NLX Chassis Design Suggestions

Version 2.0

IMPORTANT INFORMATION AND DISCLAIMERS

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

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

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

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

Release 2.0, November 1998

Contents

List of Figures

List of Tables

1. Introduction

This paper provides the information necessary to design a low-profile NLX chassis. (Net PC and tower NLX chassis are also possible; these chassis types are not directly addressed in this paper.) In this paper, "NLX chassis" means a chassis that supports a motherboard which complies with the NLX motherboard specification, also called "the NLX specification." Therefore, most of the information in this paper relates to the mechanical design of the motherboard and other system components that interface to it, such as the motherboard mounting mechanisms, riser card, rear chassis I/O shield, and related mechanical elements.

As fully described in the NLX specification, NLX is a new motherboard form factor designed to improve upon current form factors and to adapt to new market trends and PC technologies. NLX does the following:

- Supports current and future processor technologies
- Supports new Accelerated Graphics Port (A.G.P.) high performance graphics solutions
- Supports tall memory modules and DIMM and SIMM technology
- Provides more system level design and integration flexibility; for example, the new design flexibility allows system designers to implement a motherboard that can be installed quickly, in most cases without using screws, thus lowering the PC’s total cost of ownership

NLX compliance requires that an NLX chassis design can accommodate any NLX-compliant motherboard. Also, an NLX motherboard design includes a standard I/O RFI shield for any size motherboard using that design.

The NLX specification and other information on NLX are available through a public web site whose URL is http://www.teleport.com/~nlx.

1.1 NLX Form Factor Overview

Figure 1.1 shows an example of an NLX motherboard, riser card and I/O shield.

- The add-in card riser is located at the left edge of the motherboard (as viewed from the rear).
- Tall components such as the processor and memory are typically located on the opposite side from the I/O slots. This allows full length add-in cards in many system configurations.
- The rear I/O connectors are stacked single and double high to support more connectors.

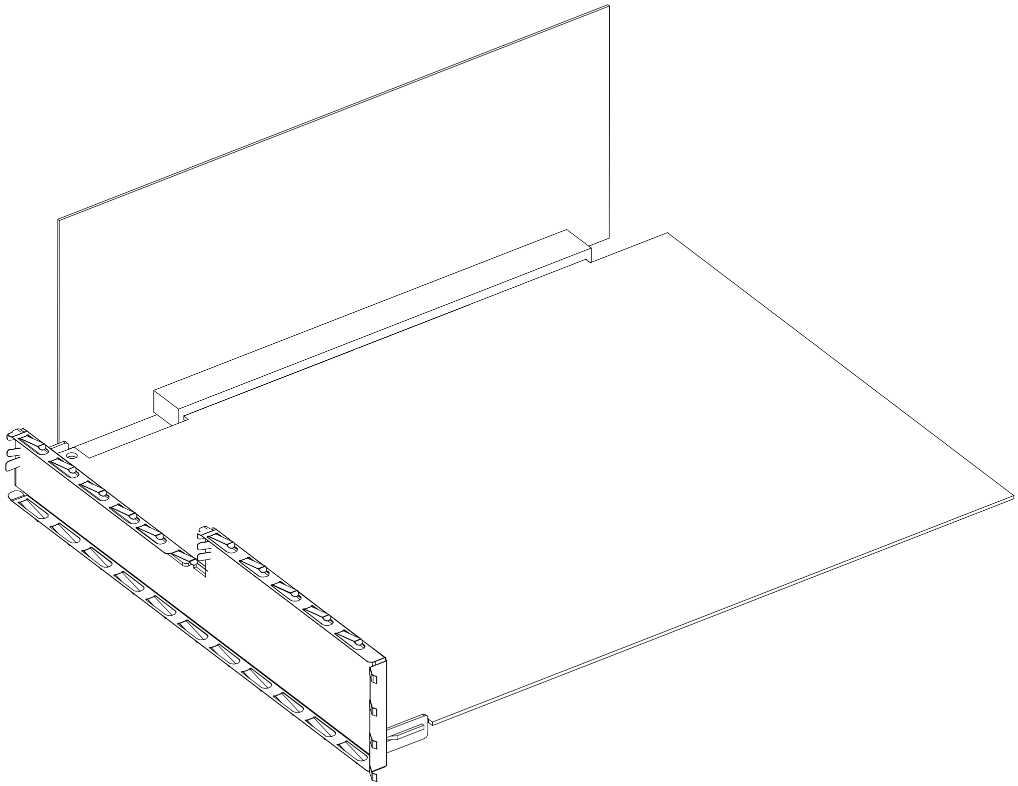

Figure 1.1: NLX Board and Riser Example (as viewed from the rear)

1.2 NLX System Configuration Example

The NLX specification defines a motherboard and riser with a basic and optional extended connector pin-out. Figure 1.2 shows an NLX motherboard in a system layout that highlights NLX form factor features (the specification offers many possible system configurations).

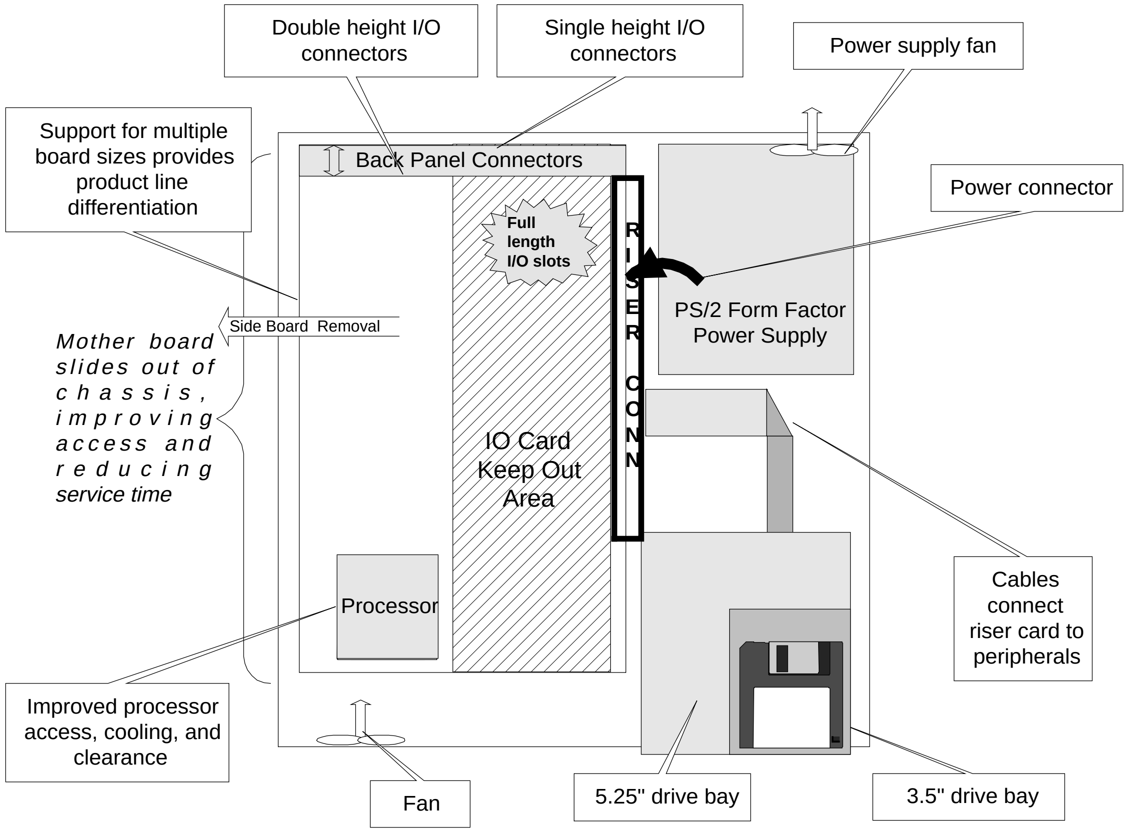

Figure 1.2: NLX System Layout Example

Summary of features shown in Figure 1.2:

- Use of double high connectors across half the rear I/O and single high across the whole 9” maximum board width
- Power supply with internal cooling fan
- Power supply harness that connects directly to the riser card, not to the NLX motherboards
- Floppy and IDE-based peripheral signal cables that attach to the riser card
- Drive bays that can all fit to one side of the riser in a convenient 17.5" wide system form factor
- Processor located toward the front of the system, close to the fan, for optimum cooling
- Dockable motherboard that slides into the system and docks with the riser card

For more information on NLX-compliant system configurations, basic and extended NLX riser cards, and related topics, see the NLX specification.

1.3 About This Document

1.3.1 Revision History

Version	Revision History	Date
1.0	First release of NLX Chassis Design Suggestions	May 1997
1.x	Combined Version 1.0 Sections 2 and 3. Renumbered Version 1.0 Sections 4 and 5 as Version 2.0 Sections 3 and 4. Added new Sections 1.4, 5, 6, and 7.	November 1998

1.3.2 Contents of Version 2.0

These design suggestions were created to assist NLX chassis designers and fabricators. In addition to this Introduction, Version 2.0 has six major parts:

- Section 2 provides information about the dimensions and mechanical features of NLX motherboards that are relevant to chassis design and the integration of the motherboard into the chassis. (Motherboard designers should consult the NLX motherboard specification.) The integration sections describe the recommended motherboard mounting mechanism, the motherboard/riser card interface, and the rear chassis I/O shield.
- Section 3 discusses power supply mounting.
- Section 4 outlines other mechanical design considerations.
- Section 5 introduces general thermal design principles and discusses fan types and placement, power supply considerations, and chassis airflow patterns.
- Section 6 discusses the thermal test methodology to evaluate the system thermal design.
- Section 7 demonstrates several thermal design examples for NLX form factors.

For the full specification of the NLX motherboard and other special topics, chassis designers should consult the documents available at http://www.teleport.com/~nlx.

1.4 Related Documents

1.4.1 Chassis and Motherboard

- *PC 98 & PC 99 System Design* Guides from Intel Corporation and Microsoft Corporation
- *NLX Motherboard Specification*
- *PCI Local Bus Specification*
- *PCI Bus Power Management Interconnect (PCI-PM) Specification*
- *Audio Codec '97 Component Specification*
- *Data sheet – Intel Pentium® II Processors At 233 MHz, 266 MHz, 300 MHz, and 333 MHz (Order Number 243335)*

- *Data sheet – Intel® Celeron™ Processors At 266 MHz, 300 MHz, 300 MHz, and 333 MHz (Order Number 243658)*
- *Application Note (AP-586) Pentium® II Processor Thermal Design Guidelines (Order Number 243331)*
- *Direct Rambus™ - Technology Disclosure (WWW.rambus.com)*
- *Advanced Configuration and Power Interface Specification*
- *Accelerated Graphics Port (A.G.P.) Design Guide*

1.4.2 EMC

- **U.S.:** *FCC Code of Federal Regulations (CFR) 47 Part 2 & 15, Class B*
- **Canada:** *DOC CRC c, 1374 Class B*
- **Europe:** *EN55022, Class B & EN50082-1 (Complying with the European Union's EMC Directive, 89/336/EEC)*
- **International:** *CISPR 22, Class B*

1.4.3 Acoustics Sound Pressure

Test setup and method are based on ISO 7779 "noise emitted by computer and business machines."

1.4.4 Safety

- **U.S.:** *Underwriter Laboratories Inc. (UL-1950)*
- **Canada:** *Canadian Standards Association (CSA C22.2-950)*
- **Europe:** *EN60950 (Complying with the European Union's EMC Directive, 89/336/EEC)*
- **International:** *International Electrotechnical Commission (IEC 950)*

2. Motherboard Integration and Mechanical Requirements

This section describes the mechanical dimensions of the NLX form-factor motherboard, including physical size, mounting hole placement, back panel I/O shield opening, card-edge connector placement, and component height constraints.

The most unique feature of an NLX motherboard is the "gold finger" card edge (shown in drawings in this section). This card edge mates with the riser card edge connector. The riser card provides interfaces for PCI, ISA, IDE, floppy drives, USB, power, audio, and front panel features.

2.1 Motherboard Dimensions

2.1.1 Definition of Terms Used in Mechanical Diagrams

Term/phrase	Definition
Tolerances	All dimensions are mm/inches unless otherwise specified. +/- tolerances are: Sheet metal / chassis = +/- 0.010" Board mounting hole to edge = +/- 0.010" Board hole to hole = +/- 0.005" Board hole size = + 0.003" -0.001" Plastic = +/- 0.005"
Motherboard thickness	A nominal motherboard thickness of 0.062" is used; deviation from this nominal value will affect most of the dimensions and must be considered.
Keep-out areas for EMI clips, required	The EMI clip location area is reserved for EMI clips. The chassis must provide keep-out areas in the EMI clip location areas to provide an electrical contact to the chassis. The EMI clips protrude from the secondary side of the motherboard in these areas. Required.
Keep-out areas for mounting holes, required	Primary side areas of the motherboard with no traces and components, to assure the motherboard can accommodate the attachment of rail support. ***Required.***
Keep-out zones, recommended	Secondary side zones of the motherboard with no pin through-holes and secondary side SMT components, to assure that the motherboard can slide into the chassis and be mounted without rails. ***Recommended.***

2.1.2 Motherboard Sizes

The NLX specification supports motherboards with overall dimensions of 9.0" x 13.6" (maximum) to 8.0" x 10.0" (minimum). An NLX-compatible chassis must be able to accommodate motherboards with these two extreme dimensions, and all in between. The specification defines the maximum motherboard width as 9.0", so an NLX chassis must be designed to accept boards up to 9.0" wide. The minimum motherboard width is 8.0"; an NLX chassis can support board widths down to 8.0" because of the way the NLX mounting holes are placed. The specification defines motherboard lengths from 10.0" to 13.6".

Motherboard dimensions with keep-out areas are illustrated in Figures 2.1 to 2.6. The figures show the primary and secondary sizes of 8.0"-9.0" board widths, in lengths of 10.0", 11.2", and 13.6".

2.1.3 Mounting Hole Placement

To simplify the design of both motherboards and chassis, the specification details three sets of mounting holes, which are located to allow motherboard optimization. **Any given motherboard needs to have only one set of four mounting holes, whereas the chassis is required to support all three sets.** The chassis designer must determine how the motherboard is mounted to the chassis with the specified set of mounting holes.

The exact method of how the chassis uses these holes is left to the system designer. The chassis designer must adhere to all requirements and keep-out area details of the NLX specification for the chassis to be NLX-compatible.

The motherboard can be mounted with any set of the mounting holes provided but must support the four holes in the defined set. The chassis needs to support only one set of four holes at a time. Three sets of holes are defined to support the most common motherboard sizes, as listed in Table 2.1 with hole sets and locations.

2.1.4 Motherboard Keep-out Specifications

- Primary side: **Required,** four keep-out areas; these are 0.390" diameter centered around the four mounting holes (plated and grounded, Mounting Hole Required “keep-out area A”).
- Secondary side: **Required,** six keep-out zones; these are 0.400" square centered around the four mounting holes (plated, required “keep-out zone A”), and the two additional keep-out areas (required “keep-out zone B”) for rail bumpers. No traces, through-holes devices, or SMT devices are allowed within “keep-out zone A”. No through-holes devices or SMT devices are allowed within required “keep-out zone B.”

Table 2.1: Mounting Hole Placement

Board length and hole set	Position of hole along length dimension (inches)	Position of hole along width dimension (inches)
13.6" board length Hole set A	13.400	7.800
	13.400	4.400
	7.175	7.800
	7.175	4.400
11.2" or greater board length Hole set B	11.000	7.800
	11.000	4.400
	4.775	7.800
	4.775	4.400
10.0" or greater board length Hole set C	9.800	7.800
	9.800	4.400
	3.575	7.800
	3.575	4.400

The NLX specification encourages motherboard manufacturers not to use through-hole and secondary side SMT devices along the suggested “keep-out zones B” (see Figures 2.1 to 2.6 for notes on zones). If these zones are kept free, the user has the option of sliding the motherboard in directly and mounting it to the chassis without rails. For this situation, board traces are allowed along the recommended keep-out zones, but chassis manufacturers are required to ensure that the motherboard traces are not damaged during the motherboard slide in/out process.

The secondary side “keep-out zone B” can have signal traces. It is the responsibility of the chassis manufacturer to ensure that the secondary side traces are not damaged by the chassis. (For example, for screw-mounting the motherboard, the chassis should use nondestructive standoffs.)

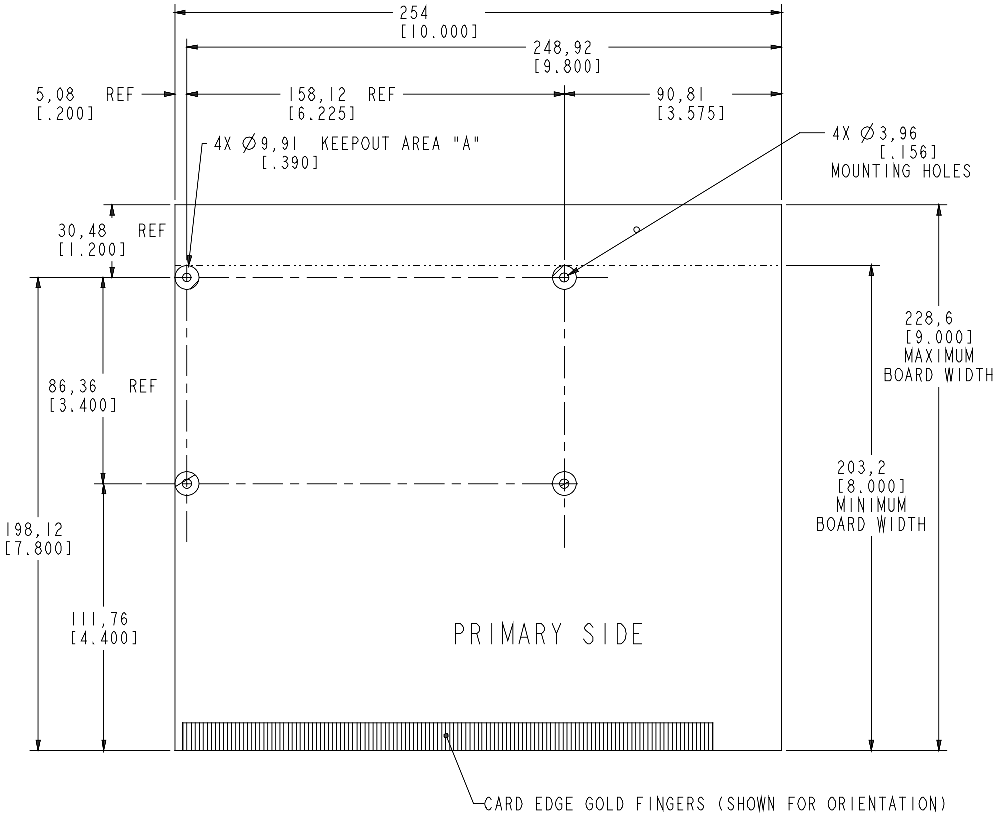

Figure 2.1: NLX Motherboard Dimensions—10.0" Primary Side

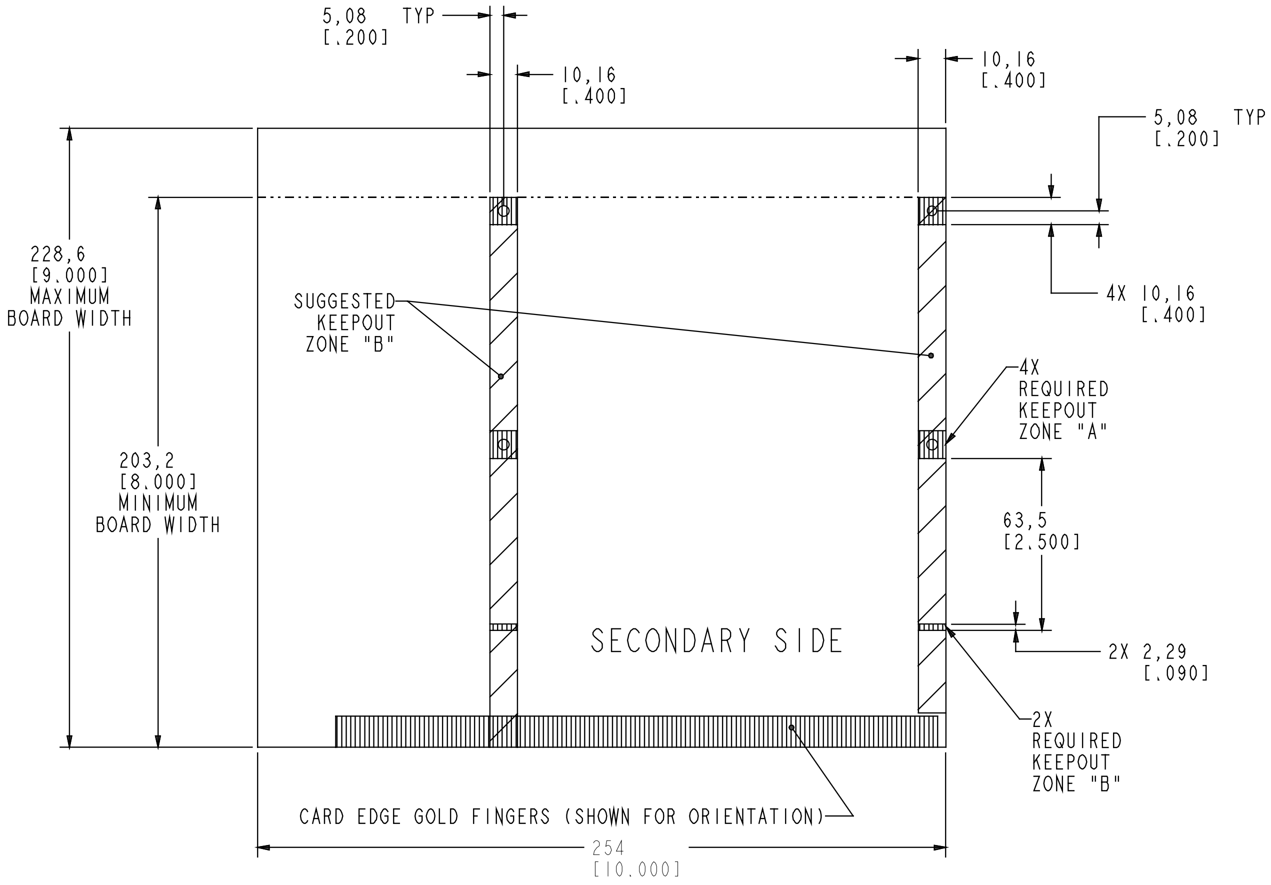

Figure 2.2: NLX Motherboard Dimensions—10.0" Secondary Side

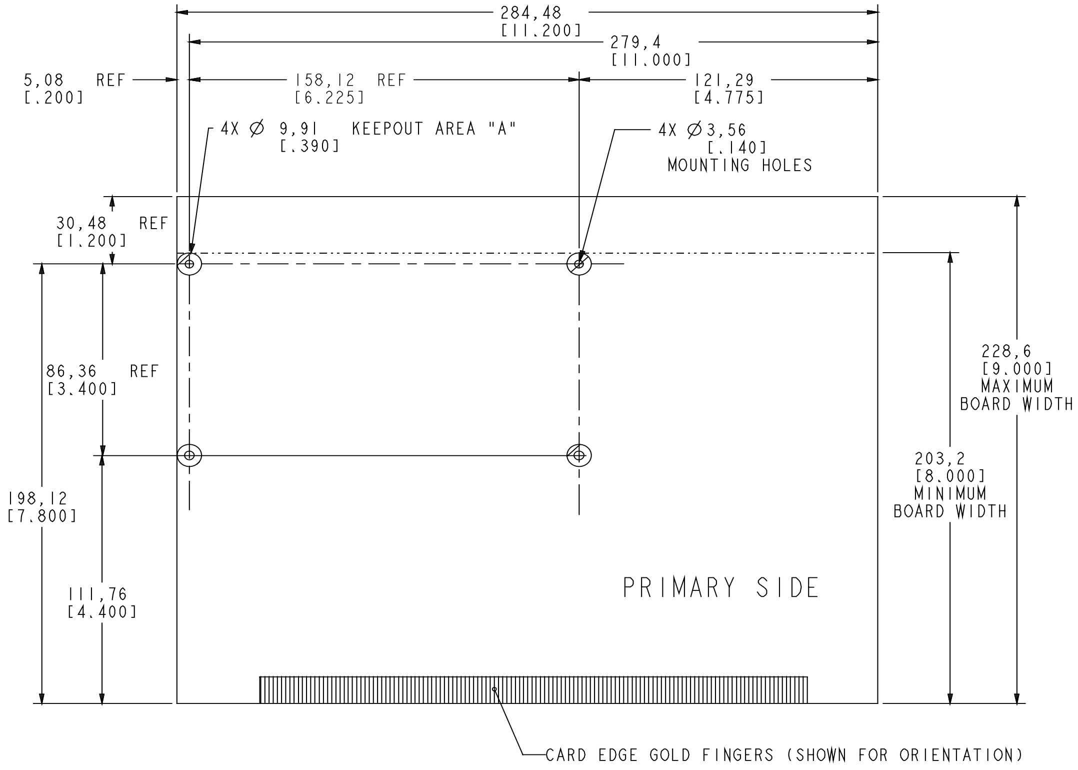

Figure 2.3: NLX Motherboard Dimensions—11.2" Primary Side

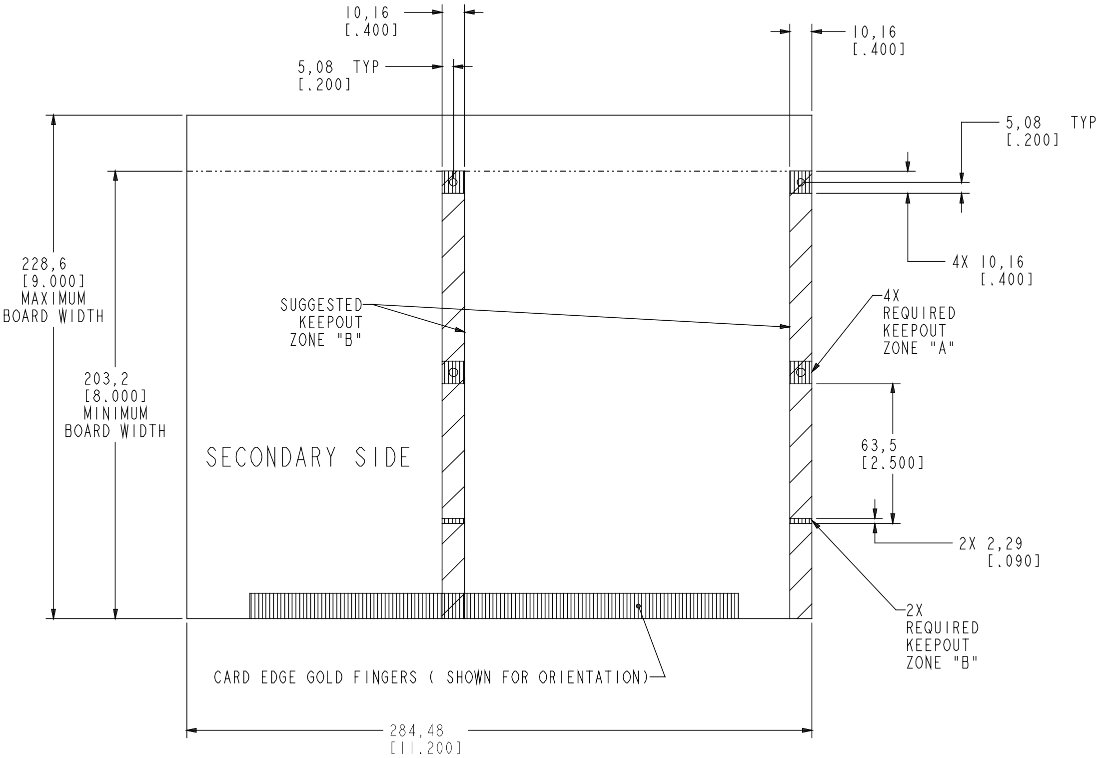

Figure 2.4: NLX Motherboard Dimensions—11.2" Secondary Side

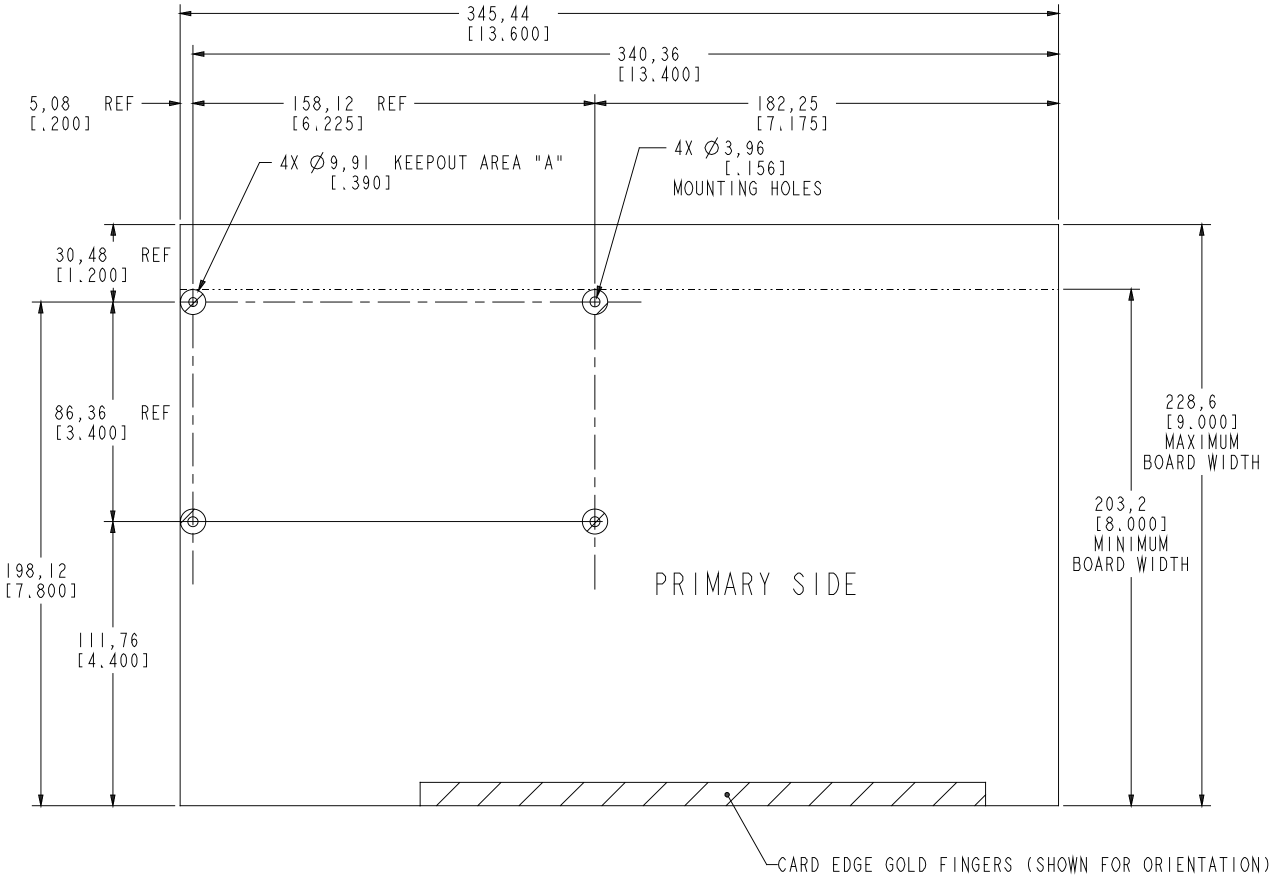

Figure 2.5: NLX Motherboard Dimensions—13.6" Primary Side

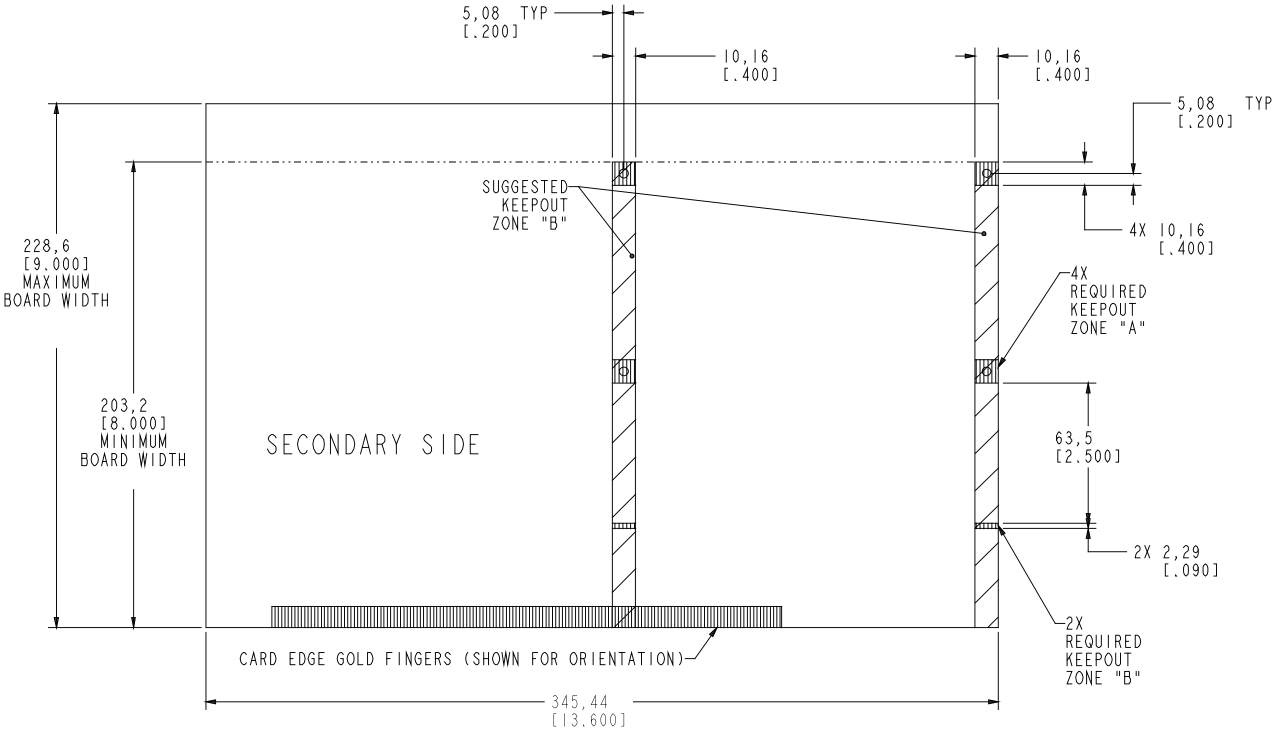

Figure 2.6: NLX Motherboard Dimensions—13.6" Secondary Side

2.2 Motherboard EMI Clip Locations with Chassis Keep-outs

Multiple grounding mechanisms are defined in the NLX specification to allow for EMI grounding between the motherboard and chassis. The chassis should allow for proper grounding in all indicated areas by providing flat surfaces that are free of paint. The following sections describe possible motherboard grounding schemes.

2.2.1 Grounding Through the I/O Shield

Grounding of the motherboard through the I/O shield is required in all NLX designs. The I/O shield attaches to the motherboard and should make contact with all of the connectors on the motherboard back panel to allow for proper grounding to the chassis along its perimeter.

This single point ground scheme may not be sufficient in all cases to suppress EMI radiation. Secondary grounding can be achieved through the use of clips snapped to the motherboard or through clips in rails. See Section 3.3 for details.

2.2.2 Grounding Clips Snapped on Motherboard

In some cases, EMI clips may be necessary to suppress EMI radiation hot spots on the motherboard. These clips are snapped on to the motherboard and protrude through the secondary (bottom) side, to make contact with a clean flat surface at the bottom of the chassis.

Figures 2.7 and 2.8 illustrate these motherboard grounding clips.

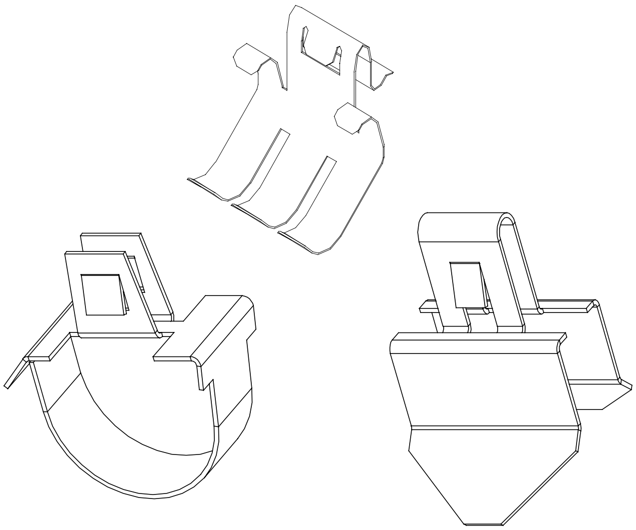

Figure 2.7: Examples of Motherboard Clips

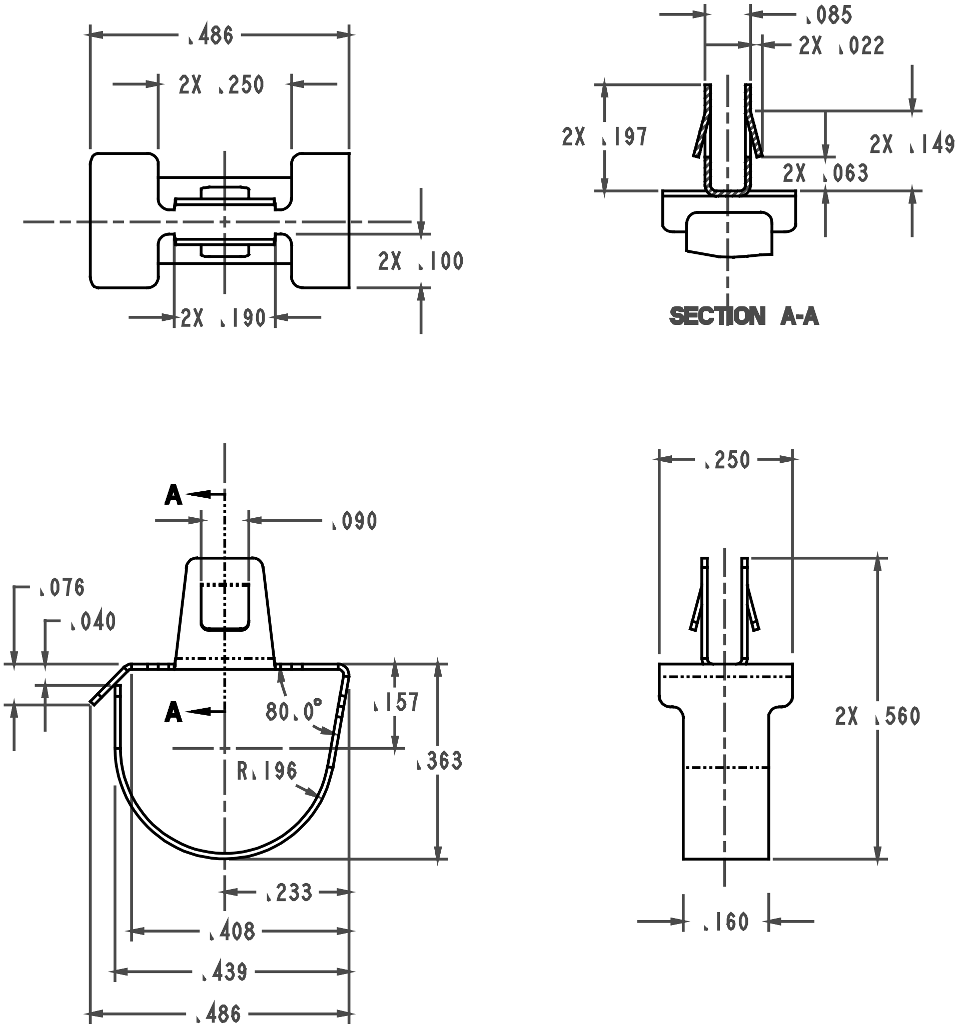

Figure 2.8: Sample Clip Designs

The NLX specification defines these clip areas. The chassis design must provide the contact surfaces. Figure 2.9 shows the chassis keep-out areas intended for this purpose.

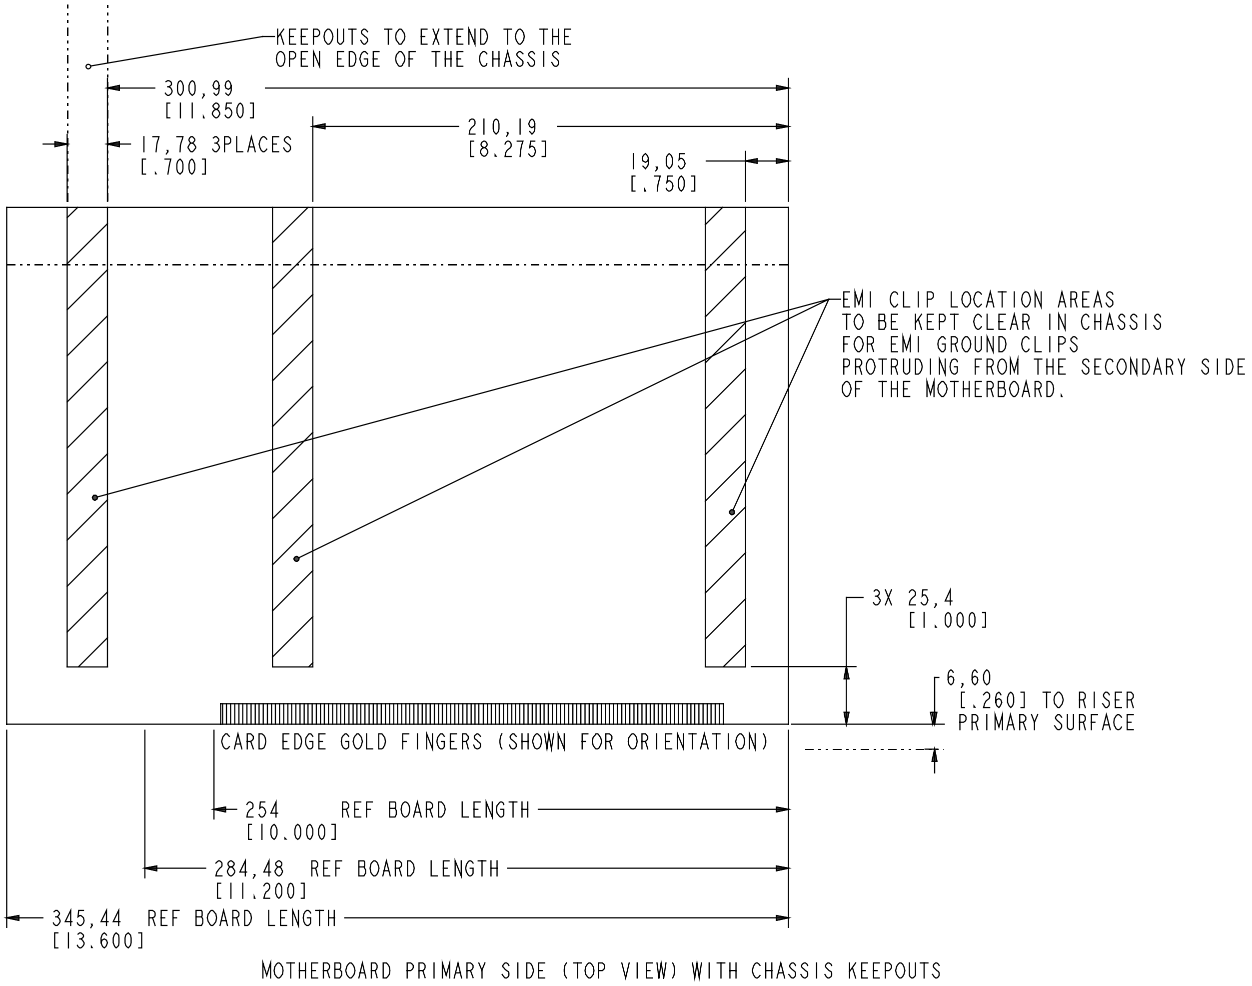

Figure 2.9: Chassis Keep-out Areas and Motherboard EMI Clip Location Areas

2.2.3 Grounding Clips in Rail Slides

If the chassis design uses rails for motherboard installation, clips in the rails make grounding contact with the chassis. The clips in standard NLX rails are located such that if grounding pads are provided, good EMI contact can be made. The grounding pads are at least .40-inch by .40-inch square and are located .34" behind the centers of the mounting holes. Figure 2.10 shows the grounding pad locations on the chassis. See section 3.1.5 for details.

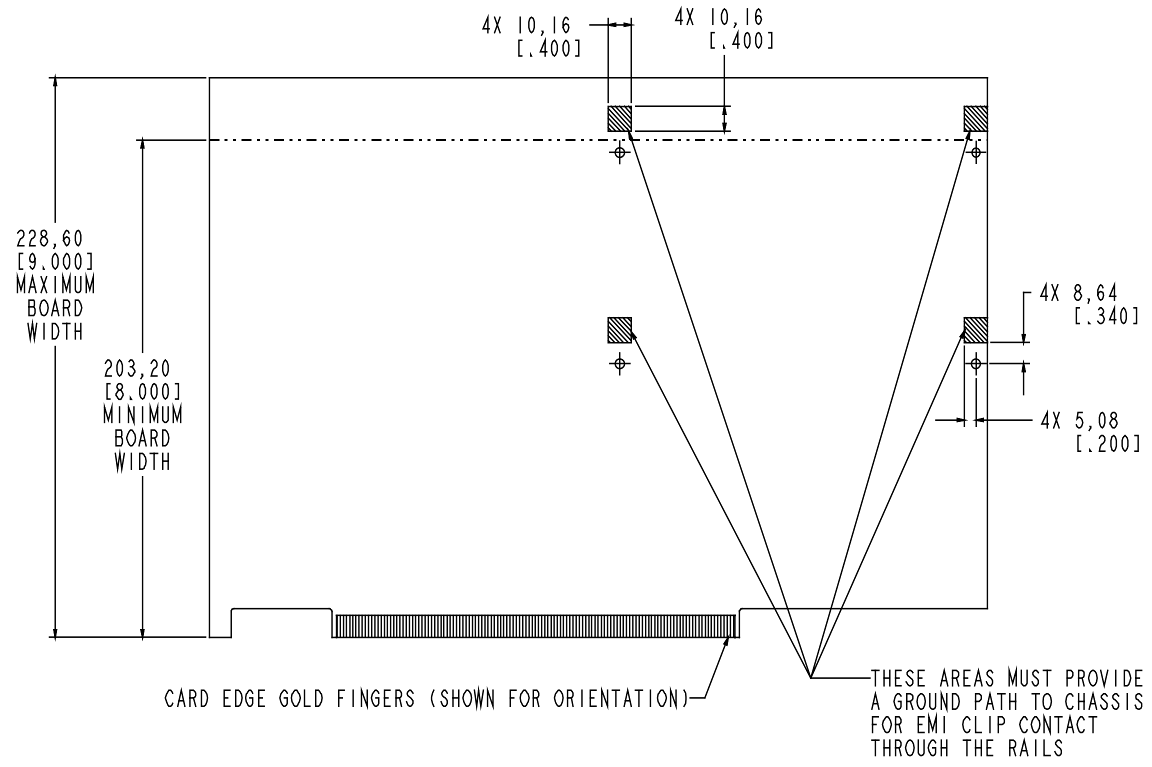

Figure 2.10: Ground Pad Locations

2.3 Motherboard Component Height Limits and Chassis Design

Figure 2.11 shows the maximum component height for the primary side (top) of the motherboard. Figure 2. shows the maximum component height for the secondary side (bottom) of the motherboard. The chassis designer should dimension the chassis with sufficient margin to accommodate vibration movements of taller components that is often seen during shipping and handling of computer systems.

2.3.1 Motherboard Top Component Height

The far left side of the motherboard accepts a component height of up to 3.60", supporting an NLX A.G.P. “notched” card. This maximum height area applies to 9.00" wide motherboards that have the option of supporting an NLX A.G.P. add-in card with a connector on the motherboard.

The front of the mid-left side of the motherboard accepts a component height of up to 2.80", supporting placement of extra-tall components such as very tall DIMM/SIMMs, and processor heat sinks or modules.

The middle of the mid-left side of the motherboard accepts a component height of up to 1.75”, supporting placement of tall components, such as DIMM/SIMMs.

The back of the mid-left side of the motherboard accepts double-height stacked I/O connectors of up to 1.430”.

The right side mid-to-back section of the motherboard has a 0.70" height restriction, which allows the use of half length add-in cards in all of the NLX expansion slots.

Chassis/system designers should assume a height restriction of 2.8" exists in the right side mid-to-front section of the motherboard (zone C), caused by tall components. Zone C is highlighted (crosshatched) in Figure 2.11

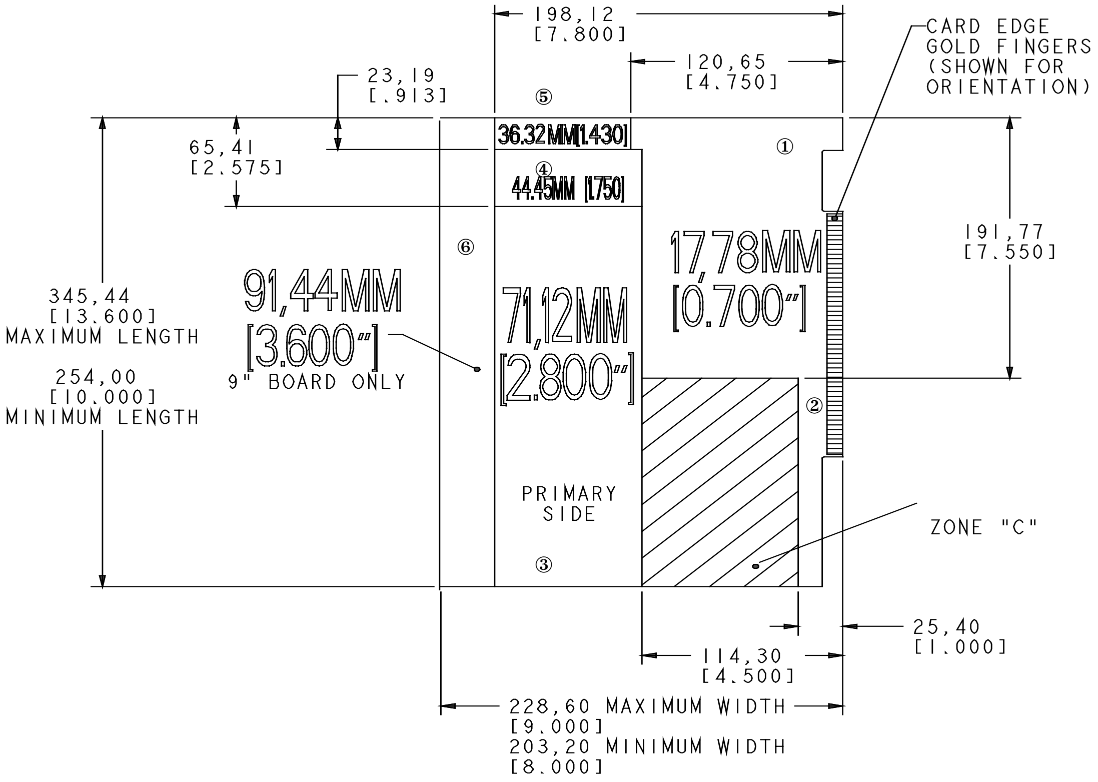

Figure 2.11: NLX Motherboard Primary Side Height Restrictions

2.3.2 Motherboard Bottom Component Height

Figure 2.12 shows the maximum component height on the motherboard bottom side. Restricted zones D are highlighted.

Overall, the motherboard has a bottom component height restriction of 0.150" except in the restricted areas D, where the component height is limited to 0.120". The restricted zones D allow sufficient board-to-chassis clearances to allow for chassis features such as rail guides to be built directly into the chassis.

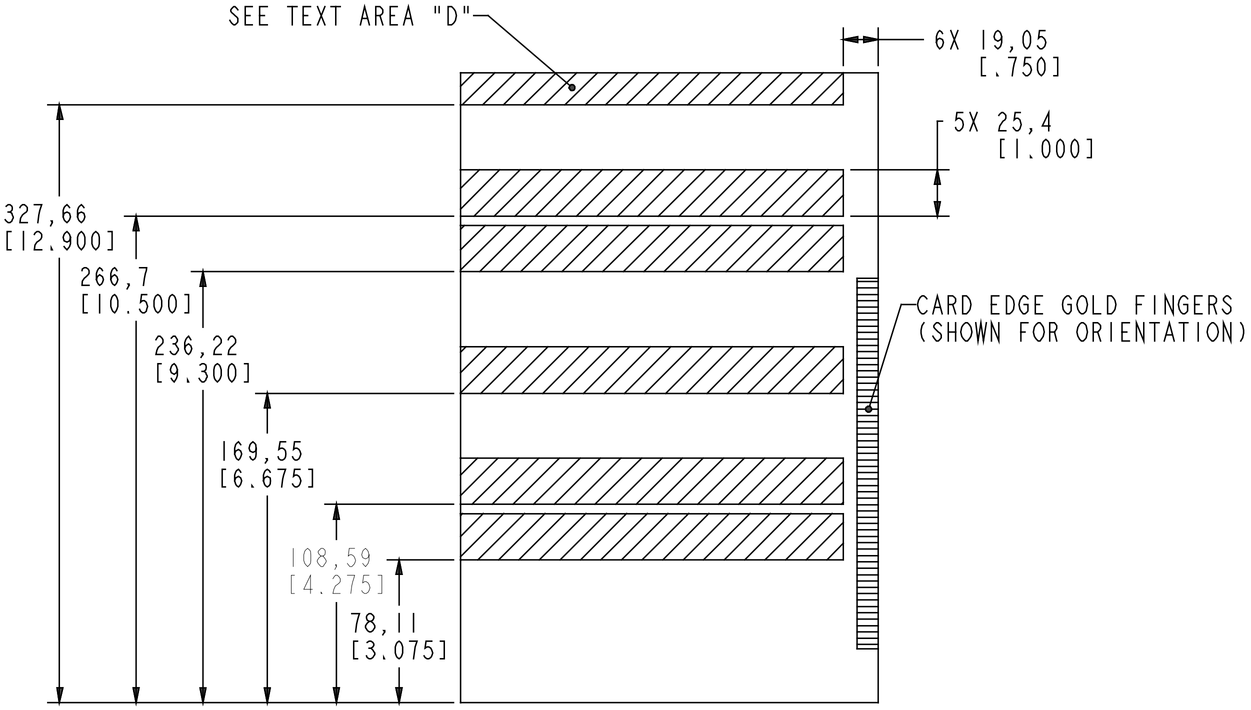

Figure 2.12: NLX Motherboard Bottom Side Component Height Restrictions

2.4 Motherboard Mounting Features

It is preferable in an NLX chassis to use a rail design to mount the motherboard. Other methods, such as standoffs, can also be used. An NLX chassis must provide three sets of mounting features, to accommodate all sizes of NLX motherboard. Figure 2.13 shows the required features.

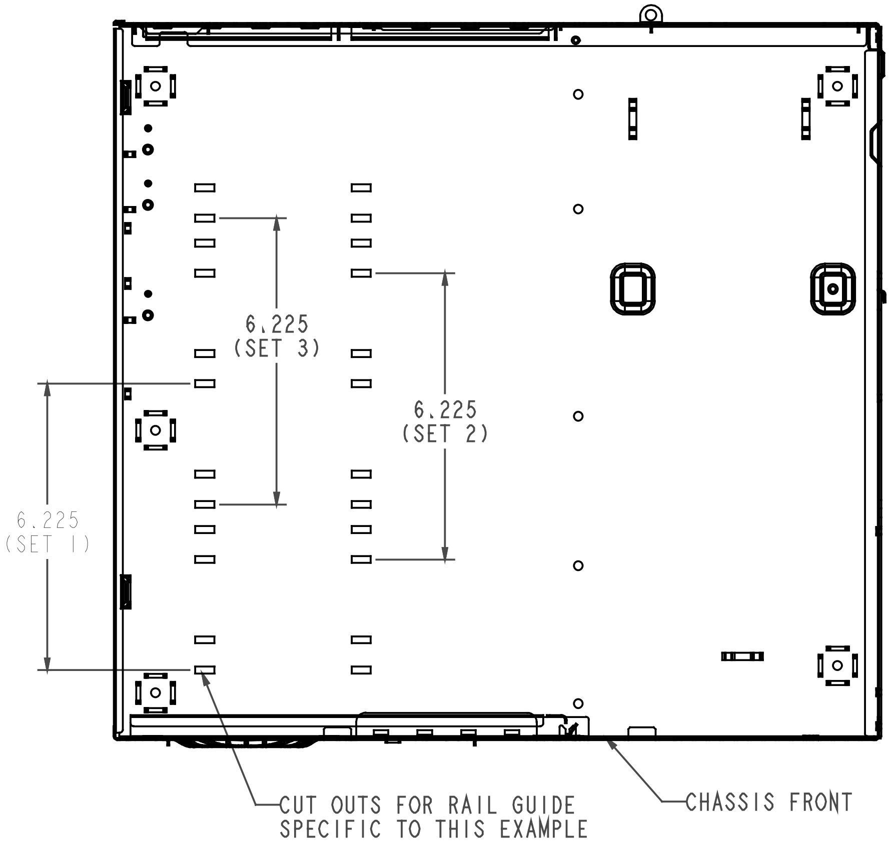

Figure 2.13: Mounting Features in Chassis

2.4.1 Motherboard Rails—Mounting the Motherboard to Chassis

The NLX motherboard rail system enables the mechanical support required for new processors and the slide-in installation and removal of the motherboard. To ensure that any rail-mounted NLX motherboard can be inserted into any chassis design for NLX, the motherboard rail system described in the NLX specification is highly recommended. Since each OEM is free to determine where in the assembly process the rails are installed, the rails are part of the chassis kit.

The rail design shown here is for plastic rails that contain a “ridge bumper” that contacts the motherboard but is not attached to it. The ridge bumper prevents the plastic rails from striking the motherboard during shipping. Metal rails, being stiffer, do not and must not have this ridge bumper, because it would damage secondary side motherboard traces.

Figures 2.14a and 2.14b show the rail features for attaching the rail to the motherboard and to the chassis. The rail flange geometry and position relative to the motherboard mounting holes is fully specified and fixed. This lets chassis manufacturers design chassis mounting schemes that will accept the rails on any given NLX board. A 0.140" clearance is defined between the top of the rail and the bottom of the motherboard for the length of the rails, except in the mounting hole locations, where the screw boss and support tabs exist.

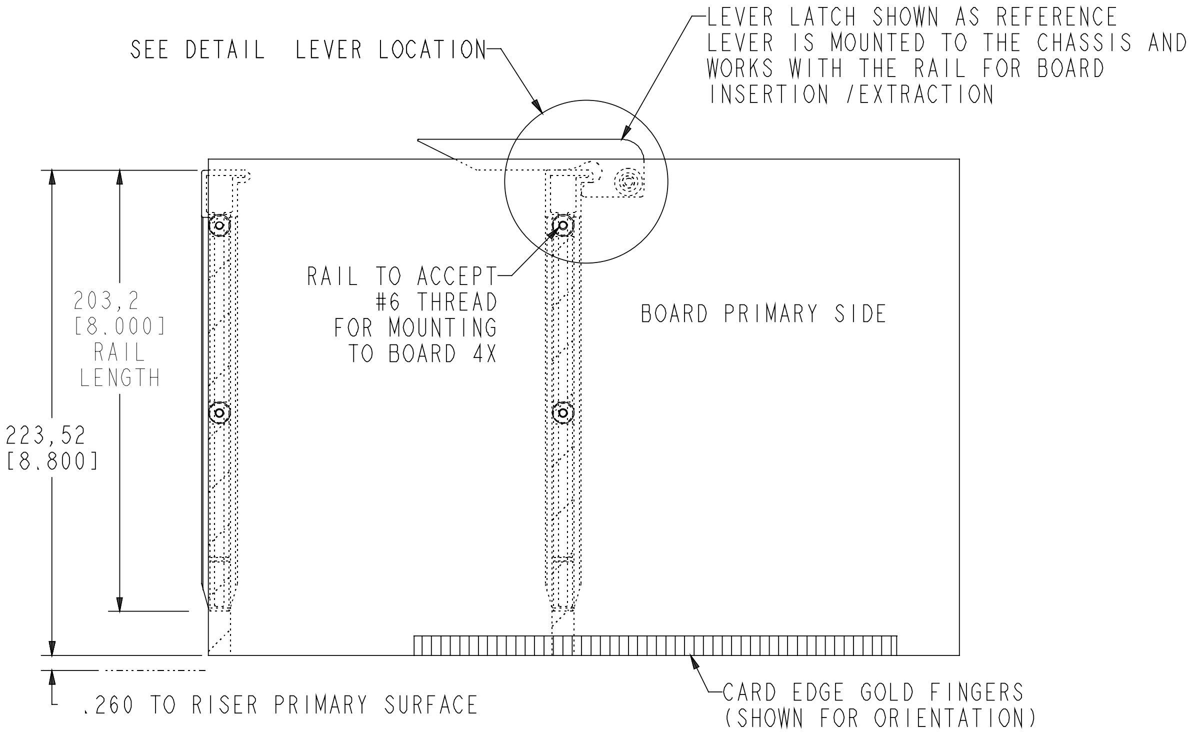

Figure 2.14a: Typical Rail Design Example (Part I)

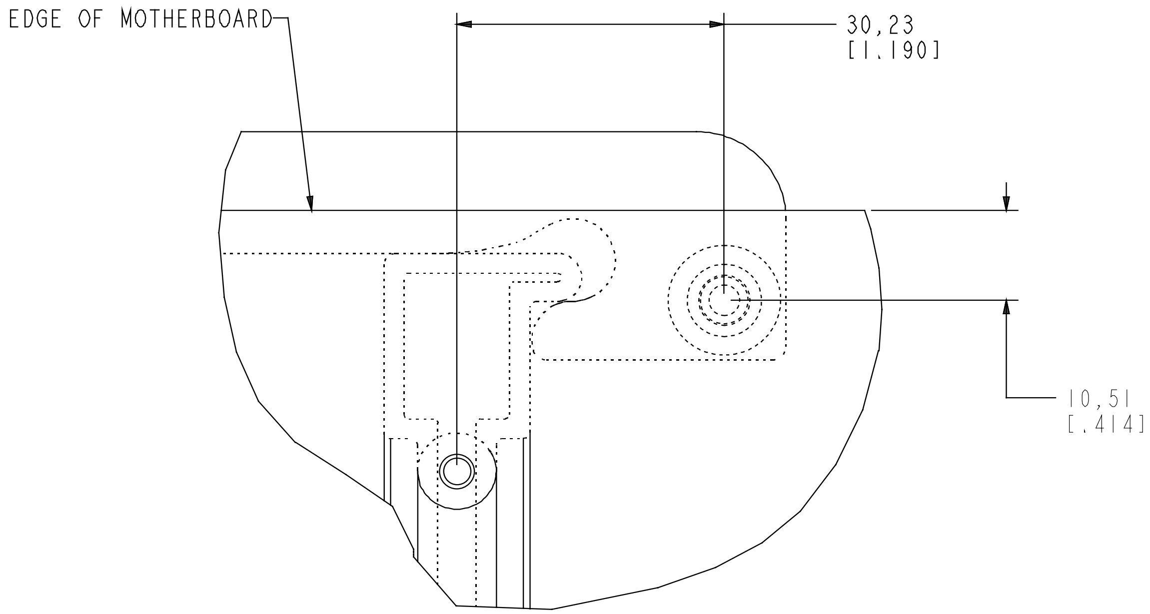

Figure 2.14b: Typical Rail Design Example (Part II)

2.4.2 Motherboard Insertion and Extraction

The NLX motherboard is installed by sliding it into the riser already mounted to the chassis. The rear of the motherboard is supported by the rear panel I/O shield. The two rails guide the motherboard to the riser connector. It is recommended that NLX chassis designers include a visual alignment marking, such as an arrow or stamped line, on the base of the chassis such that it will align with the “connector key datum” on the motherboard in its proper insertion location. Motherboard designers may want to put a mark next to the “connector key datum”, on the motherboard to facilitate the users alignment during insertion. This will help the installer guide the rails into the chassis mounts. The rail system and a single lever latch work together to insert and remove the motherboard from the riser connector. The lever latch must be installed on the middle rail to provide sufficient leverage and to not interfere with the required chassis keep-out areas for EMI clips.

2.4.3 Chassis Tab Features as Rail Guides

Tab features can be formed in the chassis bottom to serve as rail guides. Figure 2.15 illustrates this. If this method is used, consideration should be made for the size of the openings created to ensure compliance to regulatory Standards.

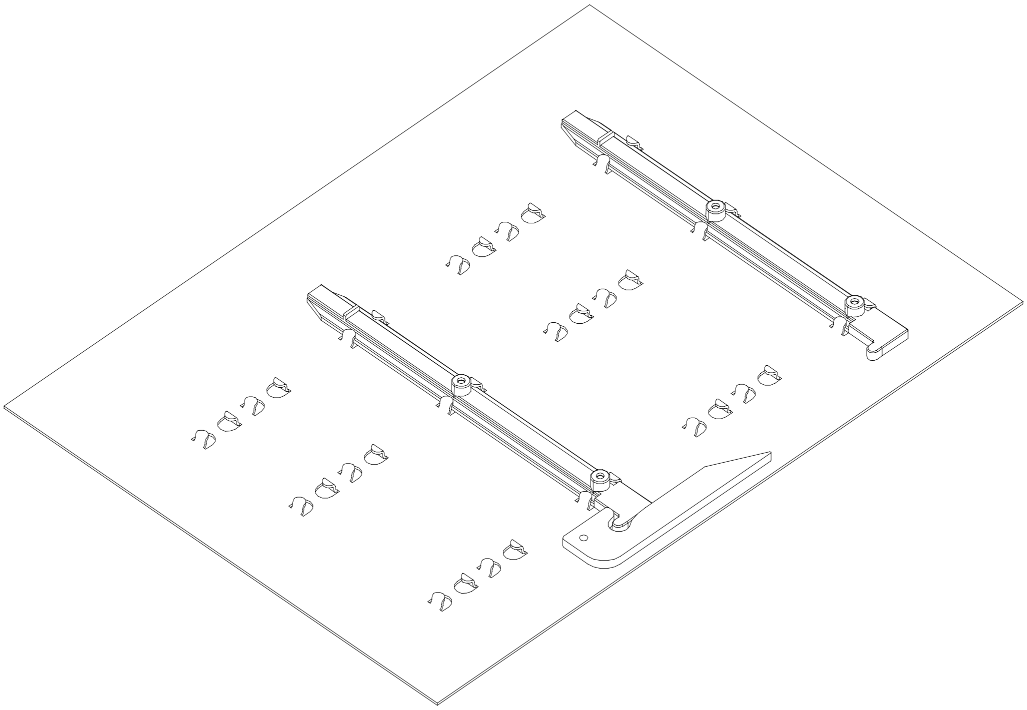

Figure 2.15: Chassis Tabs As Rail Guides

2.4.4 Clearance in System Rails

A build-up of tolerances in a motherboard rail guide system and the motherboard-to-riser card connection may occur. This can cause misalignment of the motherboard to the riser card. Therefore, sufficient clearance must be provided in a rail system design to allow the motherboard and riser connector to mate properly.

2.4.5 Rail, Clip and Lever Latch Design Examples

Figure 2.16 shows an example of a plastic injection molded rail. The overall length of the rail is 8.00 inch (203.20 mm). The rail has a tapered shape in the front to act as a lead in for the guides and a hook features at the rear for the lever-latch to interact with the rail during insertion and extraction of the motherboard. The center to center distance between mounting bosses is 3.400 inch (86.36 mm). The rail height at the bosses is .305 inch/7.75 mm, which leaves .008 inch/0.20 mm for the clip thickness. The clearance between the motherboard and the rail is .140 inch/3.56 mm except at bosses and support rib. The rib height is .313 inch/ 7.95 mm. The rail has molded in features so that EMI/RFI ground clip can be snapped on to it.

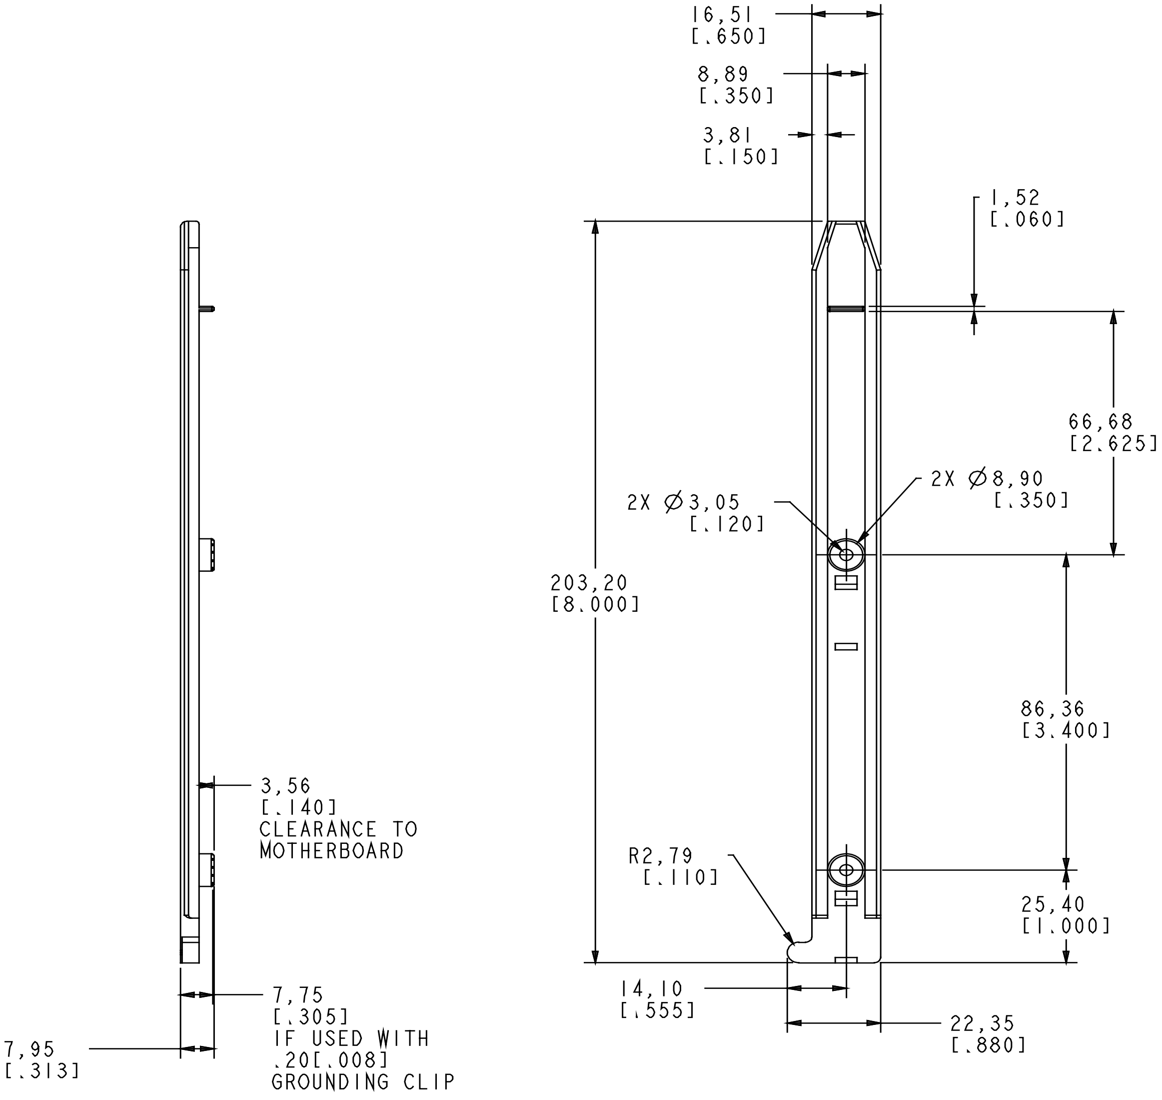

Figure 2.16: Molded Plastic Rail Example

Figure 2.17 shows an example of the EMI/RFI ground clips used in the rails. The clip is snapped on the rail before being attached to the motherboard. In general the material used for the clip should provide an electrical path to ground and should be of a resilient material such as stainless steel to ensure good contact. The clip can be made out of .008 inch/0.20 mm, half hard 301 stainless steel material. The 0.158 diameter hole provides clearance for a #6-32 screw. The radius portion of the clip assembly provides a sufficient contact surface for EMI/RFI grounding of the motherboard. There are four contact surfaces per assembly in this design method.

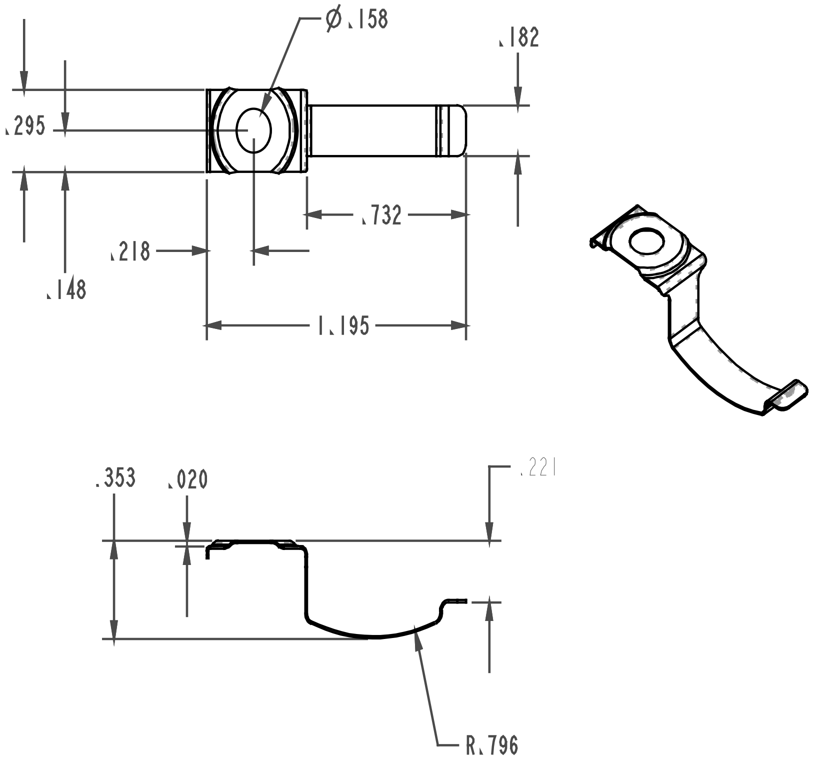

Figure 2.17: Example Rail Ground Clip

Figure 2.18 shows an example of an injection molded lever-latch. The lever-latch is used to give a mechanical advantage over the connector during insertion and extraction of the motherboard as well as latch the motherboard in during normal operation.

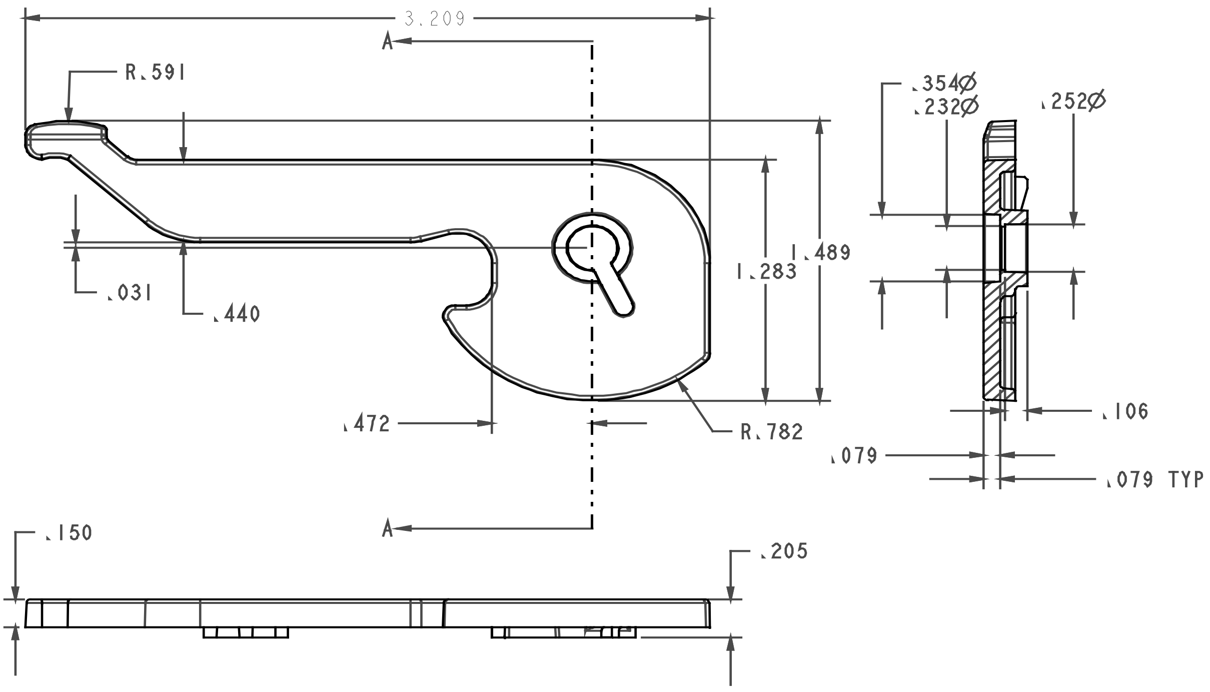

Figure 2.18: Example Molded Lever Latch

2.5 Motherboard / Riser Card Interface

The motherboard card edge mates with the riser card edge connector. The motherboard pushes the riser card at insertion and pulls it at removal. This causes stresses on the riser.

2.5.1 Mechanical Support Behind Riser Card

Chassis design must provide features for mechanical support directly behind the card connector to prevent any bowing of the card during insertion and extraction of the motherboard. There is no specific NLX requirement of how a riser card should be mounted in a chassis.

Figures 2.19 and 2.20 show two examples of the location of the card edge and other expansion slot connectors. The riser card can extend up to the front of the chassis and have connectors to interface with the front bezel.

On the riser card there are component height restrictions, shown as areas A and B. Areas A (two places) have a component height restriction of .200 inch maximum. This ensures that riser card components do not interfere with motherboard components. Area B component height restriction is .600 inch maximum.

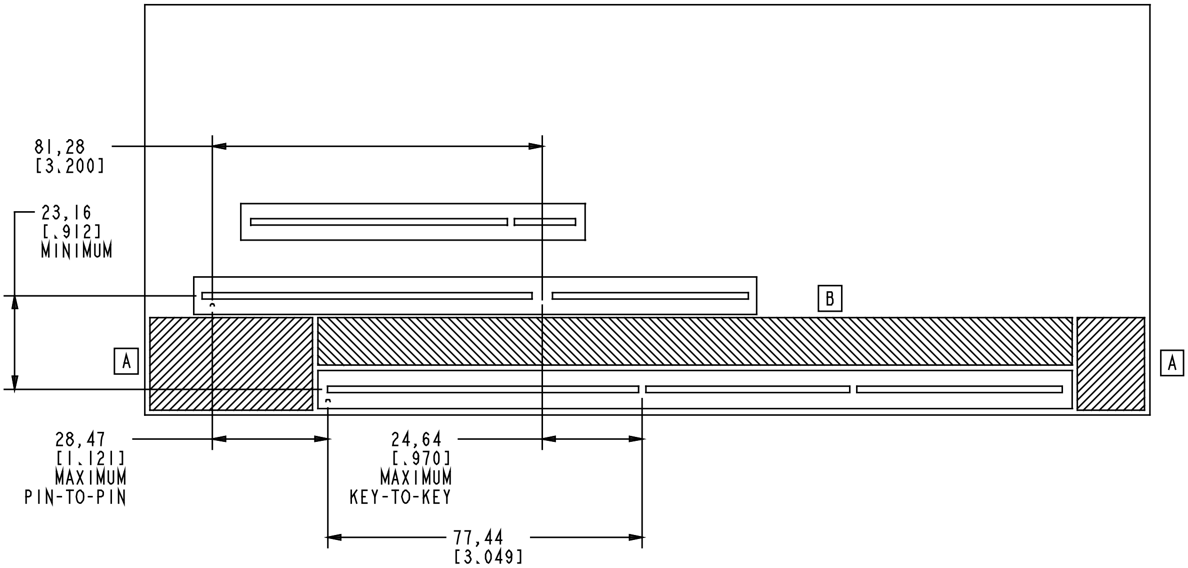

Figure 2.19: Riser Card with Keep-Out Areas Example #1

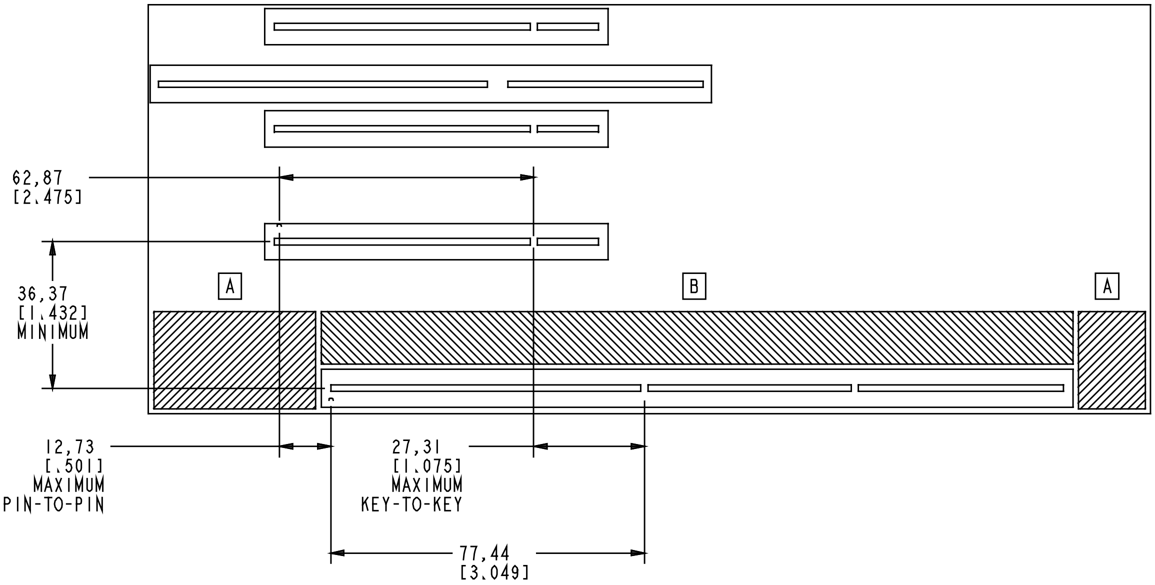

Figure 2.20: Riser Card with Keep-Out Areas Example #2

2.6 Back Panel I/O Shield

An NLX chassis must provide the specified I/O shield opening to ensure that any NLX board will fit into any NLX chassis. The I/O shield opening is fully defined so that all NLX I/O shields fit into the standard NLX shield opening. The I/O shield is assembled onto the motherboard; this assembly slides into a six-sided opening on the rear wall of the chassis. The shield location relative to the motherboard is fully defined to ensure compatibility between board and chassis. The I/O shield must contact clean, paint-free surfaces on the flanges of the I/O shield opening.

It is the responsibility of the system or board designer to provide mechanical support for the motherboard with the I/O shield. The I/O shield must be placed so that the motherboard card edge aligns with the riser card edge connector.

2.6.1 I/O Shield Opening Dimensions

Figure 2.21 shows the I/O shield opening dimensions from the riser card.

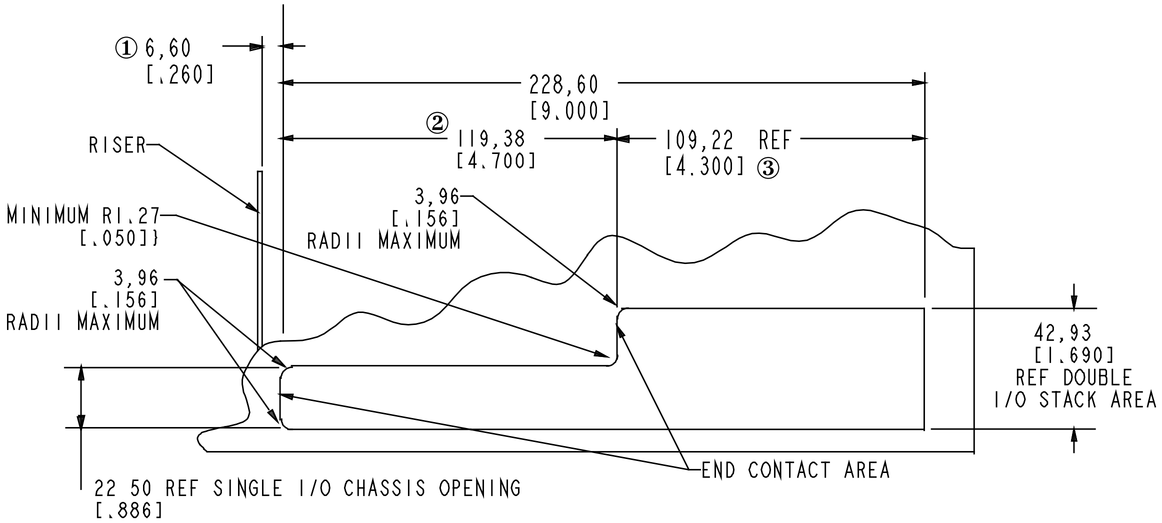

Figure 2.21: I/O Shield Opening Dimensions

2.6.2 I/O Shield Dimensions from Motherboard

Figure 2.22 shows dimensions of the I/O shield from the chassis bottom and motherboard.

2.6.3 I/O Shield Side View

Figure 2.23 illustrates the NLX I/O shield side view.

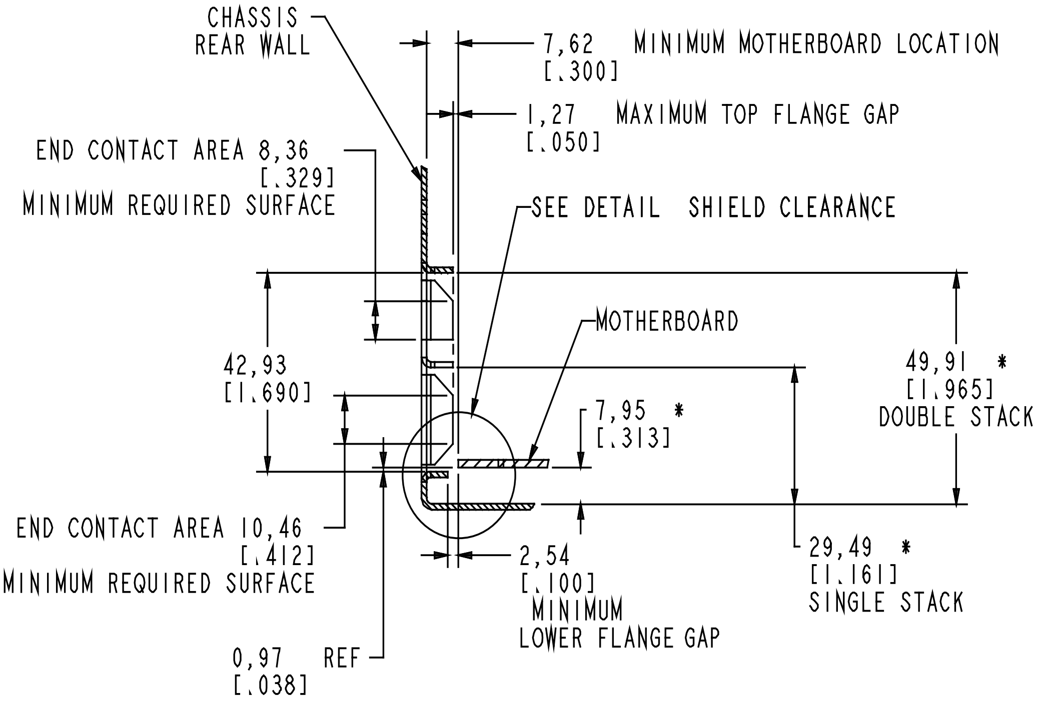

Figure 2.22: Dimensions of I/O Shield Opening from Chassis Bottom

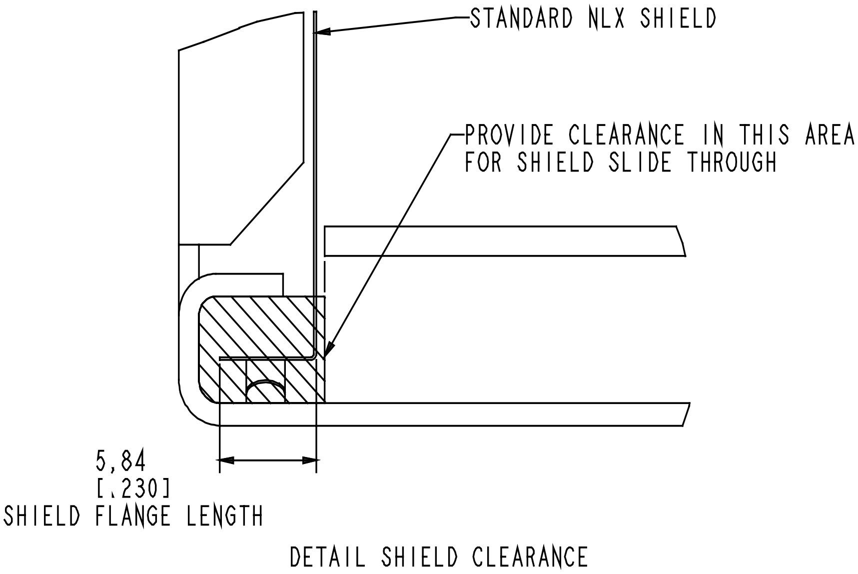

Figure 2.23: I/O Shield Side View

2.6.4 I/O Opening Flanges and Retaining Tabs

The I/O shield is surrounded by a flange used to provide EMI containment. The top of the I/O shield mates with a flat surface provided by the chassis back panel opening. The bottom of the I/O shield mates with the base of the chassis. A flat surface on the chassis is required for all surfaces of the I/O shield. This chassis surface must be flat over the maximum top flange dimension; outside of this dimension, the chassis designer is free to implement any guide or alignment features desired.

In Figure 2.24, items marked "1" show the flanges for EMI contacts, and items marked "2"show retaining tabs.

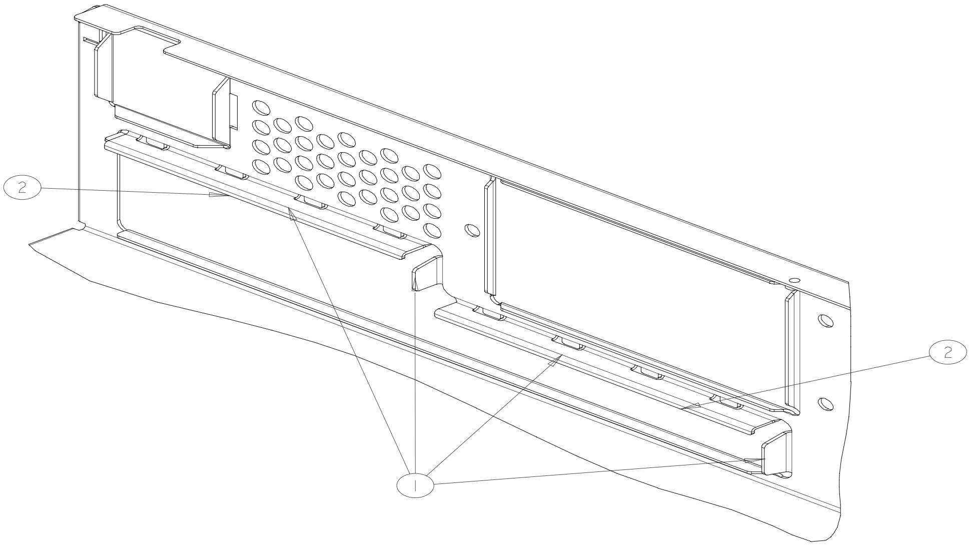

Figure 2.24: I/O Shield-to-Motherboard Snap Connection

2.6.5 Rear I/O RFI Shield Outline

The perimeter of the I/O RFI shield mates with the I/O opening of the chassis. The shield performs two functions: it provides the RFI seal around the I/O opening, and it provides mechanical support to the connector side of the motherboard.

Figure 2.25 shows an example I/O shield outline.

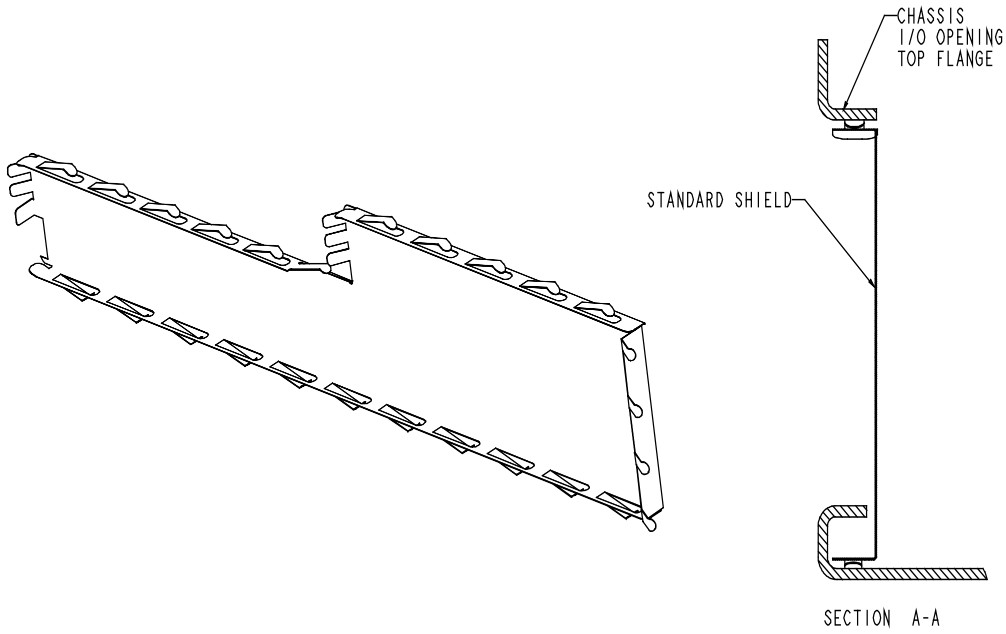

Figure 2.25: Example I/O Shield Outline

2.6.6 Alignment Mark on Chassis

A score line or stamped mark on the chassis at a location corresponding to connector key datum can simplify alignment of the motherboard to the riser card and easy mating.

2.7 Removable Side Panel

The left side panel of the NLX chassis should be removable for motherboard installation and removal.

Figure 2.26 shows a chassis design with a side panel that can be removed without a tool. There should be sufficient clearance at the opening to allow for tall components on the motherboard. A stiffener bar is added for structural strength.

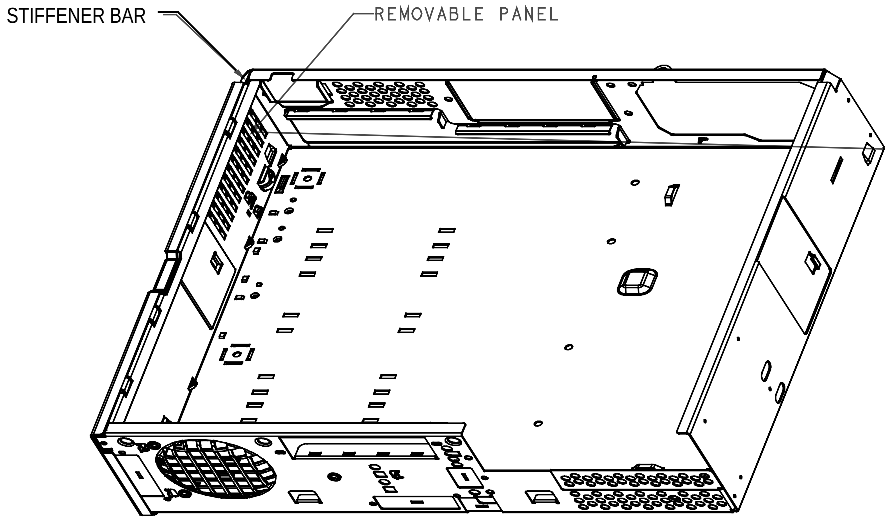

Figure 2.26: Removable Side Panel

2.8 A.G.P. Support

A fully compliant NLX chassis must provide an A.G.P. I/O opening. The opening must line up with the A.G.P. connector on the motherboard. Figure 3.15 shows the location of the connector on the motherboard. The chassis must provide an opening and a retaining method for an A.G.P. card. A full-size A.G.P. card must also have card guide support. Figure 2.27 shows an example of an A.G.P. implementation.

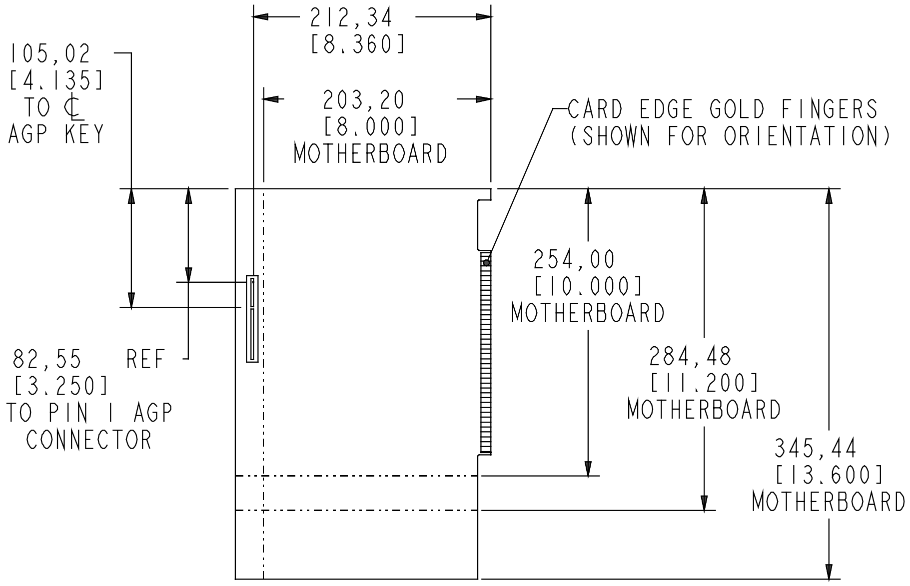

Figure 2.27: A.G.P. Connector Location On Motherboard

Figure 2.28 shows the chassis features for mounting the A.G.P. card to the chassis.

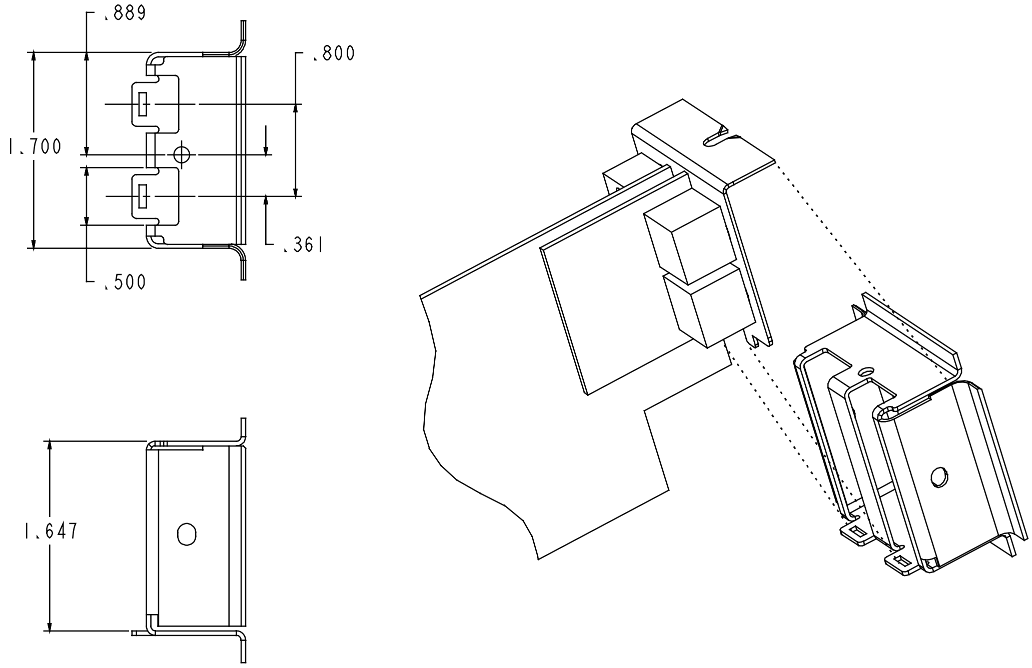

Figure 2.28: A.G.P. Card I/O Dimensions and Implementation Example

3. Power Supply Mounting

An NLX power supply has the same form factor as a PS/2 power supply. The fan is mounted inside the power supply housing and exhausts air from the rear of the chassis in most cases. The chassis must have a cut-out for the fan air exhaust and AC plug. The chassis must also have mounting holes for the power supply, per the NLX specification.

Figure 3.1 shows the power supply cut-out and mounting holes.

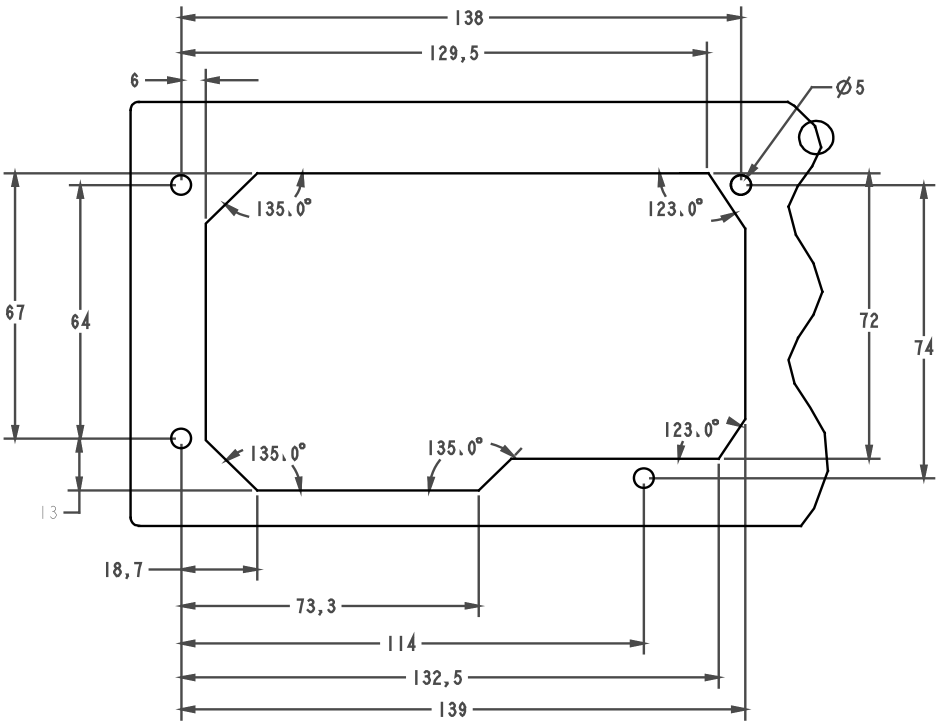

Figure 3.1: Power Supply Cut-out and Mounting Holes

4. Other Mechanical Design Considerations

4.1 Full Length Cards, Peripherals and Card Guides

Typical low-profile chassis provide three peripheral device bays and room for three full length add-in cards. Normally in an NLX chassis the peripherals are behind (to the right of) the riser card. A 3.5" peripheral can be installed in the add-in card section (to the left of the riser card), which prevents installing full length cards in this area.

A good card guide design serves multiple functions. It provides support for full-length add-in cards. It can also provide support for mounting a second fan in the front of the chassis. Care must be taken in the design of a card guide, because removal of add-in cards can be difficult if the card guide extends the full length of the add-in card.

Figure 4.1 shows an NLX card guide design.

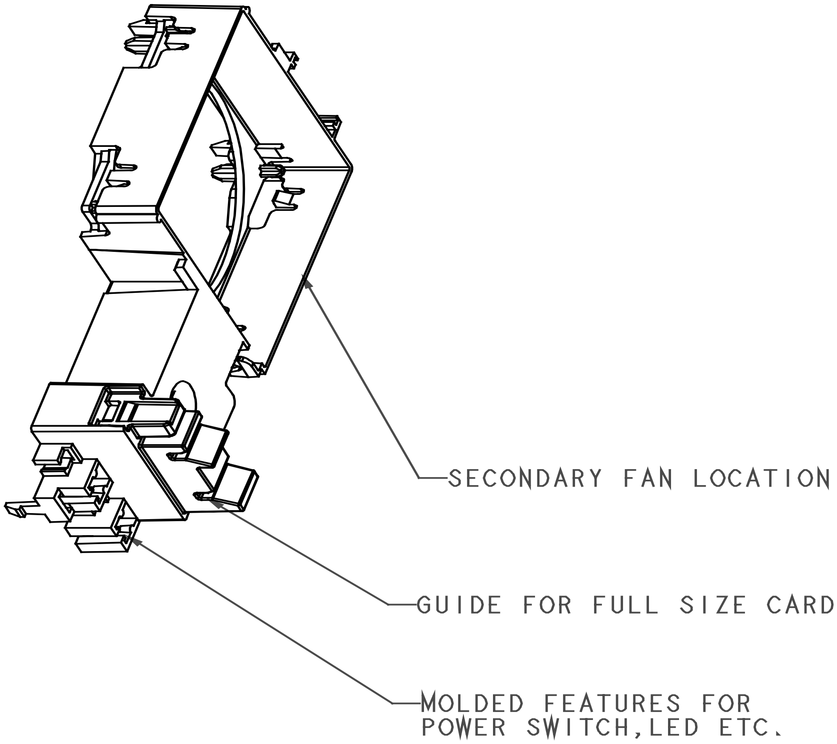

Figure 4.1: Card Guide Design Example

4.2 Fans and Vents

One fan is mounted inside the power supply. A second fan may also be required in the front of the chassis. The fan provides cooling air for the processor, peripherals, A.G.P. card, and other expansion cards.

Simulation and testing indicate that the chassis should be vented at the front, side, and rear. Front vents are provided at the front left side of the chassis. Rear venting is provided between the A.G.P cut-out and the add-in card area. Left side vents are provided in the rear one-third of the side panel to aid in cooling an A.G.P. card.

Figures 4.2 and 4.3 show the vent locations in an NLX chassis.

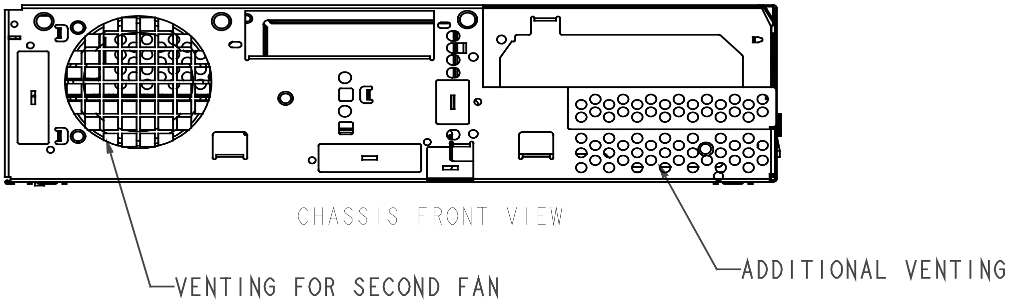

Figure 4.2: Chassis Front View

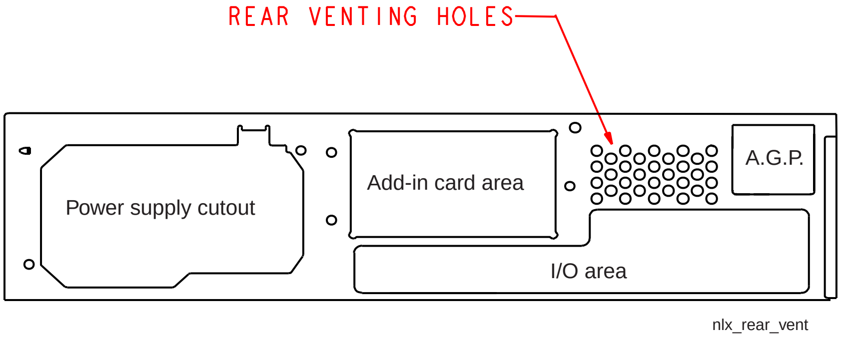

Figure 4.3: Chassis Rear View

4.3 Riser Card and Front Panel Interface

The riser card can extend to the front of the chassis. Control and indicator (e.g., power switch, LEDs) connectors can be added to the riser card to interface with the front bezel. This simplifies internal cabling and facilitates serviceability.

4.4 Cable Routing Over Riser Card

The riser card and its support in the chassis should have cut-outs so that cables can be routed over them from the left side of the chassis to the right side of the chassis. Cabling between the two sides is often required to use add-in cards that require cabling to peripherals, as in the case of SCSI devices.

5. Thermal Design Considerations

5.1 Introduction

The newest processors, chipsets, and memory pose significant thermal challenges to the system designer. As the market transitions to higher-speeds and greater bandwidths with enhanced features, the heat generated by these devices continually increases, placing complex cooling demands on the system.

The goal of this section is to provide a brief introduction to:

- general thermal design principles, including estimating the airflow required for a given heat load, and
- the concept of thermal resistance and its application to heat dissipation through integrated circuits and heat sinks.

Most systems depend on tube axial fans to cool components, so fan size, airflow and impedance, and fan location will be emphasized. The correct fan size cannot be chosen without knowing the chassis airflow impedance. A chassis with greater airflow impedance requires a larger fan to move a given amount of air and likewise a well flowing chassis requires a smaller fan to move the same amount. Minimizing airflow obstructions and optimizing airflow patterns including the power supply are important, so factors affecting chassis airflow and determining the system characteristic curve will be explained.

The power supply usually has the largest impact on system level cooling because the power supply fan is most often the only fan in system designs. Therefore, power supply selection is critical to proper system performance. Selecting a well-designed power supply can *double* the system’s airflow, allowing more flexibility in the chassis design.

Power supplies and/or motherboards with advanced thermal management techniques such as fan speed control and Advanced Configuration and Power Interface (ACPI) are becoming more popular to control both acoustics and component temperatures. If the power supply or system fan is speed controlled, the thermal design should account for various load and temperature combinations.

Although heat sink selection is explained in greater detail in the Pentium® II processor and Intel chipset application notes, passive, active, and liquid cooled heat sinks and manufacturing methods and their relative cooling ability are briefly mentioned for reference.

Add-in cards and peripheral cooling requirements are often ignored or forgotten during the chassis design. The newest A.G.P. graphics cards and DVD drives, for example, dissipate significantly higher power than previous generations. Future graphics controllers have the potential to increase to 15 W. These components now require some airflow rather than the previous natural convection cooling scheme. It is important to understand what requirements the peripherals and add-in cards have that may be included in the system.

Acoustics are an important consideration, because a cool-running system may not always operate quietly. A basic guideline requires systems operating with less than 300 W must be quieter than 45 dBA at 23 °C. Refer to ISO 7779 for complete acoustic testing details.

After the system design is complete and prototypes are available for evaluation, the system must be tested to ensure it meets specified criteria for component temperatures and required airflow. Various techniques are available to measure temperature and airflow. The most common device used to measure component temperatures is the thermocouple. Hot-wire anemometers, static pressure tubes, and flow chambers are some of the tools used to measure airflow. Whichever method is selected, careful setup is required to provide accurate results.

Because of the complexity of varying chassis designs, modifications may be necessary to some of the suggestions given here to achieve an effective cooling scheme. The system’s cooling scheme must ensure that all components and peripherals remain within their specified operating temperature ranges. In addition to the suggestions discussed in this document, the designer should be aware of the system-level requirements described in the *PC 99 System Design Guide*.

5.2 Thermal Design Principles

5.2.1 General Principles

The first step in defining an acceptable cooling solution is to estimate the total airflow required to cool the entire system (not just the processor).

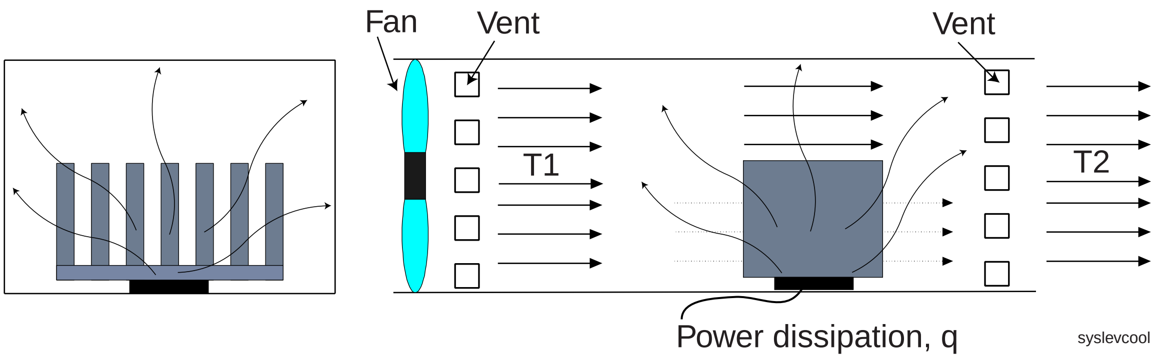

Figure 5.1: Simplified System Component Cooling

To do this we refer to the 1st Law of Thermodynamics (Conservation of Energy) for a steady state, steady flow process:

$$ {}_1q_2 \equiv \Delta h + \Delta K.E. + \Delta P.E. + {}_1w_2 \quad (\mathrm{kJ/kg}) $$

Where: q = heat dissipated in the system
Δh = change in enthalpy
$\Delta K.E.$ = change in kinetic energy
$\Delta P.E.$ = change in potential energy
and, w = work done by the system

Assume the change in kinetic and potential energy of the airflow is zero, and no work is performed by the system. Then factoring in the mass flow of the air, this equation can be rewritten as:

$$ {}_1q_2 \equiv \dot{m}\Delta h = \dot{m}c_p(T_2 - T_1) = \dot{\mathrm{v}}\rho c_p(T_2 - T_1) $$

now, solving explicitly for volumetric airflow, we have

Equation 1: $$ \dot{\mathrm{v}} = \frac{{}_1q_2}{\rho c_p (T_2 - T_1)} $$

The airflow of a system can be significantly different that the air flow of just the power supply or system fan. A system can, in reality, experience somewhere between 30% to 50% restriction of airflow due to system impedance. Therefore, a fan capable of providing even more airflow than Equation 1 indicates is needed to overcome the system impedance and cool the system. For well-designed chassis an airflow increase of approximately 30% is needed to account for the system impedance. If possible use the measured DC power of the system as the heat load in Equation 1. The AC power can be used as an approximation, but the inefficiency of the power supply makes the AC power value larger than the DC power value resulting in an inaccurate airflow requirement. Both ρ and c_p are evaluated at room temperature, a correction factor is necessary for other ambient air conditions.

5.2.2 Thermal Circuit

There are three fundamental means of heat transfer: conduction, convection, and radiation. Heat sinks are used to increase the effectiveness of the heat transfer from the hot solid surface to the cool ambient through conduction and convection. This is accomplished primarily by increasing the effective surface area that is in direct contact with the coolant, air. Remember, air is a poor conductor of heat. The increase in component surface area allows more heat to be conducted from the component silicon and then dissipated into the system environment. This lowers the device operating temperature, ultimately ensuring the device temperature remains well below its maximum allowable temperature specification.

A typical example of a thermal model, the electric circuit analog to heat conduction, is shown in Figure 5.2. The goal of this section is to provide a basic understanding of fundamental heat sink concepts. Some key terms are defined below:

T_j, T_c, T_s and T_a - The temperatures at the junction, component case, heat sink base and ambient air.

q - The maximum electronic power dissipation of the electronic component.

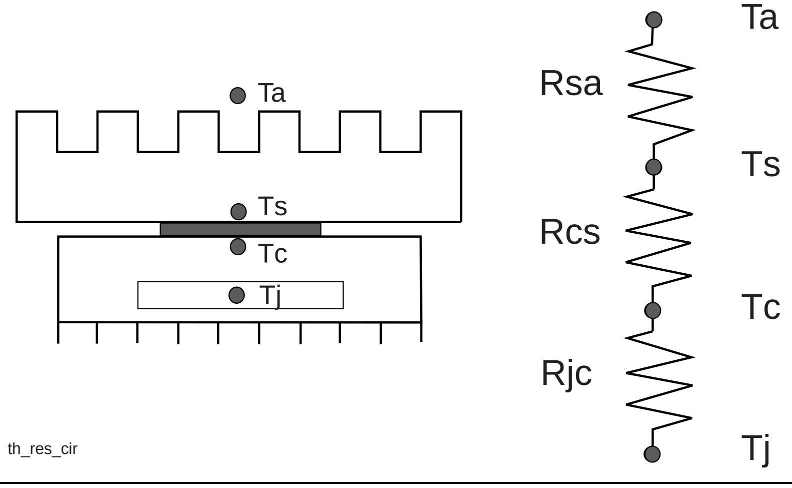

Figure 5.2: Thermal Resistance Circuit

The best figure of merit for heat sink performance is the overall thermal resistance R_{ja}. In this model the heat is assumed to flow serially from the junction, through the case and heat sink, to the ambient air. The overall thermal resistance is defined as follows.

$$R_{ja} = R_{jc} + R_{cs} + R_{sa} = \frac{(T_j - T_c)}{q} + \frac{(T_c - T_s)}{q} + \frac{(T_s - T_a)}{q}$$

R_{jc} represents the thermal resistance between the junction and the case of the component and is typically not under the control of the system designer. R_{cs} represents the resistance of the thermal interface material and R_{sa} represents the heat sink thermal resistance between heat-sink and air. The smaller the resistance value, the more power a device can dissipate without exceeding its junction temperature. In reality, the flow of heat is three dimensional (not one dimensional as shown above), but the model shown above still works fairly well.

The *interface resistance* (R_{cs}) depends on many factors including the surface flatness, the surface finish, the mounting pressure, the contact area, and the type of interface material used along with its thickness.

The interface resistance is often expressed as

$$R_{cs} = \frac{t}{\kappa\ A}$$

where "t" represents the material thickness, "A" represents the contact area, and "κ" is the material thermal conductivity.

The *heat sink resistance* is a function of the airflow and the effective surface area of the heat sink.

$$R_{sa} = \frac{1}{h_c A} \text{ in } \frac{^\circ\mathrm{C}}{\mathrm{W}}$$

The value of the heat sink resistance can be reduced by increasing the effective surface area of the heat sink, A, or by increasing the convective heat transfer coefficient, h_c. Care must be exercised here as h_c is very sensitive to airflow. If the added surface area chokes off the airflow through the heat sink, the value of h_c may be reduced so much that the product of h_cA actually decreases!

A quick estimate of the necessary heat sink performance can be found in the following manner. For the systems we commonly deal with, the maximum processor case (or heat plate) temperature is specified along with the maximum processor dissipated power. Initially, ignore the thermal interface resistance and estimate the heat sink resistance using:

$$R_{sa} = \frac{T_c - T_a}{Q}\left(\frac{^\circ\mathrm{C}}{\mathrm{W}}\right)$$

Radiation effects are typically ignored initially for forced convection cooling schemes. However, radiation is important in natural convection cooling schemes and may be responsible for up to 25% of the total heat transfer. Unless the component is facing a hotter surface nearby, the heat sink surface should be painted or anodized to enhance radiant heat transfer.

5.3 Fans

Fans implement the forced convection approach to cooling. Stated simply, the greater the air velocity over the surface of a component, the greater the heat transfer from that component. Fans may differ in their characteristics, and therefore a prudent choice of fans can optimize both airflow and acoustics.

Fans can be used to blow air into (pressurize) or out of (evacuate) the chassis depending on which direction they are installed.

Pressurizing the chassis with a fan delivers cool room ambient air onto any location where it is needed to enhance heat transfer.

Evacuating induces a negative pressure (relative to room ambient) inside the chassis, which draws air in through the vents. This inflow of air from the vents is pulled through the chassis across hot components and is exhausted out the fan. Fan Types

There are several types of fans to consider for system cooling; *tube axial* and *radial*. Tube axial is the most commonly used type throughout the computer industry. Axial fans typically cost less and generally push more air at a common back pressure. Radial fans, however, are much less susceptible to variations in back pressure and often have restricted openings which can focus needed cooling air directly at hot components. When power dissipation is highly concentrated, a radial fan (blower) may be a reasonable option. Figure 5.3 shows a typical axial fan characteristic curve, and the effect of running the fan at different speeds (or voltage levels).

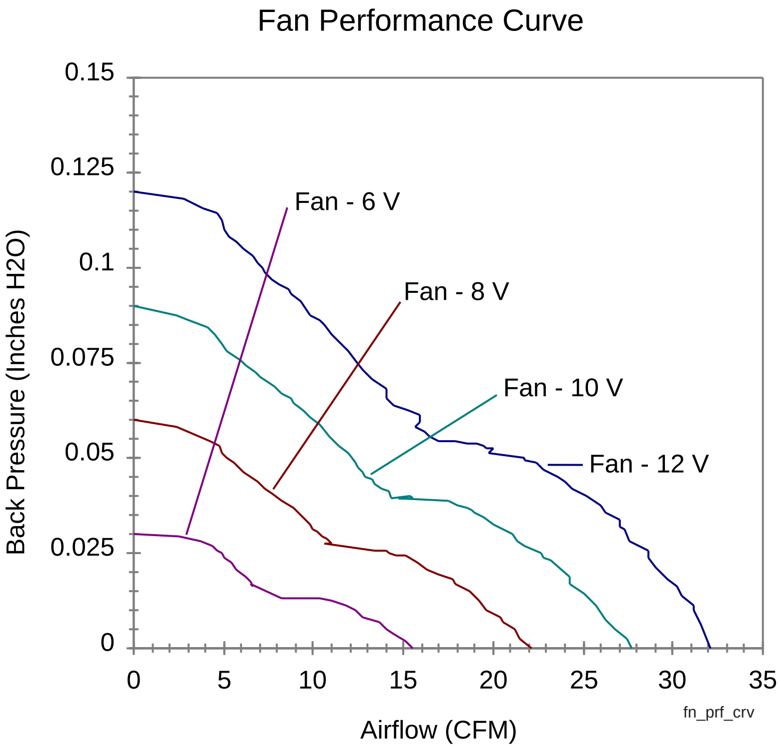

Figure 5.3: Typical Axial Fan Characteristic Curve (for Various Voltages)

5.3.1 Parallel and Series Fan Combinations

Multiple fans can be utilized in two combinations, parallel and series.

- Two fans in parallel, $Q = Q_1 + Q_2$ at zero back pressure
- Two fans in series, $p = p_1 + p_2$ at zero airflow

An example of a parallel fan combination is a system fan and a power supply fan both either pressurizing or evacuating a chassis. Ideally, a parallel fan configuration doubles the system airflow. An example of a series fan combination is a system fan blowing air into the chassis and a power supply fan exhausting air from the chassis. Ideally, a series fan configuration doubles the system's ability to overcome built up backpressure. In reality, due to venting, leakage and design compromises, when we employ multiple fans we often are implementing a combination series/parallel configuration. The effect of employing series/parallel fan configurations is shown in Figures 5.4 and 5.5.

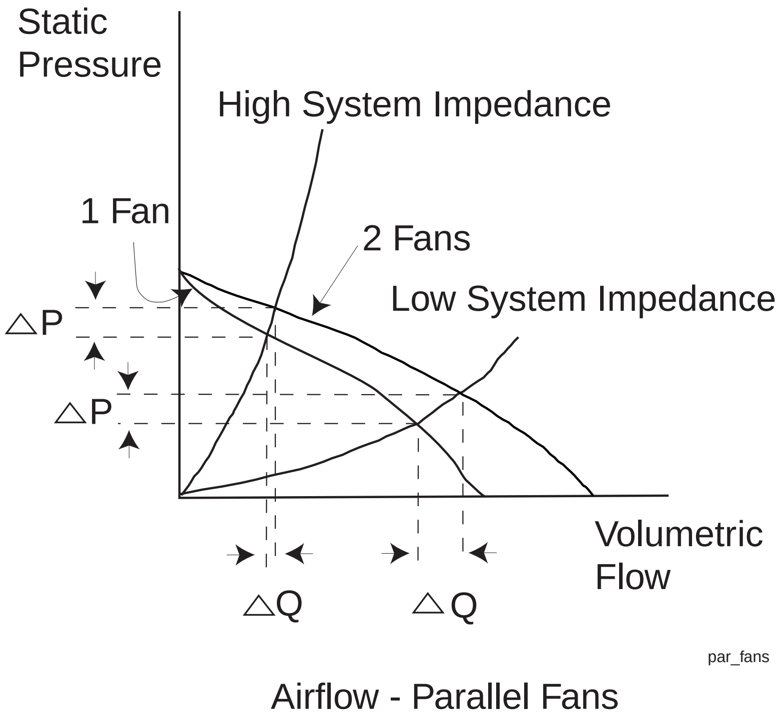

Figure 5.4: Performance Curves for Parallel Fan Combination

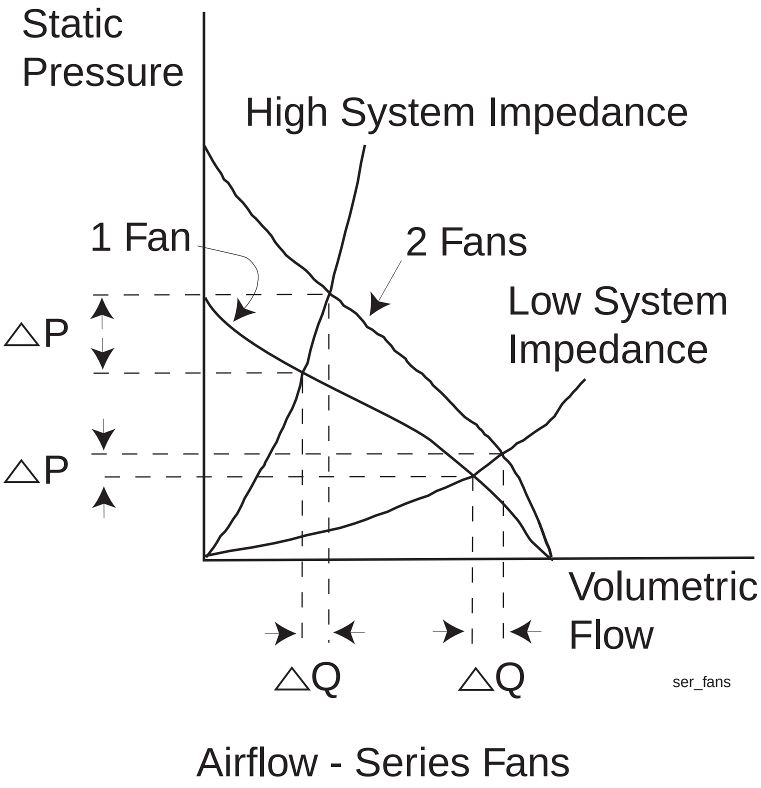

Figure 5.5: Performance Curves for Series Fan Combination

Employing multiple (identical) fans in a system does provide some marginal increase in airflow. The exact amount depends on many factors, including fan speed and configuration, as well as chassis airflow impedance. If the fans are not identical, then the figures will change slightly, but the trends will be the similar. The general rule is, if the chassis has high impedance, place the fans in series. If the chassis has low impedance, place the fans in parallel.

5.3.2 Fan Relationships

The relationships in Tables 5.1 and 5.2 below describe how airflow, pressure, and power vary with fan speed and with variations in both fan speed and fan diameter.

Table 5.1: Fan Laws: Variable Speed - Constant Diameter

Description	Relationship
Airflow vs. R.P.M.	$Q_1/Q_2 = n_1/n_2$
Pressure vs. R.P.M	$p_1/p_2 = \left(n_1/n_2\right)^2$
Power vs. R.P.M.	$h.p._1/h.p._2 = \left(n_1/n_2\right)^3$

Table 5.2: Fan Laws: Variable Speed - Variable Diameter

Description	Relationship
Airflow vs. R.P.M.	$Q_1/Q_2 = \left(D_1/D_2\right)^3 \left(n_1/n_2\right) = \left(D_1/D_2\right)^2 \sqrt{p_1/p_2} \sqrt{\rho_2/\rho_1}$
Pressure vs. R.P.M	$p_1/p_2 = \left(D_1/D_2\right)^2 \left(n_1/n_2\right)^2 \left(\rho_1/\rho_2\right)$
Power vs. R.P.M.	$h.p._1/h.p._2 = \left(D_1/D_2\right)^5 \left(n_1/n_2\right)^3 \left(\rho_1/\rho_2\right)$

Obviously all tube axial fans used in systems today are of constant diameter from the front of the fan to the back of the fan. Key points of constant diameter fan relationships to remember are:

- Airflow increases linearly with speed
- Pressure increases with the square of the speed
- Power increases with the cube of the speed

Understand that increasing the fan speed to increase airflow results in a much larger increase in pressure. If increased airflow is desired, consider increasing the fan diameter from 80 mm to 92 mm instead of increasing the speed. Cost must be considered because generally 92 mm fans are more expensive than 80 mm fans. However, a 92 mm fan operating at the same flow rate as an 80 mm fan is approximately 6 dBA quieter.

5.4 Airflow Impedance

Air flowing through a computer chassis encounters frictional resistance, known as airflow impedance. This impedance creates a pressure drop in the chassis which roughly obeys Bernoulli's principle and is found to vary approximately with the square of the velocity, or since $Q = A \cdot v$, with the squared volumetric airflow. Plotting pressure loss versus volumetric flow rate, which results in the system characteristic curve, can show the relationship. The point about this behavior is that if one data point on the curve is known, the system's overall performance can be predicted. When the system characteristic curve is superimposed on the fan performance curve, the operating point of the system is specified explicitly. The concept is demonstrated in Figure 5.6 where different power supplies were compared with different chassis.

Here are some additional guidelines to consider when assessing system airflow issues:

- The operating point should be chosen to the right of the local maximum peak on the fan pressure curve to avoid pressure and volumetric flow rate fluctuations.
- Choose a fan with a steep characteristic curve to maintain constant volumetric flow even with variable system impedance.
- Avoid obstructions near the inlet and exhaust of the fans as these tend to decrease airflow and increase system noise.
- Use fan speed control whenever possible and cost effective. This yields adequate thermal margin and provides a significant acoustic advantage.
- Power supply cables and drive signal cables should be kept short and properly folded.

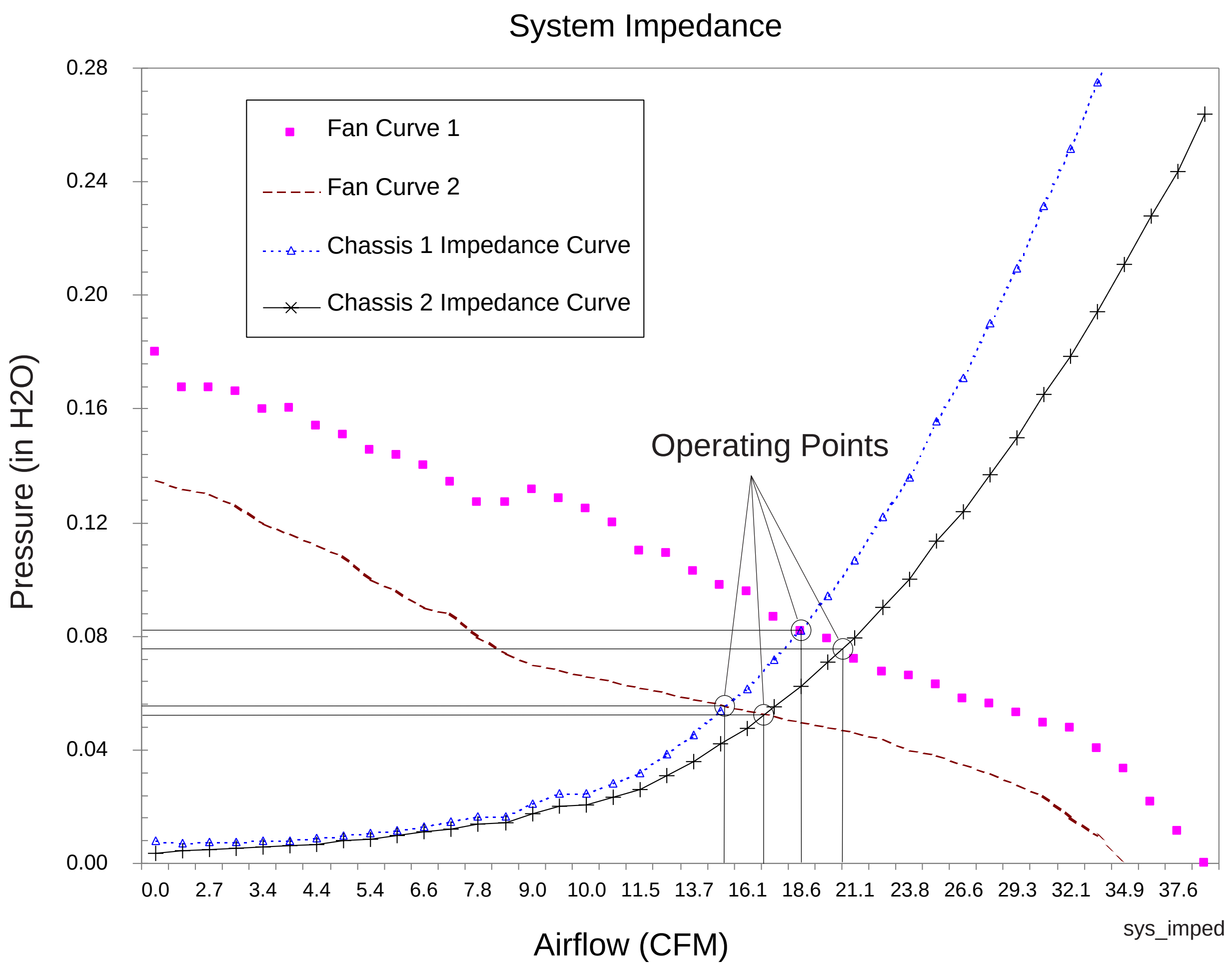

Figure 5.6: System Characteristic Curve

5.5 Power Supply Characteristics

The power supply is among the most influential components in the system cooling design. The chassis venting scheme may be well designed, but if the correct power supply is not selected, the system will not cool the processor, chipset, memory, and/or the peripherals. The power supply and any system fans must provide enough airflow to cool the system heat load as outlined by Equation 1 in Section 5.2.1.

Key considerations when selecting/designing a power supply:

- Evacuate the chassis (rather than pressurize it) with the power supply fan. The advantage of evacuating the chassis is that cool room ambient air can be delivered (via vents) to any location where it is needed to enhance heat transfer. Evaluation has shown evacuating produces greater cooling than pressurizing using the same fan with proper implementation.
- All vents should have a minimum free area ratio of 60%. Consult the EMI design guidelines to ensure vent designs comply with all applicable regulations.
- Implement a wire fan grille rather than the common stamped sheet metal designs because the airflow impedance is reduced.
- Minimize the power supply component height to keep their profile low and streamlined. This reduces the overall supply impedance while still maintaining effective power supply cooling.
- Keep the power supply cables short to reduce their airflow obstruction.
- Select a power supply with the highest airflow possible. A well-designed power supply has lower airflow impedance, allowing a smaller, quieter fan for cooling. The poorly designed supply requires a larger, louder fan to maintain the same airflow due to its greater airflow impedance.

Figure 5.7 depicts the power supply impedance curve and the associated fan curve of three different power supplies. The point where the fan curve intersects the power supply impedance curve defines the operating point. Power supplies 1 and 2 (ATX and PS2 style, respectively) flow approximately twice as much as power supply 3 (ATX style). Note power supply 3 has a smaller fan and higher airflow impedance resulting in the lower airflow.

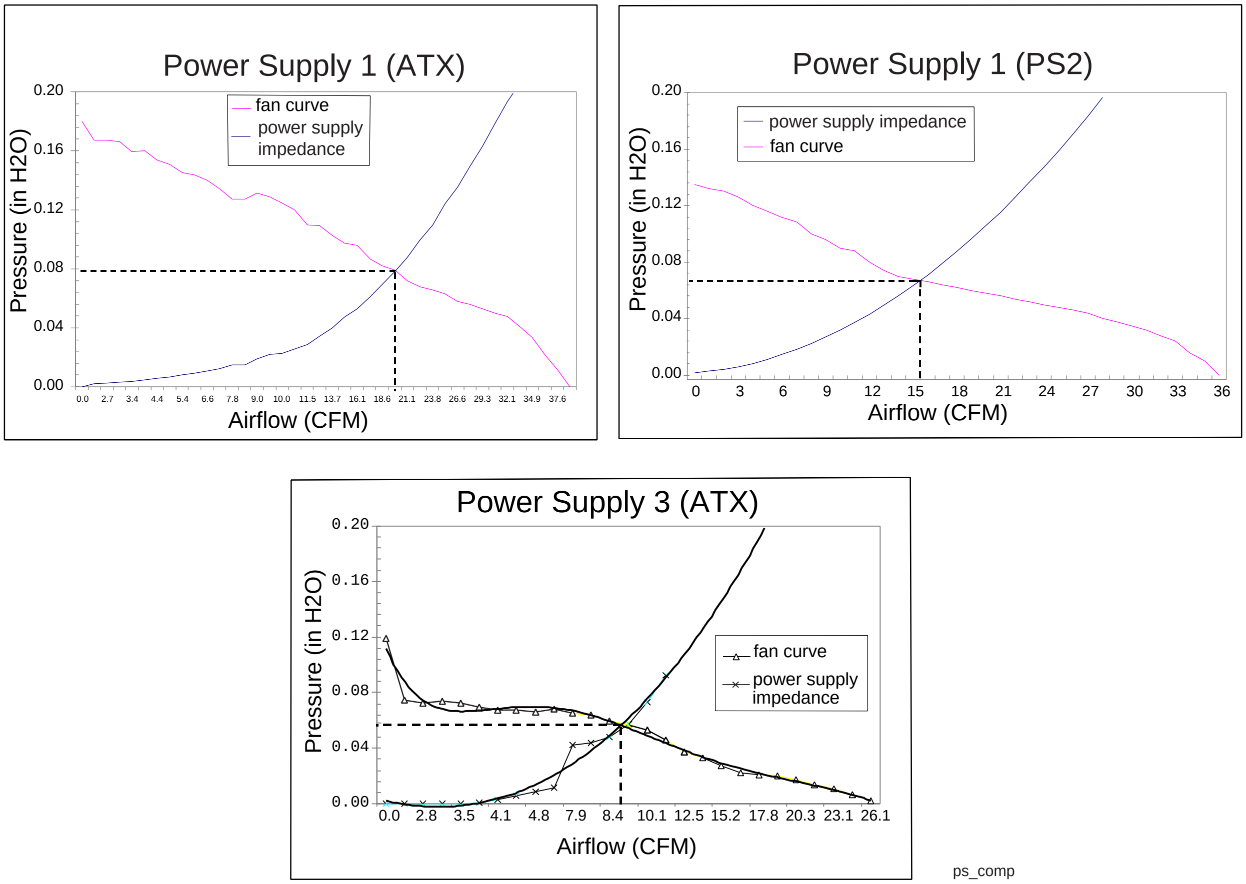

Figure 5.7: Power Supply Performance Comparison

5.6 Advanced Thermal Management

5.6.1 Fan Speed Control

Fan speed control allows a system to vary its airflow as changes in load and/or temperature occur. Fan speed control circuit ideas have been around for some time but were not always utilized possibly due to their cost and increase in system complexity. However, computers are now incorporating hotter processors and peripherals requiring greater airflow while at the same time customers are requesting quieter systems. These competing design constraints have led to a resurgence of fan speed control options. Acoustically, fan noise increases directly with fan speed and is a major contributor to total system noise. For systems that incorporate fan speed control, proper speed regulation is important since it is desirable to achieve low acoustic levels without overheating components. The fan speed control circuit should be designed such that it monitors temperature at a component (or several components) and adjusts fan speed as necessary to maintain the required thermal margin. Three distinct design options should be considered:

Discrete Digital Switches

If airflow requirements can be confined to a discrete number of fan speeds, this option is the cheapest and easiest to implement.

Analog Linear Control Between Two Guard Bands

For fans used in most systems, speed control can usually be accomplished by varying the voltage level at the fan’s power terminals (many power supplies/fans come equipped with this feature). An operating voltage range example for an 80 mm, 30 CFM, .14 amp fan might be 8 V to 12 V DC, corresponding to 1650 rpm and 2500 rpm, respectively.

Pulse Width Modulation Schemes

This is a digital variation on the second option. Consider this option if the fan needs to be varied from some minimum speed (presumably set for the system sleep state) to some maximum speed (needed for a fully loaded active state).

No matter which fan speed control method is chosen, the following issues should also be noted:

- The location where temperature is monitored is important (sensing critical component case temperatures is recommended). Whatever location is selected, it should represent the thermal state of the entire system.
- A driver circuit for the fan must be included.
- Some fans need a minimum starting voltage (see fan specification).
- Fan noise increases with fan operating voltage (speed). Minimum fan noise occurs at maximum fan power efficiency (see fan specification).
- If the fan is not speed controlled, at what voltage (speed) level is it operating? In this case since it is not possible to vary fan speed, choose the lowest rated fan speed that will cool the system under worst-case loading/temperature conditions.

If fan speed control is implemented, the thermal design should account for various load and temperature combinations. Component temperatures should be verified to ensure the thermal design meets specification under these load and temperature combinations.

5.6.2 Advanced Configuration and Power Interface (ACPI)

ACPI provides the control policy for a PC to measure temperatures of critical internal components. This allows local temperature sensing circuits to be read by the BIOS and operating system. The sensing circuits may use device-mounted thermistors, thermocouples or temperature diodes. Critical ‘triggers’ can be programmed to cause alarm events which instigate a (graded) cooling policy.

Cooling policies can be either ‘Passive’ (where performance is limited to reduce heat generation) or ‘Active’ (where fan speed or on/off control is used to limit temperature as heat increases). The policies can be mixed and can be used in any order.

In Figure 5.8 below, as the processor temperature starts to rise, the operating system senses a critical temperature and starts to limit power dissipation. If temperature continues to rise, the cooling fan is switched on (slow). Any further increases in temperature results in the operating system increasing the fan speed. If the fan speed maximum is reached and the temperature continues to rise, the system shuts down entirely in order to prevent damage.

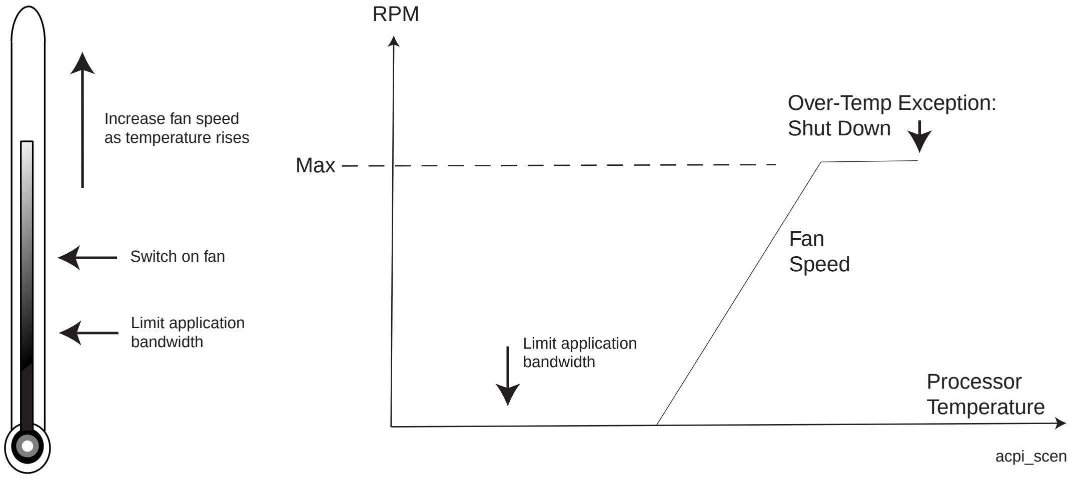

Figure 5.8: Example ACPI Scenario

5.7 Heat Sinks

5.7.1 Introduction

This section provides an overview of heat sink types and categories. Refer to the *Pentium® II Processor Application Note - Thermal Design Guidelines* and the corresponding thermal application note for the chipset for a detailed discussion concerning heat sink selection. As with other key system components, the selection process for thermal management components includes consideration of such issues as cost, ease of assembly, manufacturability, and upgradability as well as functional performance.

5.7.2 Heat Sink Categories

One way to classify heat sinks is by the cooling mechanism used to remove heat from the heat sink itself.

- *Passive Heat Sinks:* Passive heat sinks are used in natural convection applications and applications where heat dissipation does not rely on a specified local air velocity.
- *Semi-Active Heat Sinks:* These are “passive” heat sinks that leverage increased airflow from existing fans in the system.
- *Active Heat Sinks:* These heat sinks employ dedicated fans, configured for either impingement or cross flow. Since fans are an integral part of this type of heat sink, the reliability of this cooling solution is tied heavily to the fan’s reliability.
- *Liquid Cooled Heat Sinks:* These heat sinks typically incorporate tubes-in-block designs or milled passages in brazed assemblies channeling cooling liquids such as water or oil.
- *Phase Change, or Recirculating, Heat Sinks:* These include two-phase systems that employ a boiler and condenser in a passive, self driven, mechanism. Heat pipes are the most common example of this type of heat sink. They passively transfer heat from a source to a sink where the heat is dissipated. The heat pipe is an evacuated vessel that is partially filled with a minute amount of water or other working fluid. As heat is directed into the device, the fluid is vaporized creating a pressure gradient in the pipe forcing the vapor to flow along the pipe to the cooler section where it condenses, giving up its latent heat of vaporization. The working fluid is then returned to the evaporator by capillary forces developed in the heat pipe's porous wick structure, or by gravity.
- *Thermoelectric coolers (TECs):* These are solid state heat pumps that utilize the Peltier effect. During operation, DC current flows through the TEC causing heat to be transferred from one side of the TEC to the other, creating a cold side and a hot side.

5.7.3 Heat Sink Types

One can also classify heat sinks in terms of manufacturing methods and their final form.

- *Stampings:* Stamped heat sinks provide a low cost solution to low power density thermal problems. Copper or aluminum sheet metal can be stamped into any desired shape. Attachment features and interface materials can be added with ease during the manufacturing process.
- *Extrusions:* Molten metal is drawn through a die in the desired direction of the heat sink fins and then cooled. Extruded heat sinks allow the elaborate formation of two-dimensional shapes capable of dissipating large heat loads. This is the type most commonly attached to processors.
- *Bonded/Fabricated Fins:* Most air-cooled sinks are convection limited. Bonded fin heat sinks maximize available heat transfer surface area by bonding planar fins on a grooved, extruded base plate. Due to the unique manufacturing process, aspect ratios (fin height to width ratio) of 20 to 40 can be easily achieved, greatly increasing the heat sink's cooling capacity.
- *Castings:* This technology is used in high density pin fin heat sinks which provide maximum performance when using impingement cooling.
- *Folded/Convoluted Fins:* Corrugated sheet metal is bonded to an extruded base plate. The heat sink surface area is increased due to the folds, thus the overall thermal performance improves.

5.8 System Airflow Patterns

Airflow management is critical to ensure that adequate localized airflow is provided to all components in the system. The processor, chipset, memory, and A.G.P. graphics typically require more airflow than the other peripherals or add-in cards, so care must be taken to properly distribute the airflow among components.

An important consideration in airflow management is the temperature of the air flowing over the components. Heating affects from add-in cards, memory, and peripheral devices increase the internal air temperature thus reducing the cooling efficiency of the air. The recirculation of air can also contribute to increased internal air temperatures.

For example, a system with minimal venting and a low-flow power supply fan will have restricted airflow through the system. Restricted airflow results in lower system air speeds and often, stagnant air pockets. This can directly lead to an increase in the overall internal air temperature. The warm, slow air creates less effective component cooling that requires additional cooling mechanisms such as fan heat sinks rather than passive heat sinks. The well-designed system allows less expensive, passive heat sinks to be used on the processor and possibly no heat sinks on other components.

Ducts can be designed to isolate key integrated circuit devices (such as processors and chipsets) from the effects of system heating (such as add-in cards and peripherals). Air provided by a fan or blower can be channeled directly over the key integrated circuit devices, or split into multiple paths to cool multiple integrated circuit devices. This method can also be employed to provide some level of redundancy in a system requiring redundant capabilities for fault tolerance. This is accomplished by channeling air from two or more fans through the same path across a processor. Each fan, or each set of fans, must be designed to provide sufficient cooling in the event that the other has failed.

When ducting is used, it should direct the airflow evenly from the fan across the entire component. The ducting should be accomplished, if possible, with smooth, gradual turns as this will reduce airflow impedance. Sharp turns in ducting should be avoided. Sharp turns increase friction and drag and will greatly reduce the volume of air reaching the key integrated circuit devices.

The three main factors contributing to the distribution of airflow in a system are:

- Power supply characteristics, including location, impedance, and fan size can heavily influence system airflow patterns.
- Vents in the chassis must be placed to allow in-rushing cool ambient air to cross hot components and exhaust out the power supply in an evacuating configuration. Alternatively, vents can be placed to allow hot air to escape the system in a pressurized configuration.
- If a system fan is used in addition to the power supply fan, the location and the flow direction can be important. The use of a second system fan can differ depending on the form factor. Most NLX systems use the fan to cool the processor, chipset, and memory. Pressurizing with the system fan generally cools well in these form factors though acoustics may suffer.

In an NLX system, the dual chamber chassis design makes power supply choice less critical than in some other configurations, and the core logic is close to the system fan. However, the system fan adds cost and increases acoustical noise.

NLX airflow pattern key features:

- The system is evacuated with the power supply fan and is pressurized with the front system fan (series combination). This increases the system’s total airflow by increasing the air pressure in the system (see Figure 5.9).
- The dual chamber system approach relies on the power supply less to cool the core logic and more on the front system fan.

Figure 5.8 shows the airflow pattern in a typical NLX system.

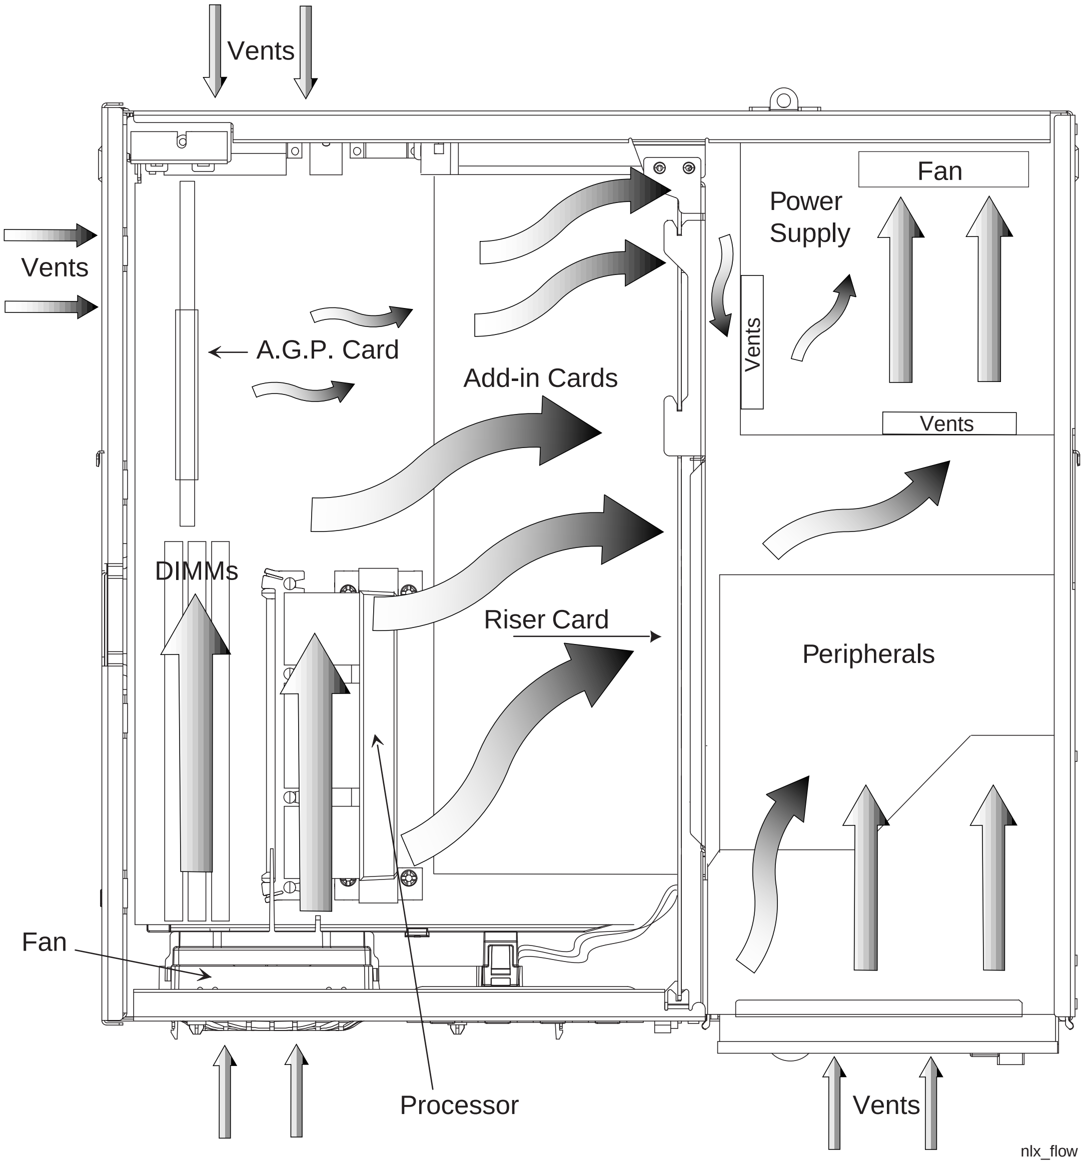

Figure 5.9: NLX Airflow Pattern

5.8.1 Chassis and Bezel Venting

Proper venting is a key element in any good thermal design. A balanced approach to vent location and pattern type is a critical factor in this design. Implementing an insufficient amount of open area to the exterior of a chassis does not allow enough air into the system for adequate cooling. Implementing too much venting can allow for air to bleed from the chassis thus decreasing the air velocity across the system's hot components. The reduction in local air velocities at these components results in less forced convective heat transfer.

To increase airflow through the system, all system accessory components (cables, wires, sheet metal, etc.) should present the lowest possible airflow impedance.

(NOTE: To eliminate possible electromagnetic compliance issues, neither the maximum vertical nor maximum horizontal dimensions of ventilation apertures, I/O ports, and open areas along chassis seams should be less than 1/20th of a wavelength of the highest harmonic frequency of interest.)

Key venting considerations:

- *Power Supply* – In the NLX system the front system fan delivers the majority of the airflow into the system. Therefore, the airflow capability of the power supply is less important. Though it is necessary to balance the volumetric airflow of both the power supply and the system fan to maximize the local air velocities at the hot components of interest.
- *Front bezel venting* – The bezel vent area should be as large as possible because it provides the main airflow source for the core logic components. **It is necessary to ensure the plastic bezel vent pattern allows air to enter freely so it does not overly restrict airflow into the system.**
- *Riser card* – Some venting at the front and back of the riser card is preferable to allow for the evacuation of the chassis and airflow through the add-in cards.
- *Side chassis venting, motherboard side* – This is required to cool the A.G.P. graphics if an active fan heat sink is not used.
- *Rear chassis venting* – This adds to the airflow capability of the chassis by allowing some release of the air delivered by the front system fan. The size and configuration of this venting should be studied to maximize the air velocities at the component level.
- *Peripheral bay venting* – Cools peripherals. Minimal venting, if any, should produce adequate results. Implementing too much venting may cause lower airflow in other areas of the chassis.

5.9 Peripheral and Add-in Card Considerations

5.9.1 Peripherals (Hard Drive, DVD, CD-RW)

The general trend for peripheral devices such as 10,000 RPM hard drives, DVD-ROM, DVD-RAM, and CD-RW is an increase in operating temperature and possibly a decrease in the required local ambient temperature necessary for cooling. General power dissipation is around 2-3 W and operating temperatures are 5 °C to 50 °C. Future generations are expected to increase maximum power dissipation and operating temperatures. The system designer must determine the thermal environment of the market segment and ensure that the peripherals are within the published thermal specifications.

5.9.2 Add-in Cards (Graphics Controllers)

The general trend for high performance graphics controllers is toward integrated and higher frequency functions. For competitive 3-D performance levels, the graphics controller power trend will increase over time. The forecasted graphics controller maximum power levels for Accelerated Graphics Port (A.G.P). introduction are 3-6 W. Future generation graphics controller power is expected to increase to 5-10 W for most controllers, with the potential of up to 15 W for some high performance versions. The system designer must determine the thermal environment of the market segment and ensure that the add-in cards are compatible within this environment.

5.10 Measurement Techniques

5.10.1 Temperature

Temperature can be measured by a diverse array of sensors. These devices all measure temperature by sensing a change in some physical parameter. The six most common types of devices are thermocouples, resistive temperature devices (RTDs and thermistors), infrared sensors, bimetallic devices, liquid expansion devices, and state changing devices. This discussion will focus on three commonly used types: resistive temperature devices, infrared sensors, and thermocouples.

5.10.1.1 RTDs and Thermistors

Resistive temperature devices (specifically thermistors) are commonly used in electronic circuits (such as power supply fan speed control circuits) to sense and control the temperature of electronic components. Resistive temperature devices take advantage of the principle that a material’s electrical resistivity varies with changes in temperature. Metallic devices are called *RTD*s, and their resistivity *increases roughly linearly with temperature*. Semiconductor devices are called *thermistors* and their resistivity *decreases nonlinearly with increasing temperature*. Thermistor devices are commonly used in temperature control circuits. Their highly nonlinear behavior poses a problem for circuit designers. Careful use of matched pairs, such that their nonlinearities offset one another, can minimize this difficulty. Thermistors are usually designated in accordance with their resistance at 25 °C; typical values are around 2 kΩ, 5 kΩ and 10 kΩ.

5.10.1.2 Infrared Sensors

Infrared sensors are non-contacting devices that can be used to provide a thermal map of a system, thus providing an indication of where to place thermocouples for system level testing. Infrared cameras are used to generate these temperature maps within the system during typical operation. These maps indicate which components should be monitored, and instrumented with thermocouples, during thermal validation and qualification testing. Several things affect the accuracy of temperature sensing via infrared devices:

- Materials radiate at various efficiencies due to differences in emissivity.
- Radiation efficiency is affected by localized oxidation, surface roughness, and other factors.
- Infrared energy may be reflected from other sources, rather than the targeted surface.
- The measured surface must completely fill the field of view of the camera.

Infrared sensing devices must account for all these issues to function accurately. To investigate thermal compliance, system components should not exceed their specified thermal limitations during testing. Infrared sensors provide an excellent qualitative indication (thermal map) of system temperatures at various components, i.e., processor, peripherals, add-in cards. If care is taken, these sensors can also provide an accurate quantitative snapshot of overall system temperatures. However, to thermally validate a system, components identified as being near their thermal limits must be instrumented with thermocouples to verify they never exceed their thermal limits under any anticipated loading or environmental conditions.

5.10.1.3 Thermocouples

Thermocouples are the most common temperature sensors used in test and development work. A thermocouple consists of two dissimilar metal wires joined as shown in Figure 5.10. The AB connection is called the junction and is attached at the desired measurement location. The opposite end is the reference end. When $T_{junction}$ is different from $T_{reference}$, a low-level voltage is generated at the terminals. This voltage is quite small and depends on the materials A and B, and $T_{reference}$ and $T_{junction}$. Thermocouples are calibrated in microvolts per degree Celsius.

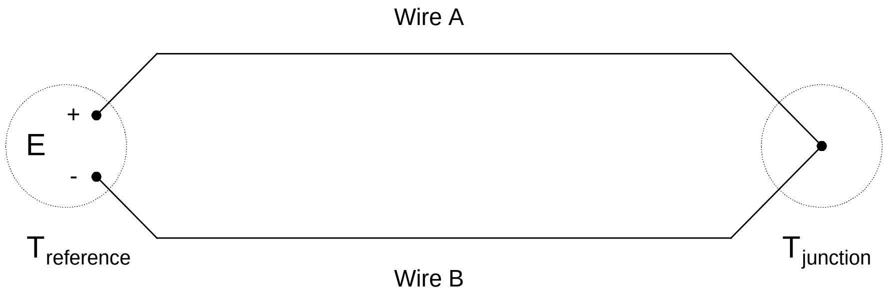

Figure 5.10: A Simple Thermocouple

5.10.1.4 Thermocouple Types

The three most common types of thermocouple used for measuring moderate temperatures are highlighted in Table 5.3.

Table 5.3: Thermocouple Types

Type	Material	Generated EMF	Temperature Range	Accuracy	Comments
J	Iron-Constantan (white-red wire)	51 µV/°C	0 °C to 750 °C	Standard: 2.2 °C or ¾% Special: 1.1°C or 0.4%	Iron wire is magnetic, susceptible to corrosion and should not get wet; not recommended for low temperature measurements.
K	Chromel-Alumel (yellow-red wire)	40 µV/°C	-200 °C to 1250 °C	Standard: 2.2 °C or ¾% Special: 1.1 °C or 0.4%	Alumel wire is magnetic, susceptible to vibration induced EMF (use strain loops), corrosion resistant.
T	Copper-Constantan (blue-red wire)	40 µV/°C	-200 °C to 350 °C	Standard: 1.1 °C or ¾% Special: 0.5 °C or 0.4%	Not magnetic; wire thermal conductivity is high making it very susceptible to conduction errors; needs large immersion depths.

Table Notes:

1 In the material column, the first named material is the positive element. The second material forms the negative wire, and is color coded red (per U.S. standards).

2 All three types are suitably linear for the ranges of temperature measured in the Intel environmental laboratory. However, Type "J" is typically used to ensure compliance with UL test requirements.

3 The accuracy column can be interpreted as the percent of the difference between $T_{junction}$ and $T_{reference}$. Type "K" thermocouples have a standard accuracy of 2% below 0 °C and Type "T" thermocouples have a standard accuracy of 1.5% below 0 °C.

4 The wires come in a variety of gauges, 30 and 36 being fairly common. The wire should be as small as is practical to prevent heat-sinking of the specimen to the outside world and to prevent disturbance of the air flow, but must be heavy enough to be reasonably durable and resist damage. The recommended wire size for general use in electronic equipment is 30 gauge.

5.10.1.5 Reference Temperature

Transition from thermocouple wires to pairs of copper wire or terminals for connection to measurement circuitry must be done in a controlled constant temperature zone. Since the signal from the thermocouple depends as much on $T_{reference}$ as on $T_{junction}$, it's important to describe this process. The most common ways to set the reference temperature include:

- *Ice baths:* stable, inexpensive and very accurate, but rather inconvenient.
- *Electronically controlled reference sources:* less accurate, more convenient, but requires periodic calibration.
- *Zone boxes:* provide uniform temperature region for connectors, need electronic compensation, most convenient.

Most thermocouple data loggers have the ability to set the reference temperature. Read the instruction manual included with the data logger to understand which method is used.

5.10.2 Airflow

It is essential to understand airflow in order to effect an appropriate thermal management solution. This includes sketching system airflow patterns, and measuring local airflow velocities, system volumetric airflow, and system pressure drops.

- *Airflow patterns:* A simple but effective way to observe system airflow patterns has been to build a clear Plexiglas cover with holes drilled at appropriate locations for inserting smoke probes into the air stream. Insert a smoke probe into the air stream, then observe and sketch the localized airflow pattern. This provides a qualitative assessment of overall system airflow behavior.
- *Localized airflow velocity:* Use a hot-wire anemometer (or similar gauge). These probes provide LFM data perpendicular to the probe head and are highly sensitive to the probe orientation in the air stream. (Note: Probe accuracy is suspect below 30 LFM for most types.)
- *Static pressures:* Insert a static pressure tube perpendicular to the airflow. If the probe is not completely perpendicular to the airflow, the readings will be biased by the dynamic pressure (velocity head) of the air stream. The dynamic and static pressures combine to give the total (or stagnation) pressure.
- *Fan characteristics:* Use an airflow chamber to verify the fan performance curve at various voltages. The chamber can be used to determine the chassis airflow impedance as well.

6. Thermal Test Methodology

6.1 Introduction

A designer must make several decisions after the preliminary chassis design is complete and the basic airflow pattern has been established. These include the location and airflow requirements for a system fan (if necessary), selection of a proper power supply, and design of the ducting to direct airflow over key components (if necessary). A designer must next classify and prioritize the design variables for feasibility, cost impact, and ease of implementation.

Once the design variables are prioritized, an experimental design tree can be constructed which outlines all design variable combinations that should be evaluated. An example design tree is shown in Figure 6.1.

This example illustrates two power supply combinations and three different fan speeds. The total number of “runs” required to evaluate this design is determined by counting all the power simulation and airflow boxes at the right edge of the figure. Not all combinations need to be evaluated to determine a robust solution. Some combinations may not be possible due to chassis design, power supply form factor, intended end use, hardware and/or software limitations, or cost implications. This example requires a total of 18 runs to completely evaluate the design.

Next, decide what components will be included in the system design evaluation such as the motherboard, processor type and speed, memory, graphics card (generally A.G.P.), additional add-in cards, and peripherals (CD-ROM, DVD, hard disk drive). To demonstrate design performance, the system should be configured with the heaviest load the system might encounter in normal use. Typically all add-in card slots are populated, multiple hard drives are installed, and the highest power processor is used. If fan speed control is implemented on either the power supply or system fan, the system should also be evaluated with combinations of light and heavy loads and high and low temperature.

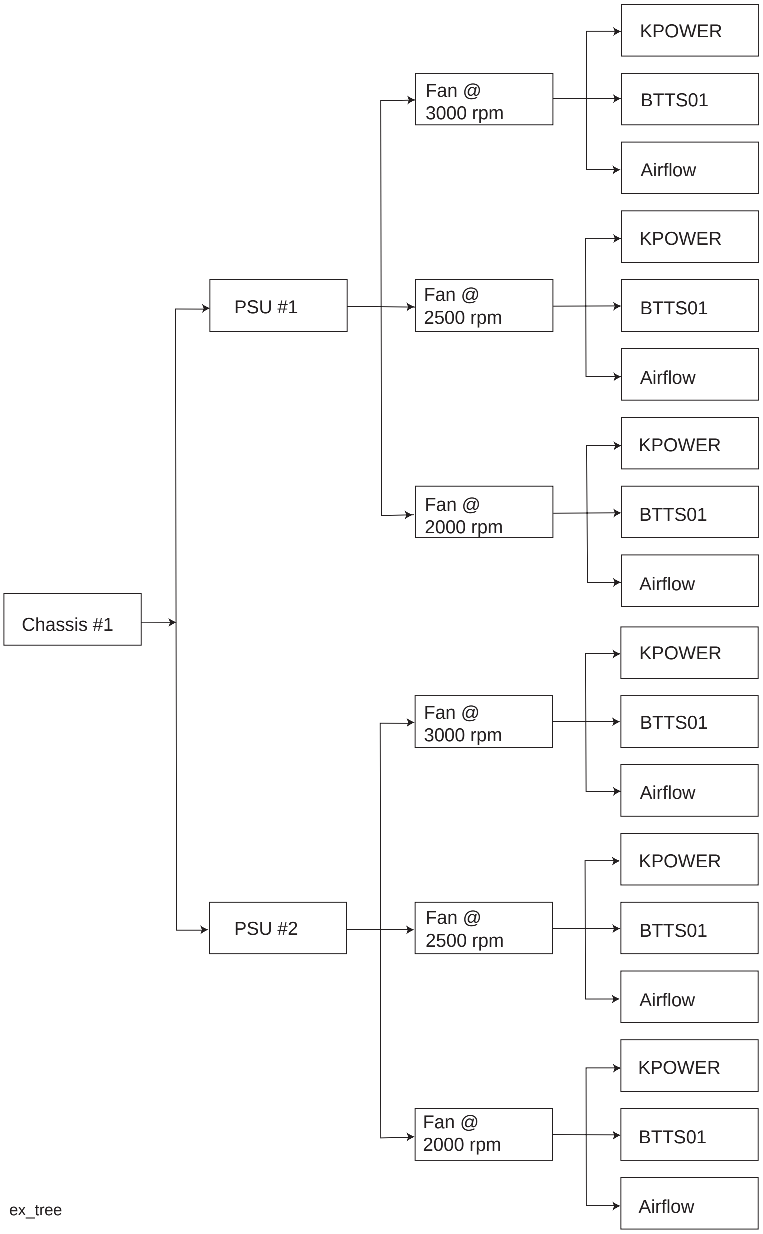

Figure 6.1: Experimental Design Tree

Different software programs tax the system components at power levels varying from the minimum to the maximum published power. Thermal design performance should be demonstrated at the maximum power dissipation of all components. Power simulation software serves this purpose. These software utilities emulate the anticipated maximum power dissipation of the processor, chipset, memory, and A.G.P. graphics accelerator. While it is impossible to emulate the maximum power dissipation of all components at the same time, simulating the maximum power of each component separately ensures unanticipated failures will not occur in the field later.

Common temperature and airflow measurement techniques should be used to collect the pertinent data for the key system components. Typical testing includes both thermocouple and hot-wire anemometer probe use to ensure key components do not exceed the published specifications for the maximum temperature.

Depending on the design, system temperatures may not stabilize for over one hour when executing power simulation software. Therefore, robust data acquisition techniques are often necessary to monitor temperature variation over time. Data acquisition equipment with the ability to monitor multiple thermocouples simultaneously and interfaced to a personal computer is highly recommended. The multiple channel capability allows the designer to monitor temperature fluctuations versus time at any location in the system. The personal computer interface provides automated data collection at predetermined time intervals to easily monitor the fluctuations and determine when temperatures have stabilized.

After data collection is complete, each design variable combination should be evaluated for thermal performance, feasibility, ease of implementation, and cost impact. Often, the lowest component temperature design has associated feasibility or cost drawbacks making the best solution a compromise among all factors.

6.2 Definitions

- **Light Load –** Basic system configuration with processor, one hard drive, and one CD-ROM drive. No add-in cards or secondary hard drives are installed.
- **Heavy Load –** Basic system configuration with all add-in card and peripheral bays populated. Some assumptions must be made for the power dissipation of the add-in cards and peripherals. Maximum add-in card and peripheral power dissipation is typically set at 10 W each.
- **Low Temperature –** Room ambient (22 °C)
- **High Temperature –** 35 °C
- **Four-corner thermal testing –** Testing at both heavy and light load, and high and low temperature. Combinations are:
  - Low temperature and light system load
  - Low temperature and heavy system load
  - High temperature and light system load
  - High temperature and heavy system load

6.3 Data Acquisition Techniques

Some type of apparatus must be used to measure the temperature of each thermocouple. If very few thermocouples need to be monitored, handheld devices and manual tracking at specified intervals can be used to determine when the temperatures stabilize. However, data collection becomes very difficult, time intensive and error prone with numerous thermocouples.

For system level evaluation where many thermocouples are required, the data acquisition technique should be automated to eliminate errors associated with manual data collection. This also makes temperature tracking and system temperature stabilization easier to determine.

Key benefits of automated data acquisition are:

- Multiple thermocouple monitoring capability. Twenty channels should be the minimum monitoring capacity.
- The unit can export collected temperatures through a serial I/O port to the computer for analysis.
- Through software, the temperature of each thermocouple over time may be displayed graphically making temperature stabilization easily identifiable.
- Data can be stored in spreadsheet format for additional data analysis and record keeping.

Many units that have these capabilities are available, such as Hewlett Packard Data Acquisition/Switch Unit Model 34970A or the Fluke Hydra Data Logger.

The steps required to collect the thermal data are listed below.

1. Determine desired system configuration: chassis, motherboard, processor, hard drives, CD-ROMs, add-in cards, and other components.
2. Determine pertinent cooling components: power supply, system fan (if used), processor active or passive heat sink, and ducting.
3. Place thermocouples on the components and key locations.
4. Boot the computer and verify proper operation.
5. Execute the power simulation software for individual components.
6. Collect temperatures with the data acquisition unit.
7. Monitor temperatures using software interface until temperatures stabilize. Depending on system design, temperatures may require up to two hours to stabilize.
8. Stop power simulation software and repeat the procedure for other power simulation software.

Airflow data acquisition is similar; however, power simulation software is not executed and a minimum ten minute stabilization period is required before data is collected.

6.4 Key Integrated Circuit Devices

6.4.1 Pentium® II Processor

Guidelines have been established for the proper techniques to be used to ensure the Intel® Pentium® II processor is within thermal specification limits. Data sheets *Intel Pentium® II Processors At 233 MHz, 266Mhz, 300 MHz, and 333 MHz* and *Intel® Celeron™ Processors At 266 MHz, 300 MHz, 300 MHz, and 333 MHz* provide information on how to accurately measure temperatures, and run the power simulation software that will emulate the anticipated maximum thermal design power.

6.4.2 Chipset

Guidelines have been established for the proper techniques to be used when measuring the chipset case temperatures. The *82443BX PCISet Application Note* provides information on how to accurately measure the case temperature, and run the power simulation software which will emulate the anticipated maximum thermal design power.

6.4.3 RIMM: RDRAM In-line Memory Modules

Typically case or junction temperatures are measured to determine component thermal compliance. The RDRAM thermal specification requires local ambient temperatures and local airflow measurements to determine thermal compliance instead of the typical junction or case temperature measurement. These local measurements are used to determine the corresponding junction temperature for RDRAM with heat spreaders dissipating 6 W, per Figure 6.2 (max spec 100 °C)[1].

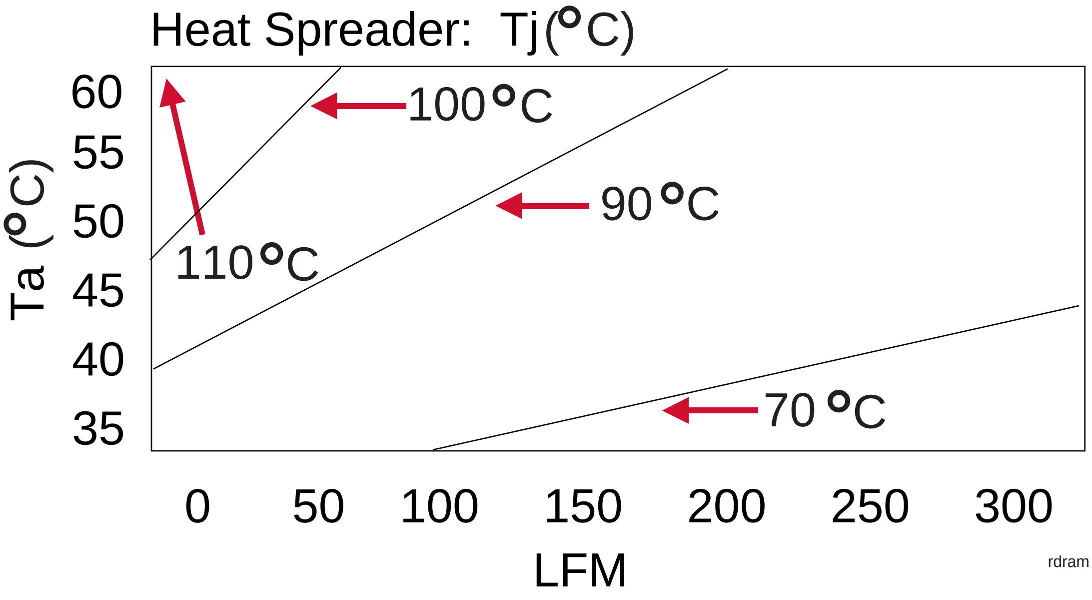

---

[1] Specification as of September 1998.

Figure 6.2: RDRAM Junction Temperature

6.5 NLX System Configuration

The NLX form factor allows motherboards ranging from 8.0 in x 10.0 in (minimum) to 9.0 in x 13.6 in (maximum) to be used. Motherboards at each end of the size spectrum are used for evaluation.

- "Short" motherboard, a 8.25 in x 10.0 in board
- "Full size" motherboard, 9.0 in x 13.6 in board

These two motherboard sizes are chosen because of the their processor, chipset, and memory location with respect to the front fan in the chassis. The Intel JN440BX motherboard locates the processor, chipset, and memory close the front fan, whereas the Intel KU440EX motherboard locates these components near the middle of the chassis making it more difficult to provide adequate airflow for cooling. Each motherboard is evaluated in an NLX chassis from Intel. **Use of Intel motherboards and chassis for evaluation does not constitute approval/disapproval of the use or design of these products.**

Thermal design performance should be demonstrated with the heaviest load the system might encounter in normal use. All available add-in card slots and peripheral bays are populated. Add-in cards vary widely in power dissipation and temperature requirements so resistive load cards designed to dissipate a user adjustable heat load are used to simulate typical add-in cards. SCSI and IDE hard disk drives, and IDE CD ROM drives populate all available drive bays. For all evaluations, the power supply fan is operated at 12 VDC. The typical heavy system configuration is listed in Table 6.1.

Table 6.1: System Configuration

Component	Type	Quantity
Chassis	Intel NLX	1
Motherboard	Intel JN440BX Intel KU440EX	1 1
Processor	Pentium II processor, 300 MHz	1
Memory	32 MB, 100 MHz SDRAM	3
Graphics – A.G.P. Add-in card	Intel740™ chip based	1
PCI	10 W Resistive Load Card	2
ISA	10 W Resistive Load Card	1
Floppy Drive	Standard	1
Primary Hard Drive	IDE	1
Secondary Hard Drives	SCSI (idle)	1
CD ROM	IDE CD ROM	1
Power Supply	PS2, Delta DPS-200PB-87A	1

6.6 Power Dissipation

The processor, chipset, memory, and graphics are the major contributors to the thermal load on the system. The power dissipation and temperature specifications for the components used in the evaluations are listed in Table 6.2.

Table 6.2: Power Dissipation

Component	**Thermal Design Power (Reference Only)**	**Temperature Spec (Reference Only)**
Pentium® II processor, 300 MHz	41.4 W	72 °C
SDRAM	6 W	Not Applicable for RDRAM Simulation
82443BX PCI/A.G.P. controller (PAC)	4 W	105 °C
Graphics (A.G.P. 2X)	6 W	109 °C (Intel740 chip)
Add-in Cards	-	T_{amb}<55 °C
Peripherals	-	T_{case}<55 °C

Note: Specifications as of September 1998.

Please refer to the Application Notes for the corresponding components to obtain the published thermal design power and maximum temperature specifications. The next generation processor, chipset, and RDRAM memory are not currently available for evaluation. Each of these components is simulated using current technology. The Pentium II processor operating at 300 MHz simulates the next generation processor, the 82443BX chipset simulates the new chipset, and SDRAM memory physically simulates the RDRAM memory.

The next generation processor is not expected to exceed the maximum power dissipation of the Pentium II processor. Future testing with the next generation processor is needed to evaluate its true thermal performance. The next generation chipset specifications are not available currently; however, the heat transfer mechanisms remain the same so the 82443BX is used for approximate power simulation. SDRAM dissipates approximately the same power as RDRAM even though the temperature specifications differ. SDRAM specifies case temperatures whereas RDRAM specifies local ambient temperatures and airflow velocities. For RDRAM simulation purposes, the local ambient temperatures and airflow velocities are measured instead of the SDRAM case temperature.

Add-in cards and peripherals vary significantly in power dissipation. The manufacturer's specifications should be consulted for the thermal design power and maximum temperature specifications.

6.7 Power Simulation Software

Power simulation software includes utilities specifically written to test the thermal design power for Pentium II processor, 82443BX chipset, and Intel740 chip based A.G.P. graphics cards. This software is used to monitor the thermal performance under "worst-case" system component conditions. Future applications may exceed the thermal design power limit for transient or all time periods. **Power simulation software is intended to simulate "worst-case" conditions but does not guarantee thermal compliance for future software applications.**

The processor is exercised using the utility KPOWER.exe. The exact power dissipation of the processor while executing KPOWER.exe is unknown but estimated at 80% of the published processor power.

The chipset is exercised using the utility BTTS01.exe. The power dissipated with BTTS01.exe depends on which switches are activated. The "/u2" option is most commonly used and will dissipate between 6.5 W and 7 W depending on the number of DIMM slots populated.

The Intel740 A.G.P. graphics accelerator maximum thermal design power was simulated using the software utility THERM740.exe. The exact power dissipation is unknown but is approximately 6 W.

Power simulation software for future processors, chipsets, and graphics accelerators may change. Ensure the correct simulation software is used for the component in question.

Please contact your local Intel representative for more information about the power simulation programs.

6.8 Thermocouple Placement and Type

Key components and additional thermal points of interest are monitored for each desired cooling configuration in the NLX chassis. Type J, 30 gauge thermocouples are used at these locations. Please refer to the *Intel Pentium® II Processors At 233 MHz, 266Mhz, 300 MHz, and 333 MHz* data sheet for thermocouple attachment procedures. Thermocouple locations are listed in Table 6.3 and shown in Figure 6.3.

Table 6.3: Thermocouple Locations

Name	Location	Number
Processor	Processor plate temperature, center of thermal plate	T1
82443BX PAC	Chip case temperature, center of chip	T2
Memory Ambient – Front	0.1-0.2” above top of memory toward front of chassis	T3
Memory Ambient – Rear	0.1-0.2” above top of memory toward rear of chassis	T4
Memory Exit – Front	0.1” above motherboard toward front of chassis	T5
Memory Exit – Rear	0.1” above motherboard toward rear of chassis	T6
Add-In Card Ambient – Rear	0.75”-1” above add-in cards towards back of chassis	T7
Add-In Card Internal Ambient	½ way from riser card centered between top and middle add-in cards	T8
Add-In Card Ambient – Middle	0.75”-1” above add-in cards towards middle of chassis	T9
Add-In Card Ambient – Front	0.75”-1” above add-in cards towards front of chassis	T10
Hard Drive Case	On top of primary hard drive case	T11
CD-ROM	On top of CD-ROM case	T12
External Power Supply	0.5”-1” in front of power supply air outlet on outside of chassis	T13
Internal Power Supply	0.25-0.5” in front of internal power supply air inlet inside chassis	T14
A.G.P. Controller	Controller case temperature, center of chip	T15
Processor Local	0.25” above center of heat sink	T16
82443BX PAC Local	0.25” above center of controller case	T17
Ambient	External and away from chassis	T18

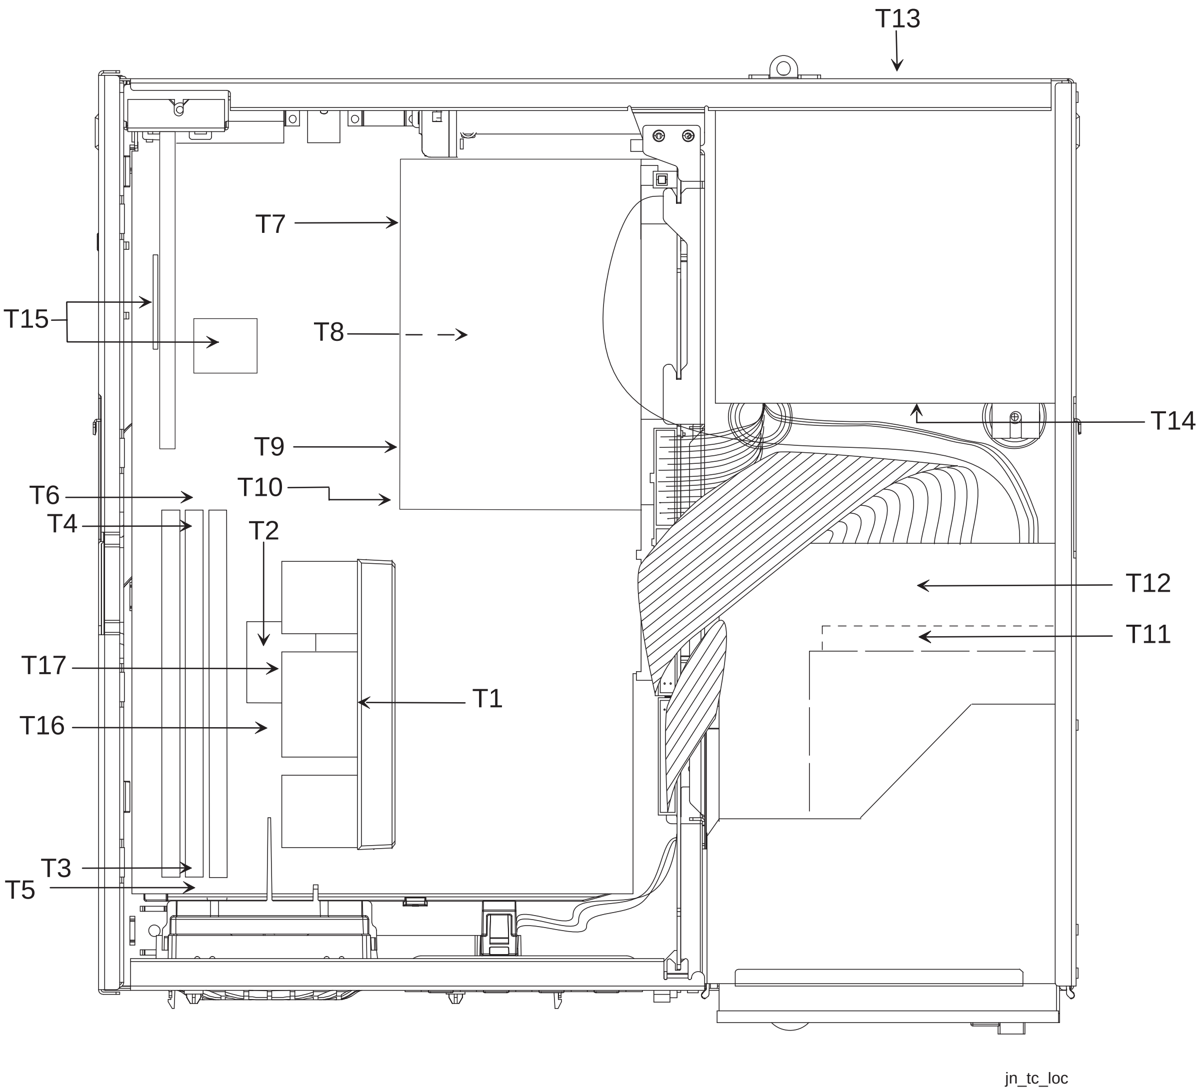

Figure 6.3: Thermocouple Locations in NLX System

6.9 Airflow Sensor Placement and Type

Key airflow pathways of interest are monitored for each desired cooling configuration in the NLX chassis. Cambridge Accusense flow sensors (hot wire anemometer) are used at these locations. Probe locations are listed in Table 6.4 and shown in Figure 6.4.

Table 6.4: Airflow Sensor Locations

Name	**Location**	**Flow Velocity Direction**	**Number**
Processor – HFHS[1]	Top of heat sink, front of chassis - HFHS configuration	Z	A1
Processor – HFHS	Top of heat sink, rear of chassis - HFHS configuration	Z	A2
Processor – HFHS	Bottom of heat sink, rear of chassis - HFHS configuration	Z	A3
Processor – HFHS	Bottom of heat sink, front of chassis - HFHS configuration	Z	A4
Processor - NLX	Top of heat sink, front of chassis - NLX configuration	Y	A1
Processor - NLX	Top of heat sink, rear of chassis - NLX configuration	Y	A2
Processor - NLX	Bottom of heat sink, front of chassis - NLX configuration	Y	A3
Processor - NLX	Bottom of heat sink, rear of chassis - NLX configuration	Y	A4
Memory	In front of DIMM 0, parallel to memory, in center of memory card	Y	A5
Memory	In front of DIMM 0, in center of memory card	Z	A6
Memory	In back of DIMM 2, in center of memory card	Z	A7
Memory	In back of DIMM 2, in center of memory card	Y	A8
Memory	Between DIMM 0 and DIMM 1 towards front of chassis as close to motherboard as possible	Y	A9
Memory	Between DIMM 0 and DIMM 1 towards rear of chassis as close to motherboard as possible	Y	A10
82443BX PAC	0.2” above chip	Z	A11
82443BX PAC	0.2” above chip	Y	A12
Power Supply	In back of external power supply vent	Y	A13
Power Supply	In front of internal power supply vent	Y	A14
Front Fan Vent	External side of vent	Y	A15
Rear Vent	External side of vent	Y	A16
Side Vent (A.G.P.)	External side of vent	X	A17
A.G.P.	Below A.G.P. processor chip	Z	A18

Notes:

1 HFHS – Horizontal Fan Heat Sink. Concept introduced in Section 7.

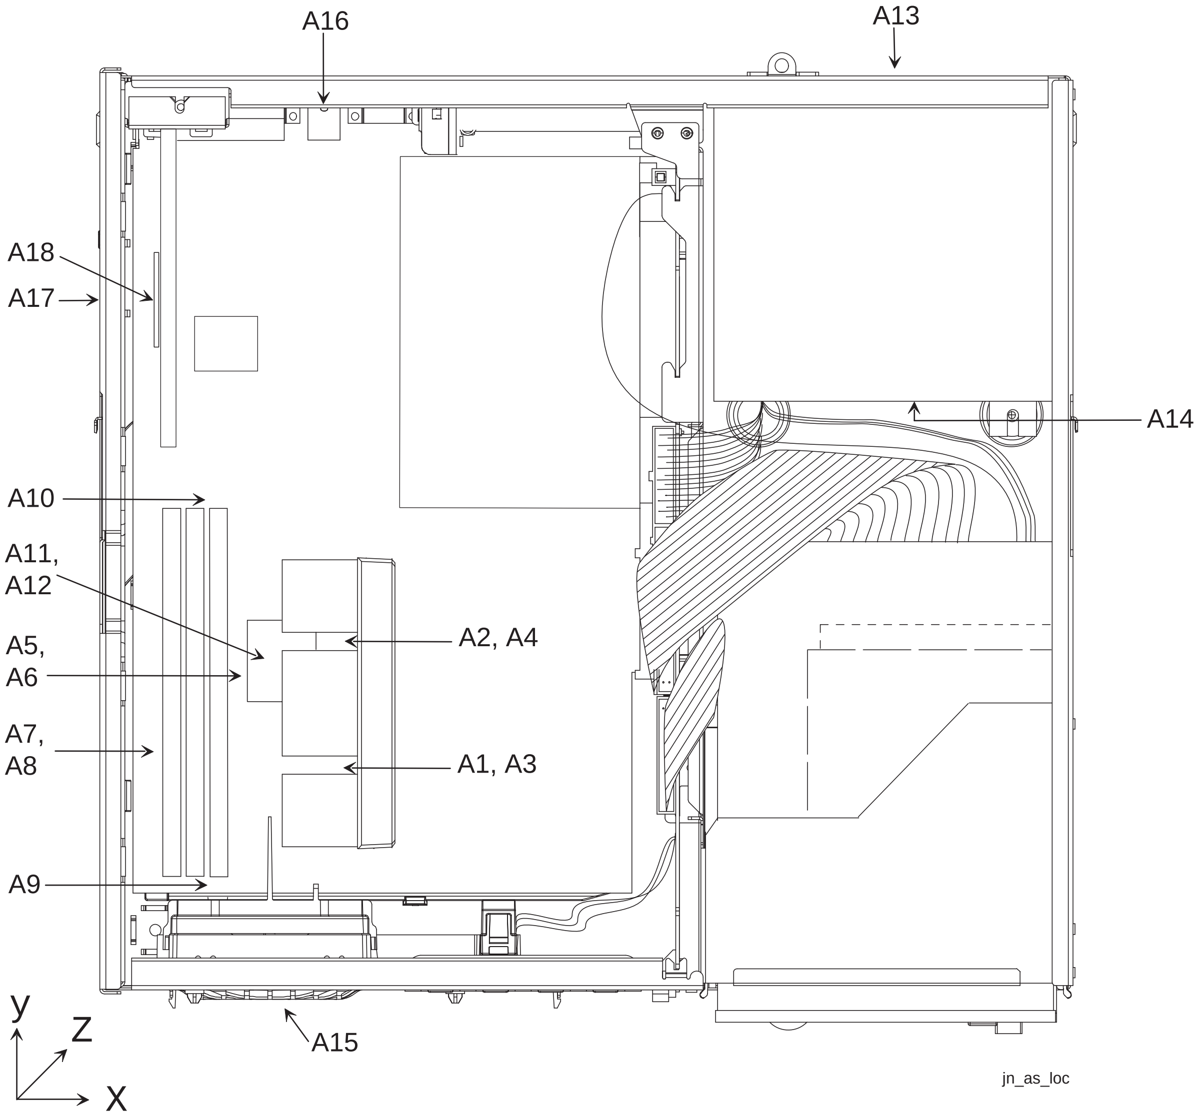

Figure 6.4: Airflow Sensor Locations in NLX System

7. System Design Examples

7.1 Introduction

The example NLX chassis (Figure 7.1) utilizes the power supply to evacuate the chassis and move air from the front to the back of the chassis. Figure 5.9 illustrates the general airflow pattern. An air intake vent in the left front panel of the chassis is designed for a second 80 mm system fan to be mounted for increased airflow across the core logic components. This vent also serves as the main system air inlet. The bottom of the front bezel incorporates small vent holes for air intake.

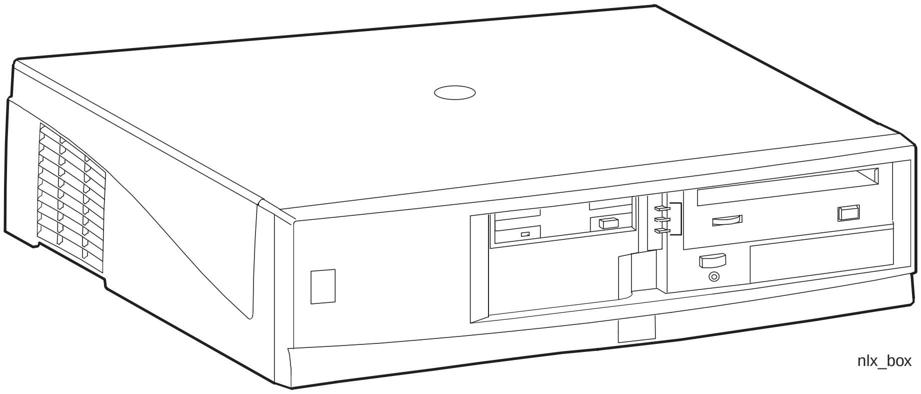

Figure 7.1: Example NLX Chassis

Based on the NLX chassis flow pattern described above, the design process started with a brainstorming session for ideas to cool the processor, chipset, memory, and A.G.P. graphics card. Considering feasibility, ease of implementation, and cost the ideas are reduced to five possibilities.

- Place an active fan heat sink on the processor.
- Implement the Horizontal Fan Heat Sink (HFHS) concept and remove the front system fan.
- Relocate the front system fan internally near the processor and possibly
  - Add a duct from front of chassis to the fan, and/or
  - Add a duct covering the relocated fan and the processor, 82443BX PAC, and memory.
- Increase the front bezel vent size.
- Replace the standard 80 mm front system fan operating at 2000 rpm with a higher flow fan.

7.2 Processor Active Fan Heat Sink

The active fan heat sink on the processor is considered because of the need for a “boxed product” solution that can be purchased separately. The active fan heat sink is a chassis independent solution allowing it to be implemented across all form factors. However, the active fan heat sink (Figure 7.2) is intended to cool only the processor and does not cool the chipset, memory, or A.G.P. graphics card. If the airflow patterns established by the power supply and the chassis layout are adequate to cool the other key components, the active fan heat sink for the processor may be feasible thermally.

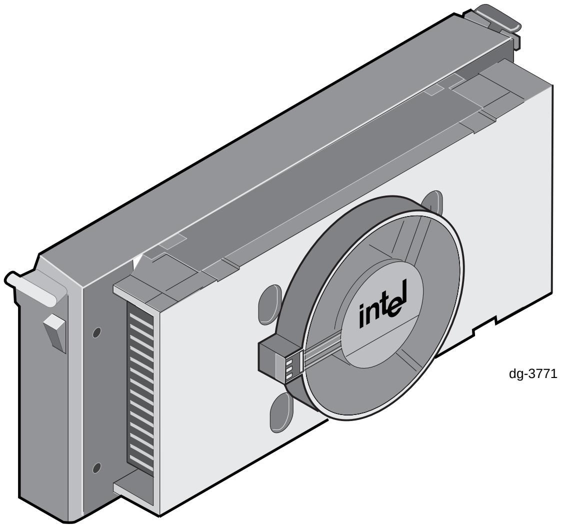

Figure 7.2: Active Fan Heat Sink

Numerous processor active fan heat sinks are available commercially. The unit used for this NLX evaluation is the Sanyo Denki model 109X1512H3036 rated at 12 VDC at 0.06 A. The fan on the unit averages 4200 rpm under normal operating conditions.

In addition to the fan heat sink on the processor, the standard 80 mm fan operating at 2000 rpm located in the front of the chassis was installed for all configurations (full size and short motherboards).

A 2.25 in x 2.25 in front bezel vent is also added to decrease the airflow restriction to the front fan. No other chassis modifications were performed.

7.3 Horizontal Fan Heat Sink (HFHS)

The HFHS (Figure 7.3) is a standardized fan heat sink technology designed for ATX, microATX, and NLX systems. The cooling concept is similar to the active fan heat sink but uses an 80 mm fan to increase the airflow to the processor, adjacent chipset component, and memory whereas the active fan heat sink increases airflow to the processor only.

Figure 7.3: Horizontal Fan Heat Sink

Key advantages of the Horizontal Fan Heat Sink include:

- The HFHS is chassis independent and therefore one design usually fits for all form factors (ATX, microATX, NLX).
- The 80 mm fan is located directly above the processor, adjacent chipset component, and memory, increasing the cooling benefit of this system fan.

Key disadvantages:

- The fan is cantilevered off the processor, adding mass, which can affect shock and vibration responses for the processor package.
- A custom processor heat sink is required for this approach because the fan is inset into the heat sink.

All NLX chassis evaluations using the HFHS are performed with a prototype fan bracket mounted to an SECC style processor, and an ATX heat sink modified to accept a 3200 rpm, 80 mm fan installed in the HFHS bracket. Two system fans unnecessarily add cost; therefore the front system fan was removed. No chassis modifications were made other than adding a 2.25 in x 2.25 in vent in the front bezel.

HFHS is a conceptual design and limited prototypes have been manufactured for thermal evaluation purposes only. Critical design issues concerning the bracket such as processor and/or motherboard attachment points must be addressed if an HFHS is adopted. Processor packaging style such as SECC vs. socketed must also be considered. If this design or a similar design is adopted, evaluation by the system designer is highly recommended to demonstrate both thermal and mechanical design performance.

When properly implemented HFHS is capable of effectively cooling all key system components, add-in cards, and peripherals. However, fan speed control or a smaller HFHS fan could possibly be implemented if proven thermally viable with further investigation.

7.4 Internal System Fan

The short motherboard poses additional cooling issues over the full size motherboard because the core logic components are located in the middle of the chassis away from the front system fan. Relocating the front system fan to the edge of the motherboard increases the airflow across the core logic components and also reduces front acoustical noise.

Relocating the fan internally increases the general local internal air velocity but not necessarily airflow on or through the core logic components. Properly designed ducting will direct the airflow over the components. This may reduce the core logic component temperatures. Therefore, “front” and “rear” duct designs are also evaluated.

Designing a duct that attaches to the front wall of the chassis and extends to the relocated fan (front duct) introduces cool external air to the core logic components rather than recirculating the warm internal air. However, the duct may also restrict the fan’s airflow, likely increasing component temperatures. Likewise, a duct covering the fan and the core logic components (rear duct) directs the air providing a focused airflow pattern. Even though this approach uses internal and not external air, the increased local air velocities may allow for sufficient convective cooling.

The following four combinations of the relocated internal fan and the front and rear ducts are evaluated.

- Internal fan without the front or rear ducts
- Internal fan with the front duct
- Internal fan with the rear duct
- Internal fan with both the front and rear ducts

These ideas are conceptual designs only. For thermal evaluation purposes, the fan is mounted directly to the chassis floor using double-sided tape and the front and rear ducts are formed from aluminum sheet metal. The front duct conforms to the front wall of the chassis and to the inlet of the relocated fan on the top and the sides. The chassis floor serves as the bottom of the duct. The rear duct is L-shaped and extends the length of the processor. This configuration is shown in Figure 7.4. The goal is to direct the flow across the processor and through the memory rather than allowing it to disperse freely through the chassis.

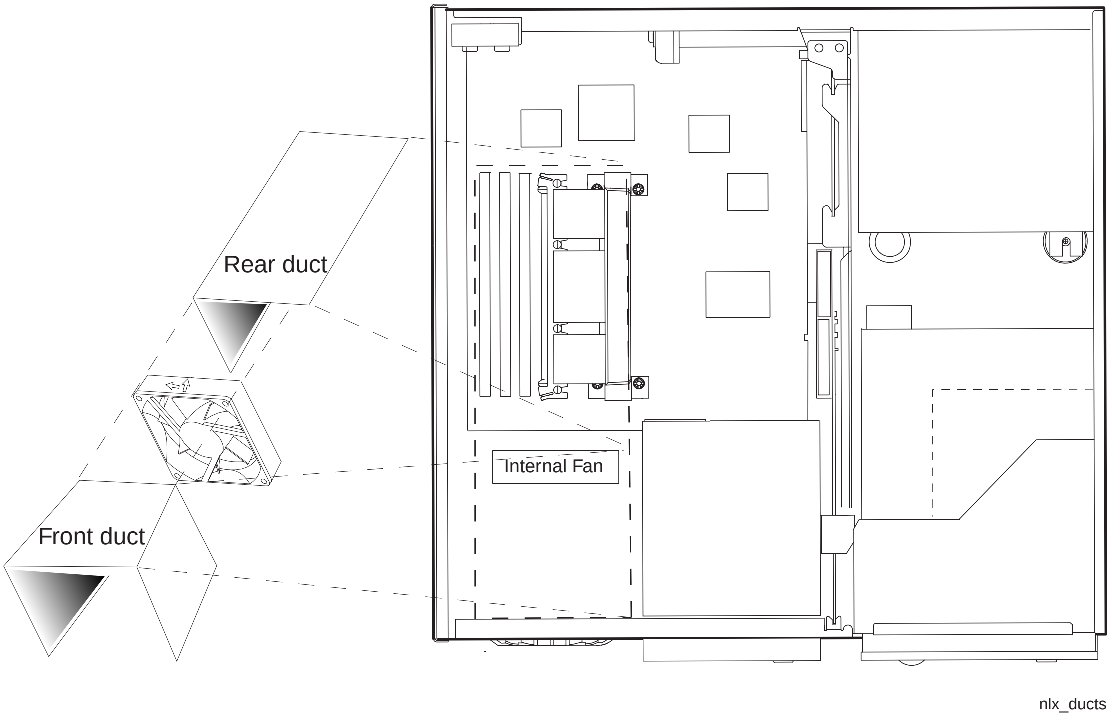

Figure 7.4: Ducted Test Configuration

For a complete, manufacturable design, the following issues must be addressed:

- Design a plastic bracket to mount the fan to the chassis floor using the motherboard mounting holes.
- Conduct a design optimization before the front or rear ducts (if implemented) are tooled for production use.

7.5 Bezel Venting

Many chassis designs place vents in the sheet metal in the proper location and with adequate size but then restrict airflow to the vent by using a bezel with little or no venting. The example NLX chassis experiences this same design difficulty. The standard bezel has three small rectangular openings in the bottom to allow air to enter the front system fan. These small holes can significantly restrict airflow to the front fan.

Airflow to the vent is improved by adding a 2.25 in x 2.25 in vent in the bezel, as shown in Figure 7.5. Understand the cost implications associated with increasing the vent size. Tooling changes can be very expensive. Increasing the bezel's front vent size may adversely affect acoustical noise and therefore should be evaluated.

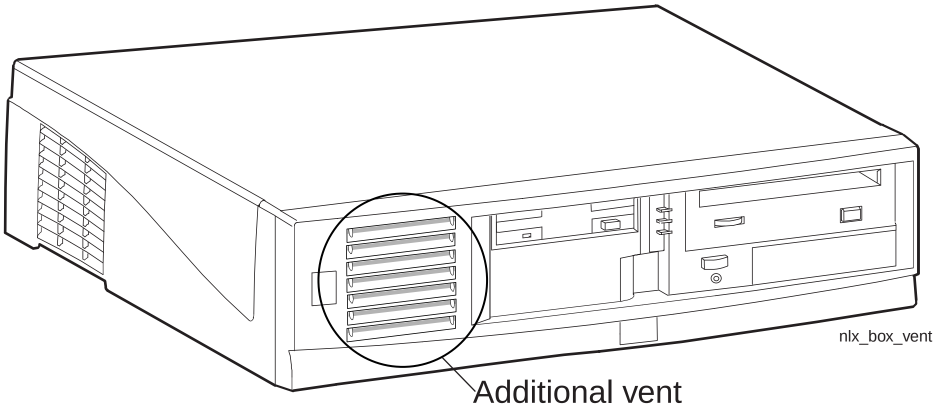

Figure 7.5: NLX Chassis with Enhanced Front Bezel Venting

7.6 System Fan Size

Currently, the Pentium II processor running at 300 MHz on the JN440BX motherboard barely satisfies the maximum specified thermal plate temperature of 72 °C in the NLX chassis with a 2000 rpm, 80 mm front system fan. Installing a higher flow fan can provide thermal margin for the processor and help cool the chipset, memory, and A.G.P. graphics accelerator. Two larger 80 mm fan flow rates, at 2750 rpm and 3200 rpm, were evaluated for cooling performance. This solution is easily implemented but can increase the system acoustical level because of the larger fan flow rates.

7.7 Results

The overall system thermal design must cool five critical components below the maximum specified temperature listed in Table 6.2.

- Processor thermal plate temperature
- Case temperature of chipset device that requires cooling
- RDRAM local ambient temperature and airflow velocity
- A.G.P. controller case temperature
- Add-in card local ambient temperature and drive bay temperature

The cooling methods mentioned above are evaluated with both motherboards. The full size motherboard locates the core logic components at the front of the chassis so only the active fan heat sink, HFHS, and larger front fan are evaluated. However, the short motherboard locates the core logic components in the middle of the chassis, requiring the relocated fan to be evaluated in addition to the other methods.

7.7.1 Full Size Motherboard

Since the full size motherboard locates the core logic components near the front system fan, this motherboard is generally more easily implemented in the cooling design. The fan provides high flow velocities over these components, thus requiring fewer changes to ensure components do not exceed their specified maximum temperatures.

For comparison purposes only, the standard unmodified example NLX chassis with the full size motherboard is evaluated thermally. The standard configuration includes a passive heat sink on the processor and an 80 mm fan operating at 2000 rpm. Neither the 82443BX device nor the A.G.P. controller have heat sinks attached. Table 7.1 compares the standard configuration with the cooling methods proposed.

In Table 7.1, all temperatures are extrapolated to 35 °C ambient air temperature; all temperatures are in °C; all velocities are in linear feet/minutes (lfm); and all sound pressure is in dBA.

Table 7.1: Example Chassis with Full Size Motherboard

Cooling Method	Front Chassis Fan Speed	Bezel[1]	CPU Tplate[2]	82443BX PAC Tcase[3]	Intel740 chip Tcase[4]	Memory Ambient[5]	Mem. Average Air Vel.	Sound Pressure[6]
Front Sys Fan	2000	Standard	74.7	105.2	99.5	44.0	**Not Measured**	N/A
Front Sys Fan	2000	Modified	66.7	98.8	99.6	37.7	381	0.4
Front Sys Fan	2750	Modified	63.0	93.5	95.8	37.8	528	2.8
Front Sys Fan	3200	Modified	59.9	89.8	93.3	37.0	620	5.5
Fan Heat Sink	2000	Modified	69.4	100.3	100.2	38.0	376	0.7
HFHS (3200 rpm)	N/A	Standard	66.4	90.4	91.2	48.8	Not Measured	Not Measured
HFHS (3200 rpm)	N/A	Modified	66.2	92.0	90.9	47.0	490	3.0

Notes: August 1998

1 Modified bezel has additional 2.25 in x 2.25 in vent.
2 CPU plate temperature running KPOWER.exe
3 82443BX case temperature running BTTS01.exe /u2
4 Intel740 chip case temperature running Therm740.exe.
5 Memory ambient taken from maximum of KPOWER, or BTTS01 /u2.
6 Delta from the standard configuration.

Each cooling method evaluated passed all component temperature specifications. However, the unmodified design marginally exceeds the processor and chipset thermal specifications. The add-in cards and peripherals passed in all configurations with maximum temperatures of 50 °C and 49 °C, respectively.

Increasing the vent size in the front bezel made the most dramatic change in key component temperatures. As shown, the processor thermal plate temperature decreased 8 °C and the 82443BX PAC and memory temperatures decreased approximately 6 °C. Acoustical noise does not increase with the additional venting indicating the vent pattern is well designed.

This confirms the concept that proper bezel venting is required to not restrict flow into the chassis.

HFHS and the active processor fan heat sink are shown to be viable cooling methods for the example NLX chassis with the full size motherboard. HFHS cools the processor, 82443BX PAC, and A.G.P. controller well below their maximum specified temperatures but the memory ambient temperature does suffer. However, increasing the bezel venting does not further reduce the component temperatures. Considering the increased cost, acoustical noise, and mechanical shock and vibration issues HFHS must overcome, HFHS may not be the best cooling method for NLX.

As mentioned earlier, the processor active fan heat sink cools the processor well but is not intended to help cool other key components such as the 82443BX PAC, memory, or A.G.P. controller. The system design must ensure these components are cooled. The example NLX chassis design and the full size motherboard locate the core logic components near the front fan, which provides high airflow velocities across these hot components. This design cools the 82443BX PAC, memory, and A.G.P. controller, allowing the use of the processor active fan heat sink. The active fan heat sink is more expensive than the passive heat sink but does not significantly increase acoustical noise (0.7 dBA increase). Both methods require the front fan to cool the key components and therefore the passive heat sink is the best choice when the full size motherboard is used.

The front fan speed greatly affects how well the components are cooled. Increasing the fan speed from 2000 rpm to 2750 rpm reduces the core component temperatures an average of 4 °C. Further increasing the speed to 3200 rpm nets an additional 3 °C. However, when the fan speed increases from 2000 rpm to 3200 rpm, the acoustical noise increases 5.5 dBA. Based on the margin the redesigned front bezel provides, the 2000 rpm fan should be retained to minimize the acoustical noise impact.

RDRAM requires local ambient temperatures and local airflow measurements to determine thermal compliance instead of the typical junction or case temperature measurement as mentioned in Section 6.3.3. The NLX form factor is well engineered to cool RDRAM because the front system fan delivers high velocity airflow parallel to the memory module(s). As illustrated in the table, the minimum local airflow provided is approximately 380 lfm and the maximum ambient temperature is 47 °C. Referring to Figure 6.2, the RDRAM junction temperature should be well below the maximum specification.

The cooling methods outlined satisfy the maximum temperature specifications published for each component. Each method has positive and negative aspects and must be evaluated for implementation. Increasing the front bezel vent size has the greatest impact on component temperatures. However, if the bezel tooling is already designed, the change can be very expensive. The active processor fan heat sink is an excellent alternative because no chassis modifications are necessary and is easily implemented but does not cool the system as well as the other methods. HFHS poses the most difficult obstacles to overcome because of the processor mounting issues but the benefits are significant. The designer must balance all the factors presented and decide which method is best for the application.

7.7.2 Short Motherboard

The short motherboard poses additional cooling challenges because the core logic components are located near the middle of the chassis. Therefore, relocating the front system fan is considered as a cooling method along with the active processor fan heat sink and HFHS.

The short motherboard uses the 82443EX chipset instead of the 82443BX chipset. While the 82443EX does not dissipate as much power as the 82443BX, it does offer a comparison between cooling methods evaluated.

Once again for comparison purposes only, the standard unmodified example NLX chassis with the short motherboard is evaluated thermally. The standard configuration includes a passive heat sink on the processor and an 80 mm fan operating at 2000 rpm. Neither the 83443EX PAC nor the A.G.P. device have heat sinks attached. Table 7.2 compares the standard configuration with the cooling methods proposed.

Table 7.2: Example Chassis with Short Motherboard

Cooling Method	Fan Speed	**Bezel**[1]	**CPU Tplate**[2]	**82443EX PAC Tcase**[3]	**Intel740 chip Tcase**[4]	**Memory Ambient**[5]	**Mem. Average Air Vel.**	**Sound Pressure**[6]
Front Sys Fan	2000	Standard	95.0	98.8	65.4	44.4	Not Measured	N/A
Fan Heat Sink	2000	Modified	66.3	78.3	70.1	38.4	141	1.1
HFHS (3200 rpm)	N/A	Modified	60.5	74.4	64.0	44.0	506	1.8
Relocated Fan No Ducting	3200	Modified	65.5	73.9	64.9	43.5	296	2.1
Relocated Fan Front Duct	3200	Modified	64.6	77.0	65.4	41.8	212	3.5
Relocated Fan Front & Rear Duct	3200	Modified	61.7	71.5	62.9	38.4	450	4.4
Relocated Fan Front & Rear Duct	2750	Modified	64.1	73.5	64.9	38.6	374	2.1
Relocated Fan Rear Duct	3200	Modified	62.7	72.6	62.7	41.6	448	2.3
Relocated Fan Rear Duct	2750	Modified	64.0	73.1	62.9	40.9	371	1.0
Relocated Fan Rear Duct	2000	Modified	68.4	76.5	65.9	40.7	274	0.3

Notes: August 1998

1 Modified bezel has additional 2.25 in x 2.25 in vent.

2 CPU plate temperature running KPOWER.exe

3 82443EX case temperature running BTTS01.exe /u2

4 Intel740 chip case temperature running Therm740.exe.

5 Memory ambient taken from maximum of KPOWER, or BTTS01 /u2.

6 Delta from the standard configuration.

Clearly the standard unmodified thermal design will not cool the system within specification because the processor exceeds the maximum thermal plate temperature by 23 °C. Therefore enhanced cooling methods must be used. The first enhancement implemented is the modified bezel. Since the increased bezel venting proved to be very effective with the full size motherboard, the bezel is immediately implemented with the short motherboard.

The active processor fan heat sink, in addition to the increased front bezel venting, reduces the thermal plate temperature almost 29 °C, indicating the capability of the fan heat sink. The chipset (i.e., PAC) temperature and memory temperatures also decrease 20 °C and 10 °C, respectively. Acoustical noise also increases 1.1 dBA when the fan heat sink is added. The active processor fan heat sink provides adequate margin for all components and is easily implemented because of the commercial availability and relatively minor cost impact.

The HFHS cools the system the best of all configurations, including relocating the front system fan. The memory ambient temperature does suffer compared to the active processor fan heat sink and the relocated front system fan. However, the airflow velocity is more than required to cool RDRAM. Surprisingly, the acoustical noise impact is relatively small at 1.8 dBA, considering that the fan is operating at 3200 rpm. Remember, HFHS has shock and vibration issues to consider if implemented along with the increased cost penalty.

Relocating the front system fan as mentioned earlier increases the airflow velocity across the core logic components and moves the acoustical noise the fan generates to the center of the chassis. The relocation reduces the component temperatures below the maximum temperature specifications. The addition of the front and rear ducts is evaluated to determine how low the core logic component temperatures can be reduced. Additionally, the fan speed is varied from 3200 rpm down to 2000 rpm to improve acoustical noise, since the 3200 rpm fan provides more than adequate thermal margin and increases acoustical noise as much as 4.4 dBA. When the fan speed is reduced to 2000 rpm the acoustical impact is a negligible 0.3 dBA increase.

The proposed cooling methods described satisfy the maximum temperature specifications published for the key components and each has its positive and negative properties that must be evaluated by the designer. The active processor fan heat sink is the easiest to implement and the lowest cost but does not cool the system as well as the other methods. HFHS cools very well but has design issues to overcome. Relocating the front system fan increases the airflow velocity and reduces acoustical noise but at a minimum a new bracket must be designed to mount the fan. Additional ducting further reduces component temperatures and, if implemented, increases the design cost. The designer must decide which parameter is most important but still consider the other parameters for the system design.