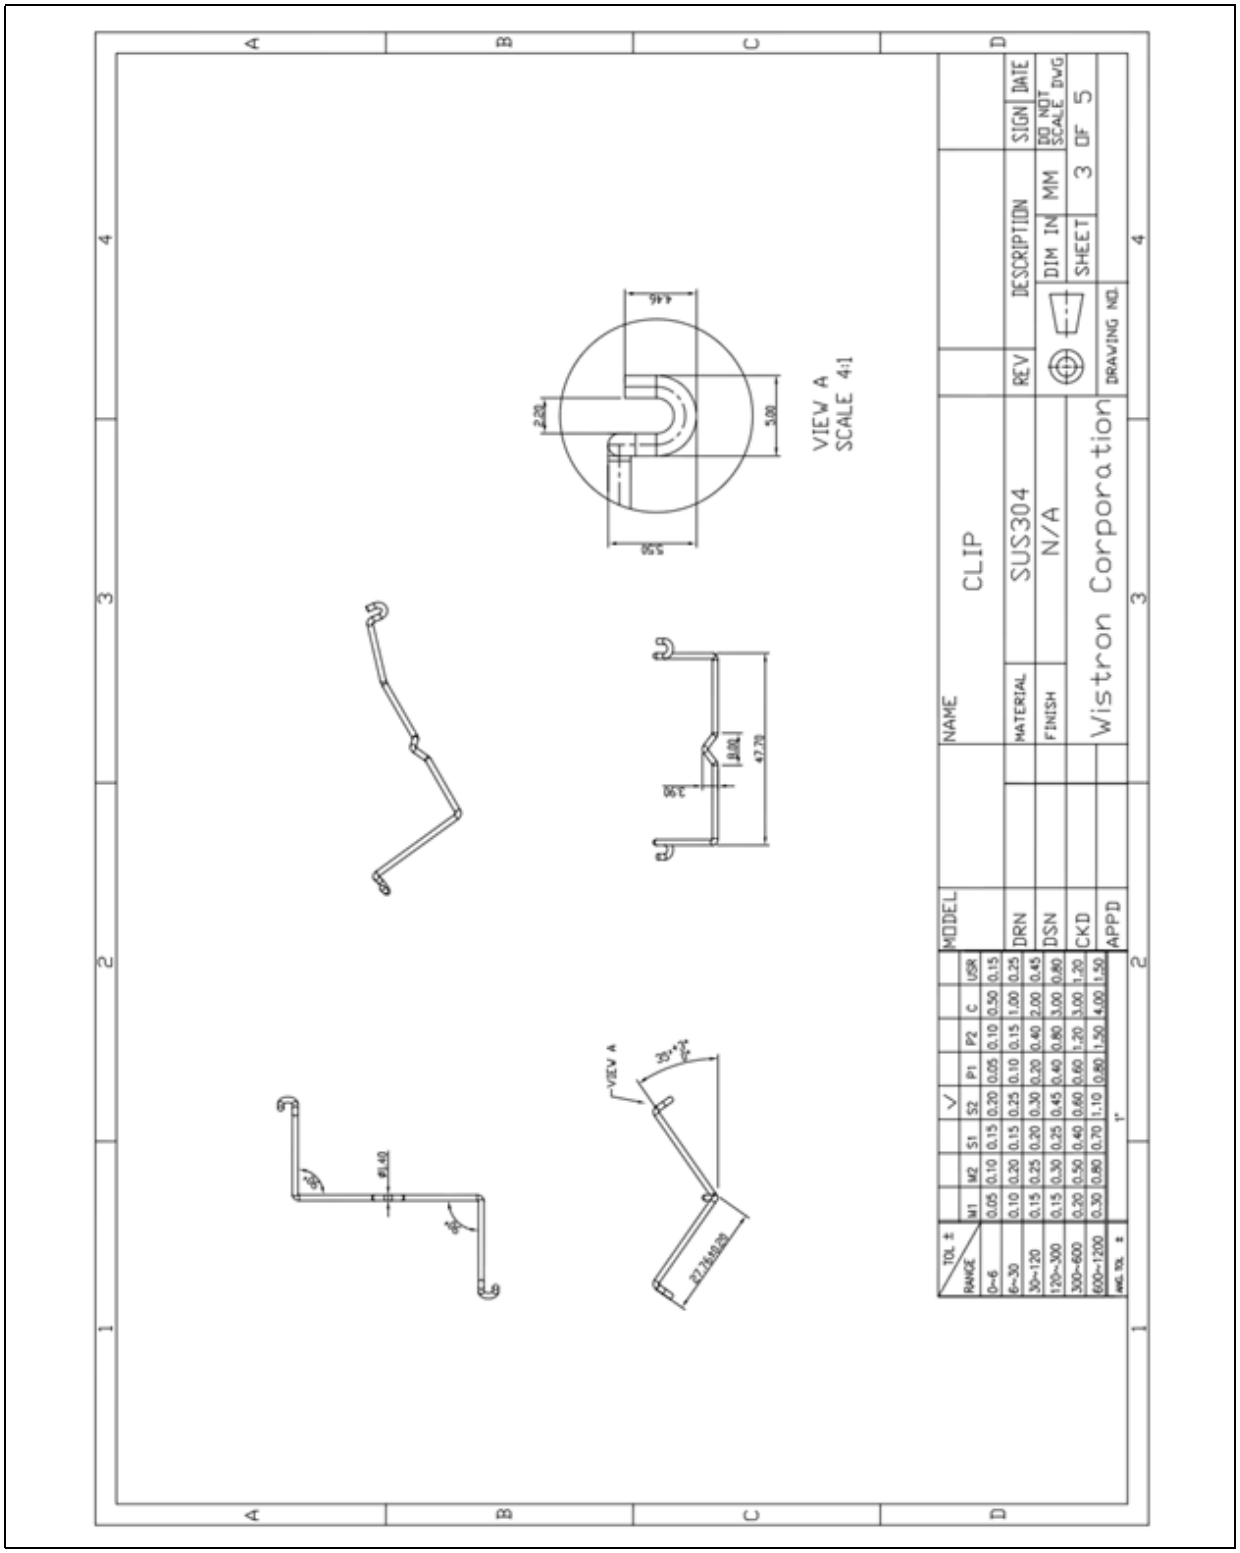

Figure A-5. Low-Profile Torsional Clip Heatsink Mechanical Interface Pad Drawing

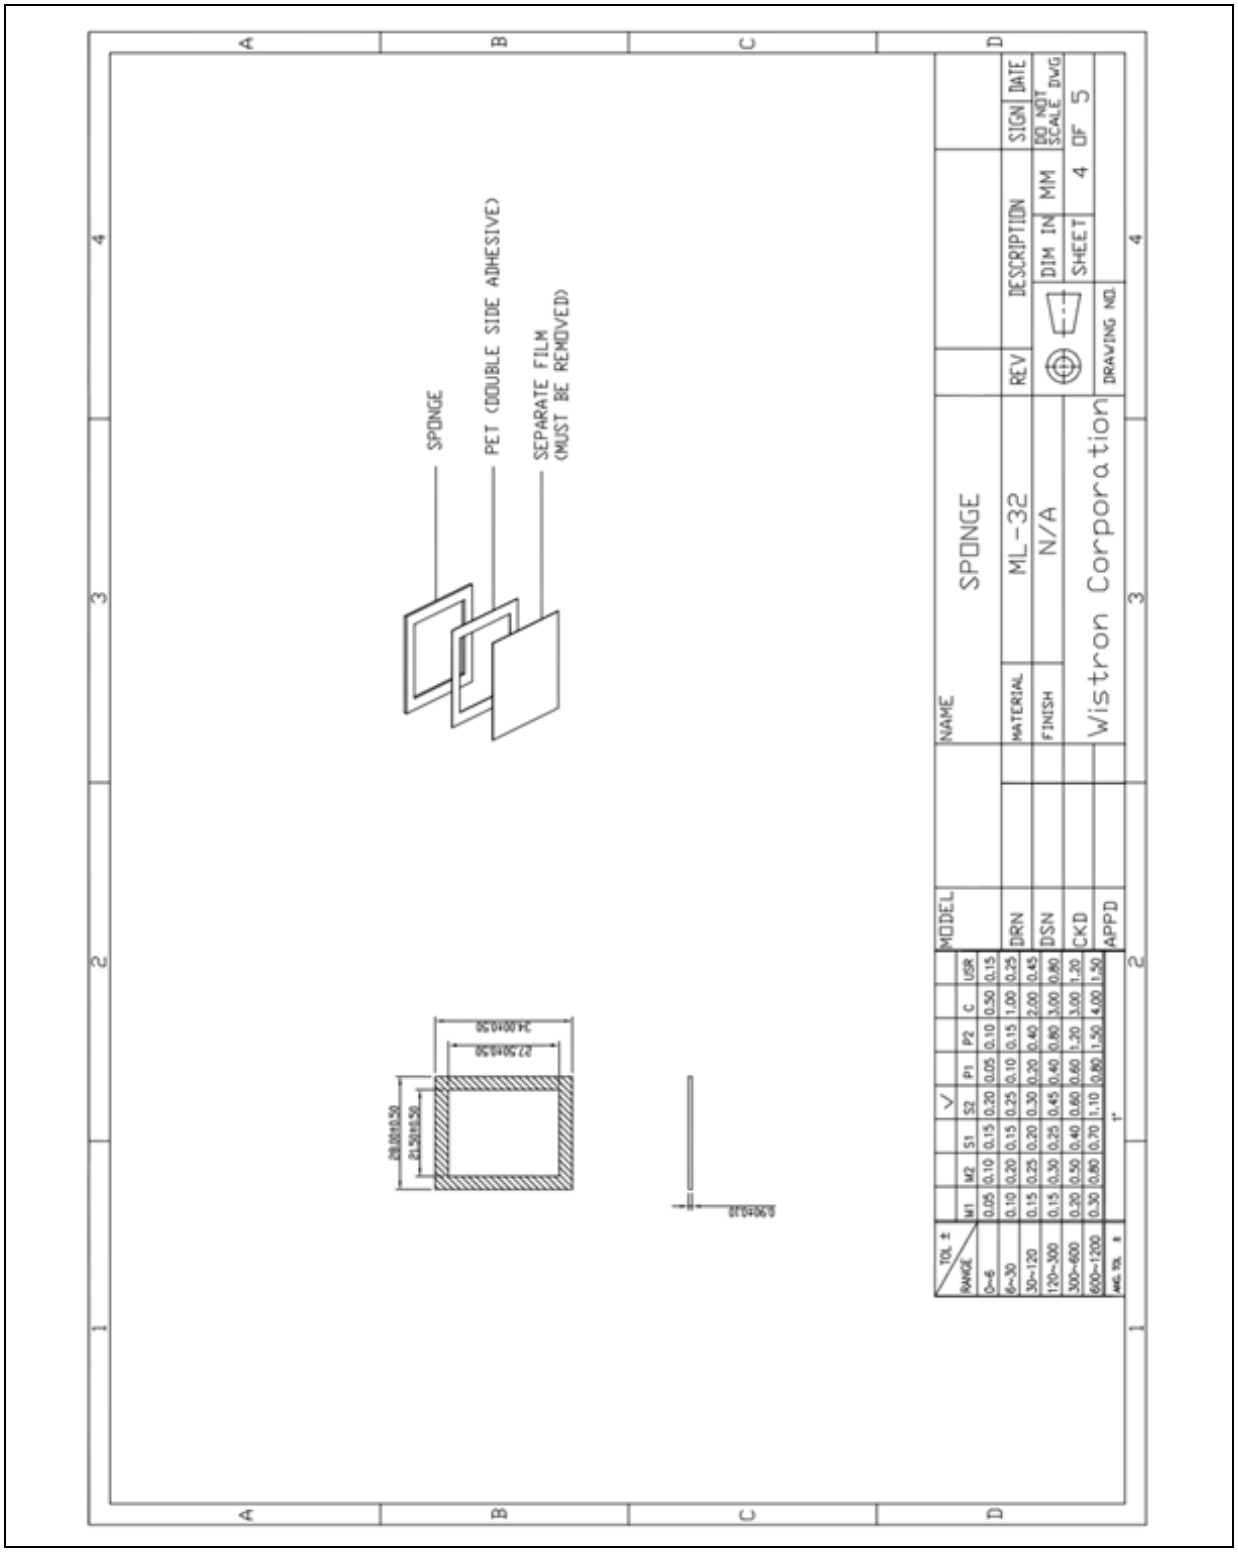

Figure A-6. Low-Profile Torsional Clip Heatsink TIM Pad Drawing

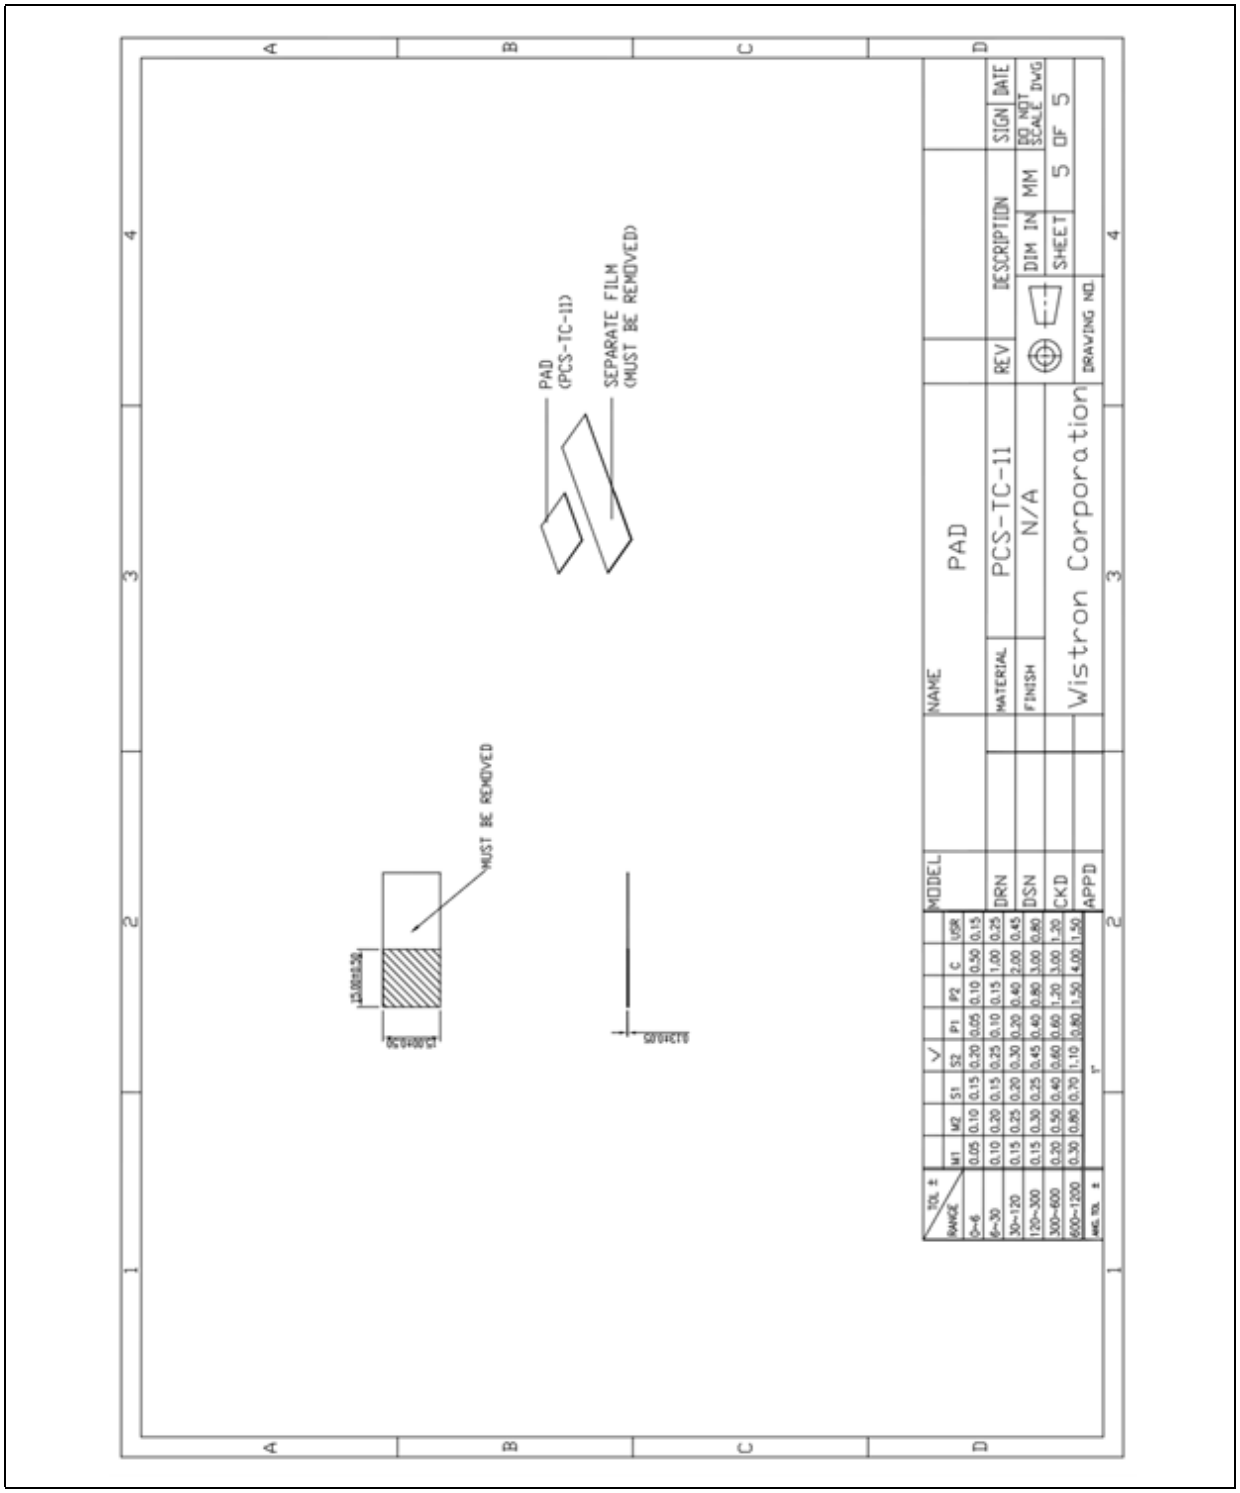

Note: Recommend using Honeywell* PCM45F.

Figure A-7. Low-Profile Torsional Clip Heatsink Keep-Out Zones

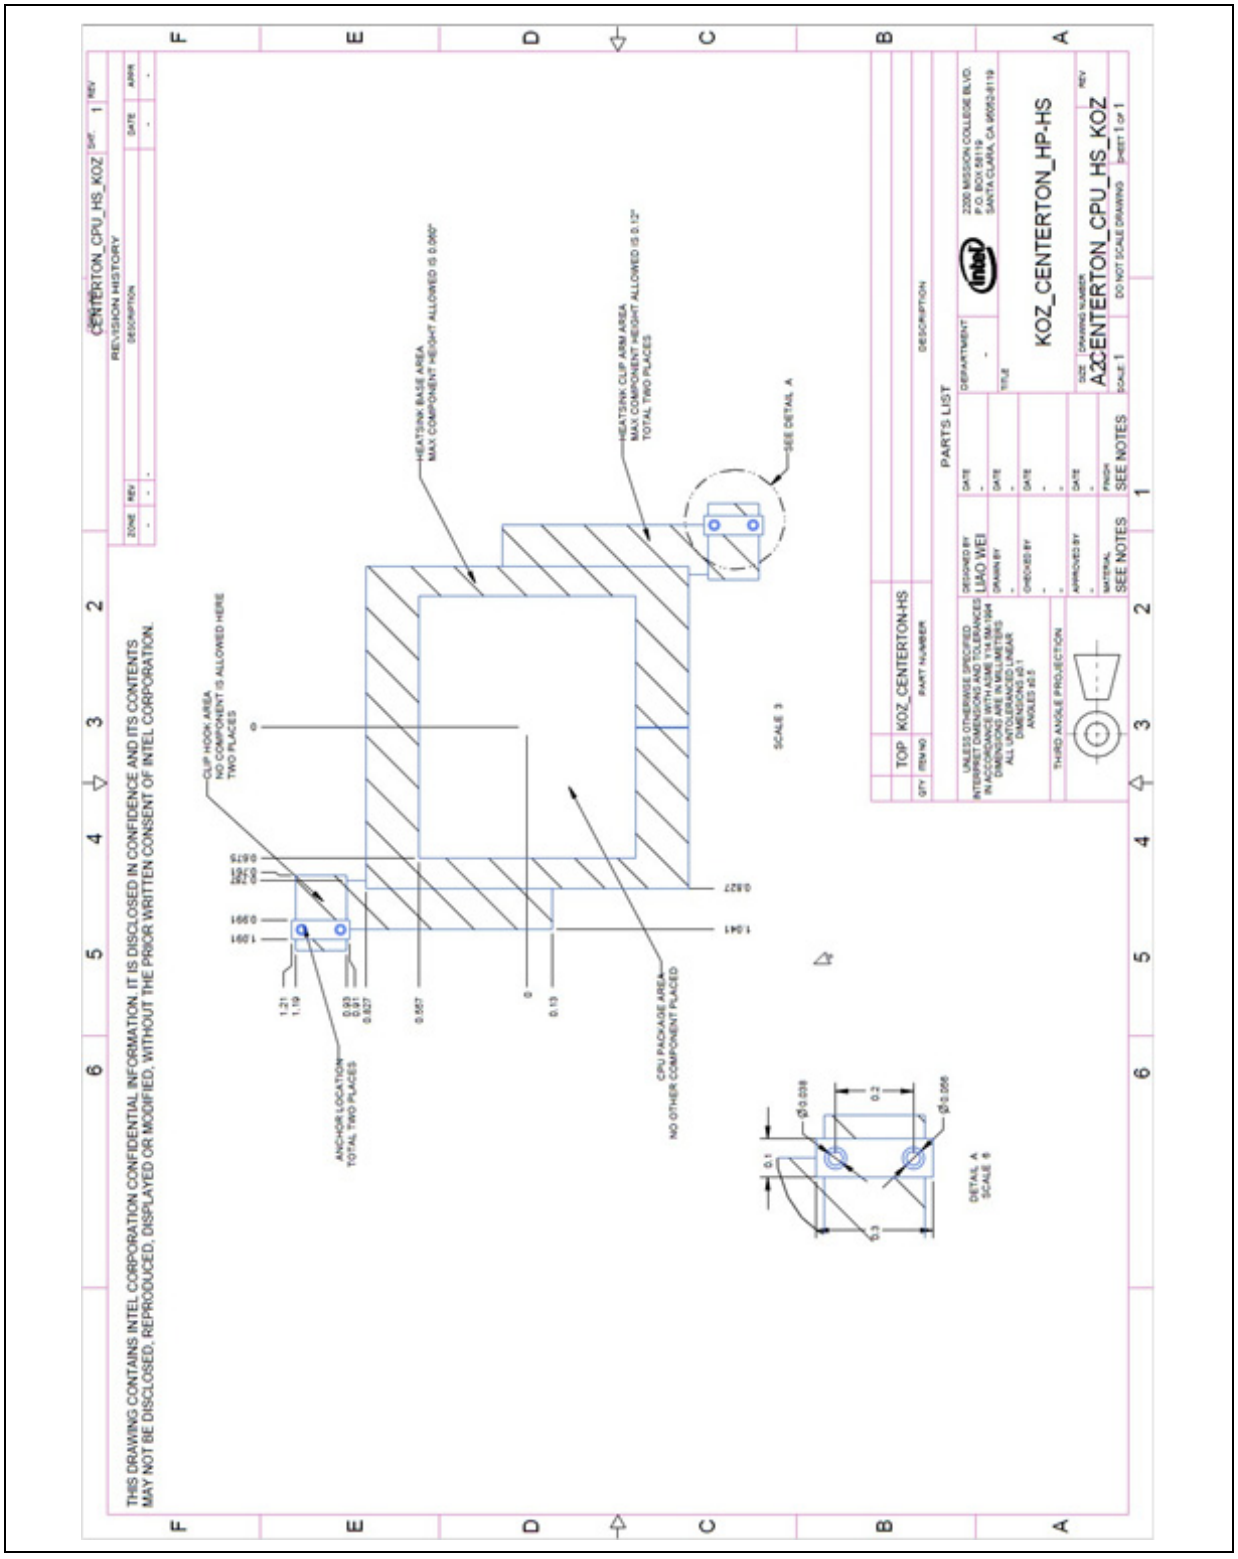

Figure A-8. Heatsink Assembly Volumetric Envelope

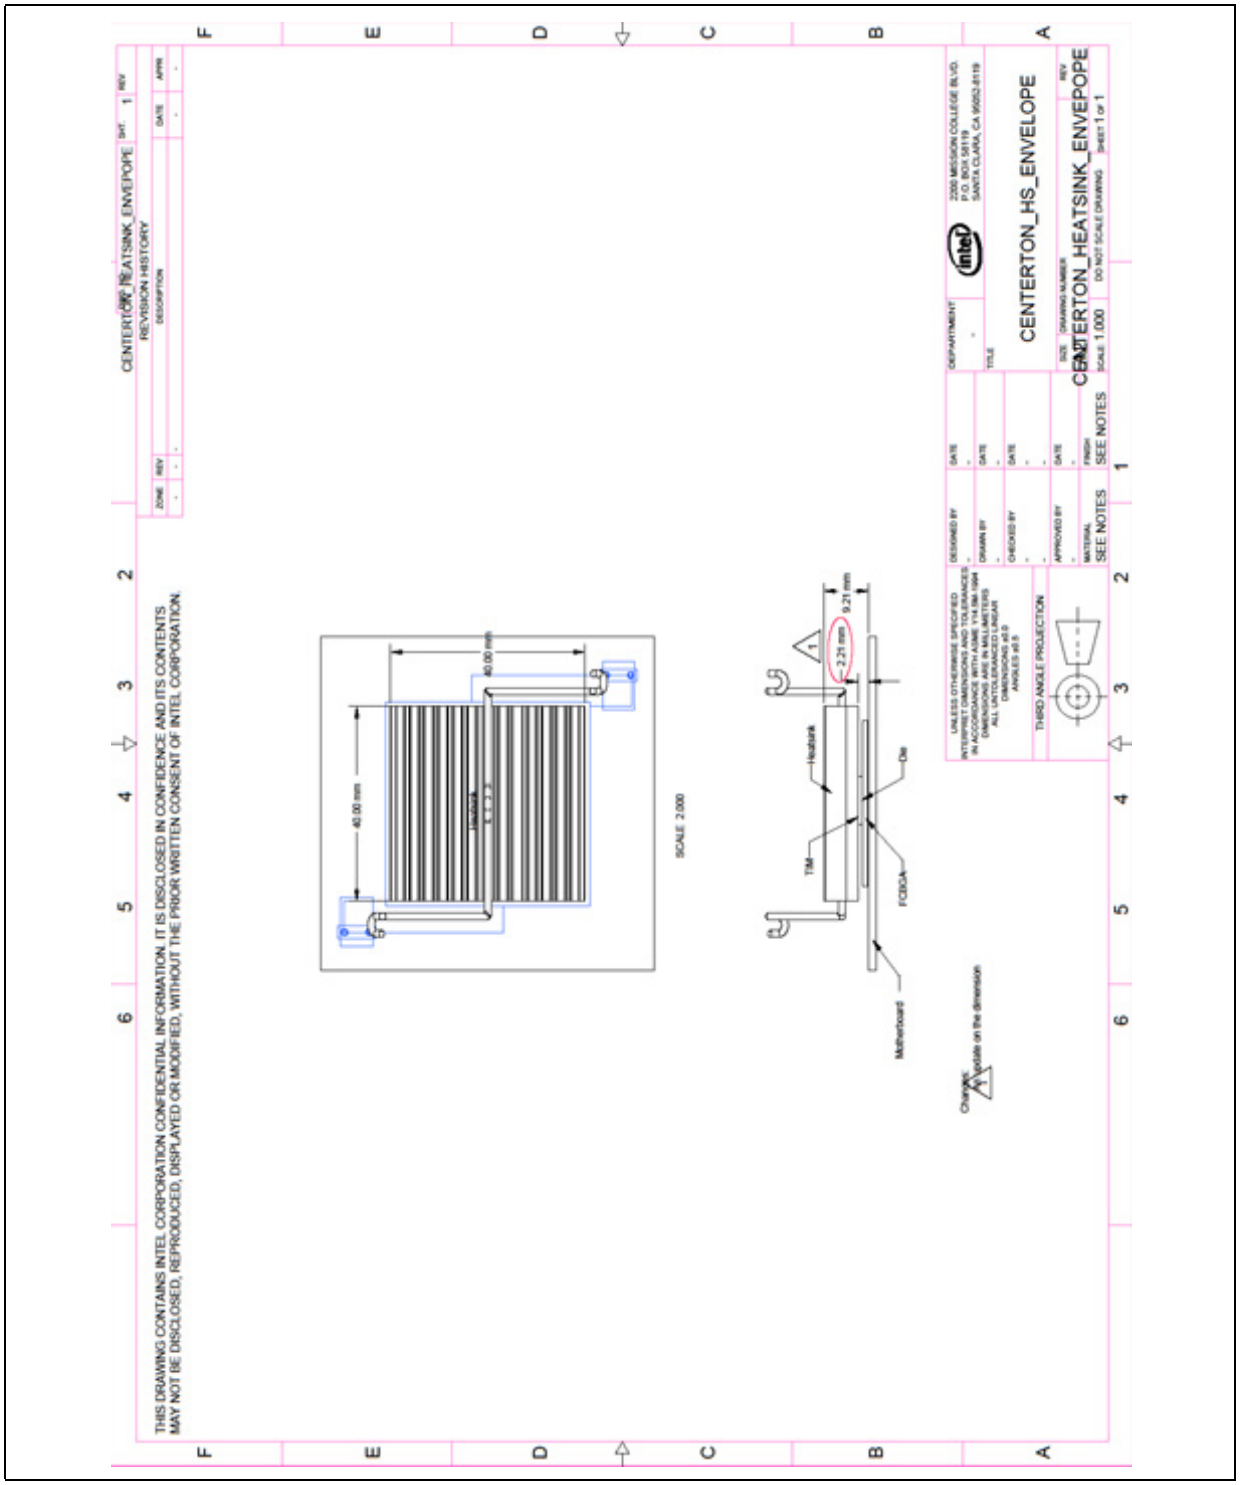

§ §

B Thermal Solution Supplier Information

This appendix includes supplier information for Intel-enabled vendors. Table B-1 lists the suppliers who can provide the processor reference thermal solution components. The part numbers listed below identify these reference components. End-users are responsible for the performance 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 B-1. Reference Heatsink-Enabled Components

Item	Intel PN	Wistron*	Foxconn*
Low-Profile Torsional Clip Heatsink Assembly	N/A	60.5P003.001	N/A
Solder-Down Anchor	A13494-008	N/A	• HB9703E-DW (0.062” thick motherboard) • HB9703E-M3W (0.085” thick motherboard)

Notes:
1. The enabled components may not be currently available from all suppliers. Contact the supplier directly to verify time of component availability.
2. Anchor is independent of heatsink assembly. Proper Anchor selection protects the chipset heatsink from shock and vibration.

§ §

C Heatsink Retention Load Metrology

C.1 Overview

The primary objective of the preload measurement is to ensure the preload designed into the retention mechanism is able to meet minimum specifications and does not violate the maximum specifications of the package.

C.2 Test Preparation

C.2.1 Heatsink Preparation

The following components are required to validate a generic fastener solution:

1. Thermal solution heat-sink (for example, PN: E51968-001 for Intel® Atom™ D500 series).
2. Fastener (for example, ITW PN: 83FT02-37-9909 for Intel® Atom™ D500 series).
3. Customized top plate to allow fastener attachment and package simulator.

C.2.2 Typical Test Equipment

For the heatsink clip-load measurement, the equivalent test equipment on the list work as a reference (Table C-1).

Table C-1. Typical Test Equipment

Item	Description	Part Number (Model)
Load cell Notes: 1, 5	• Honeywell* Sensotec Model 13 subminiature load cells, compression only. • Select a load range depending on load level being tested. • www.sensotec.com	AL322BL
Data Logger (or scanner) Notes: 2, 3, 4	Vishay* Measurements Group Model 6100 scanner with a 6010A strain card (one card required per channel).	Model 6100
Clip Force Measurement machine Note: 6	Customized machine that houses load cell for force measurement. Top-side plate can be modified to accommodate various attach pattern.	CFM-001 (Cool Star Technology)

Notes:

1. Select load range depending on expected load level. Whenever possible operate in the high end of the load cell capability. Check with the local load-cell vendor for further information.
2. Since the load cells are calibrated in terms of mV/V, a data logger or scanner is required to supply 5V DC excitation and read the mV response. An automated model takes the sensitivity calibration of the load cells and converts the mV output into pounds.
3. With the test equipment listed above, automating the data recording and control is possible with a 6101-PCI card (GPIB) added to the scanner, allowing it to be connected to a PC running LabVIEW* or Vishay StrainSmart* software.
4. **IMPORTANT**: In addition to just a zeroing of the force reading at no applied load, calibrate the load cells against known loads. Load cells tend to drift. Contact the local load-cell vendor for calibration tools and procedure information.
5. When measuring loads under thermal stress (bake for example), load-cell thermal capability must be checked, and the test setup must integrate any hardware used along with the load cell. For example, the Model 13 load cells are temperature compensated up to 71 °C as long as the compensation package (spliced into the load cell wiring) is also placed in the temperature chamber. The load cells can handle up to 121 °C (operating), but their uncertainty increases according to 0.02% rdg/°F.
6. A clip-force measurement machine is recommended to be calibrated before usage. Standard weights should be used to check for preload-cell accuracy and consistency.

C.3 Test Procedure Examples

The following procedure is for a generic fastener solution using the clip-force measurement machine at room temperature:

1. Prepare and then fasten the top plate onto the clip-force measurement machine. Place package simulator on top of the preload cell as well.
2. Place the heatsink (remove any TIM material) on top of the package simulator. Power on the clip-force measurement machine.
3. Install the fasteners and record down the measured preload. Make sure measurement is taken after the reading stabilized. Remove all fasteners and repeat two times (in total three times) to ensure consistency.
4. Repeat Step 4 for the remaining fastener samples. Recommended minimum samples are 10 sets of fastener samples.

§ §

D Thermal Metrology

This section discusses guidelines for testing thermal solutions including measuring processor temperatures.

D.1 Digital Thermal Sensor measurement (T_{J-MAX} Methodology)

The system integrator or designer should monitor the processor junction temperature (T_J) to validate their thermal solution design. To ensure functionality and reliability, the thermal solution should be able to maintain T_J in the processor at or below the maximum temperature specification (T_{J-MAX}).

The Digital Thermal Sensor (DTS) can be used to measure Tj in an operational SoC. The definition of the DTS value is given in Section 3.2.11, "Digital Thermal Sensor." Intel offers a software tool to read the DTS value in real time when conducting thermal tests. Contact the local Intel representatives for more details regarding this tool. System designers can always access the DTS value via an MSR. Check the *Intel® Atom™ Processor S1200 Product Family for Microserver and Storage BIOS Writer's Guide* for more details.

D.2 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.

Determine the local ambient temperature in the chassis around the processor to understand the effect it may have on the die 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**, avoid taking 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 D-1 on page 58 (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 acrylic, 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, graphic card, and 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, 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, the T_A measurements likely 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 D-2 on page 59. 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, 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, disable the fan regulation and power the fan directly, based on guidance from the fan supplier.

Figure D-1. Locations for Measuring Local Ambient Temperature, Active Heatsink

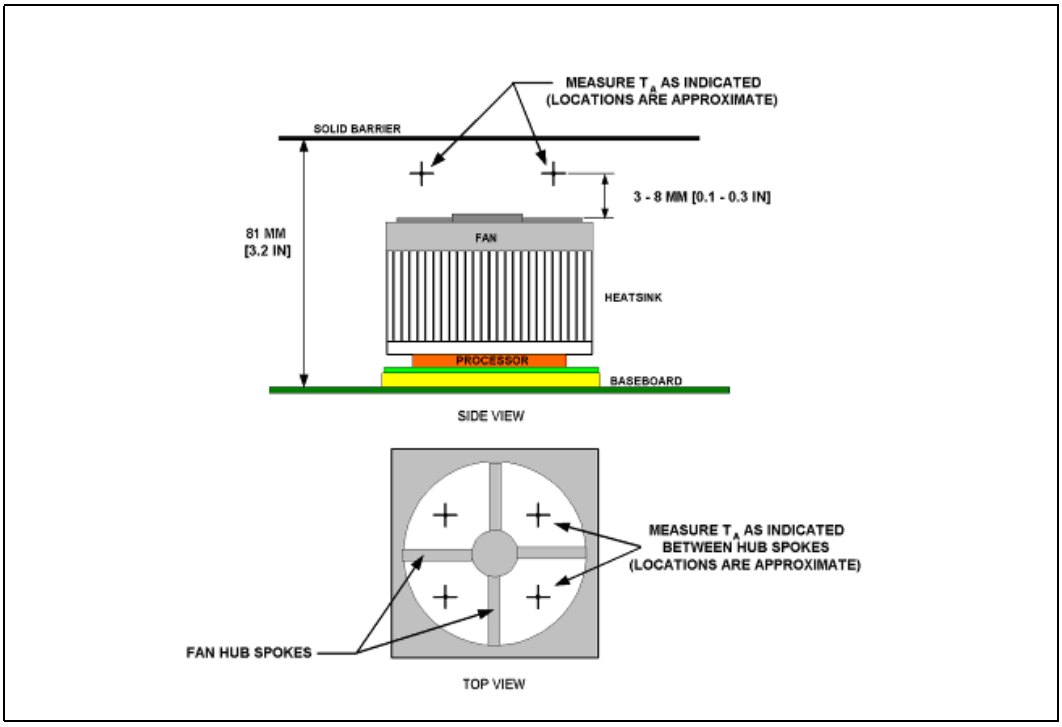

Note: Drawing Not to Scale

Figure D-2. Locations for Measuring Local Ambient Temperature, Passive Heatsink

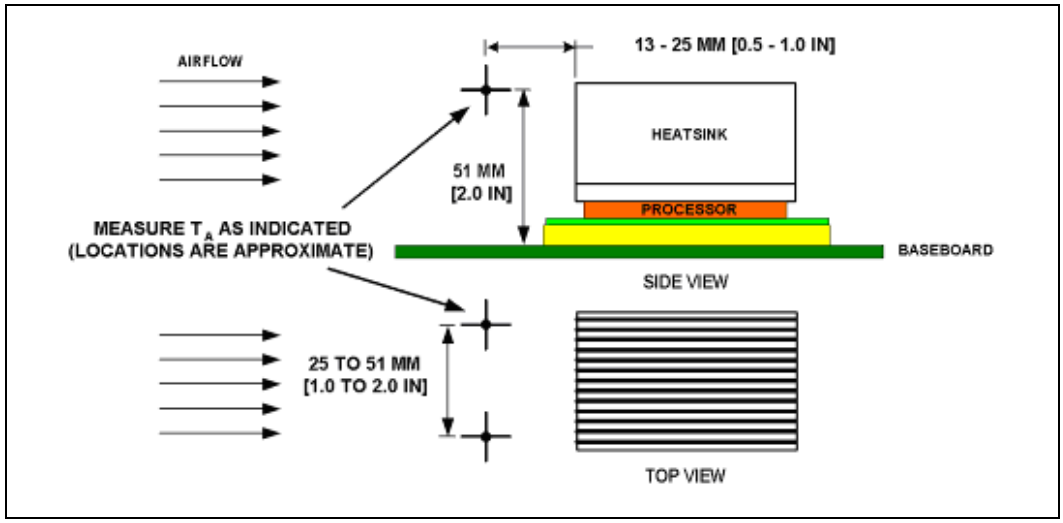

Note: Drawing not to scale

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

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

§ §