AP-830

APPLICATION NOTE

Pentium® III Xeon™ Processor/ Intel® 450NX PCIset AGTL+ Layout Guidelines

March 1999

Order Number: 245097-001

CONTENTS

FIGURES

TABLES

1.0. INTRODUCTION

The Pentium® III Xeon™ processor is a follow-on to the Pentium® Pro and Pentium® II Xeon™ processors. Like the Pentium II Xeon processor, the design of the external Pentium III Xeon processor bus enables the Pentium III Xeon processor to be "multiprocessor ready." To relax timing constraints on a bus that supports up to six loads, the Pentium III Xeon processor implements a synchronous, latched bus protocol that allows a full clock cycle for signal transmission and a full clock cycle for signal interpretation and generation. This protocol simplifies interconnect timing requirements and supports 100 MHz system designs using conventional interconnect technology. The Pentium III Xeon processor bus uses low-voltage-swing *AGTL+* I/O buffers, making high frequency signal communication between many loads easier.

The goal of this layout guideline is to provide the system designer with the information needed for the Pentium III Xeon processor and Intel® 450NX PCIset AGTL+ bus portion of PCB layout. The topology and information presented in this application note also apply to Pentium II Xeon processor/Intel 450NX PCIset system bus design. This document provides guidelines and methodologies that are to be used with good engineering practices. See the *Pentium® III Xeon™ Processor at 500 and 550 MHz* datasheet and *Intel® 450NX PCIset* for component specific electrical details. Intel strongly recommends running analog simulations using the available I/O buffer models together with layout information extracted from your specific design.

2.0. ABOUT THIS DOCUMENT

2.1. Document Organization

This section defines terms used in the document. Section 3.0. discusses specific system guidelines. This is a step-by-step methodology that Intel has successfully used to design Pentium III Xeon processor systems using the Intel 450NX PCIset components. Section 4.0. introduces the theories that are applicable to this layout guideline. Section 5.0. contains more details and insights. The items in Section 5.0. expand on some of the rationale for the recommendations in the step-by-step methodology. This section also includes equations that may be used for reference.

The actual guidelines start on Section 3.0., "AGTL+ Design Guideline."

2.2. References

- *Intel® 450NX* PCIset (Order Number 243771-001)
- *Pentium® II Xeon™ Processor Bus Terminator Design Guidelines* (Order Number 243774-001)
- *Pentium® II Processor Developer's Manual* (Order Number 243341)
- *Pentium® III Xeon™ Processor at 500 and 550 MHz* (Order Number 245094-001)
- *Pentium® III Xeon™ Processor Power Distribution Guidelines* (Order Number 245095-001)
- *VRM 8.3 DC-DC Converter Design Guidelines* (Order Number 243870-002)

2.3. Definition of Terms

Aggressor — a network that transmits a coupled signal to another network is called the aggressor network.

AGTL+ — The Pentium III Xeon processor system bus uses a bus technology called AGTL+, or Assisted Gunning Transceiver Logic. AGTL+ buffers are open-drain and require pull-up resistors for providing the high logic level and termination. The Pentium III Xeon processor AGTL+ output buffers differ from GTL+ buffers with the addition of an active pMOS pull-up transistor to "assist" the pull-up resistors during the first clock of a low-to-high voltage transition. Additionally, the Pentium III Xeon processor Single Edge Connector (S.E.C.) cartridge contains internal 150 Ω pull-up resistors to provide termination at each bus load.

Bus Agent — a component or group of components that, when combined, represent a single load on the AGTL+ bus.

Corner — describes how a component performs when all parameters that could impact performance are adjusted to have the same impact on performance. Examples of these parameters include variations in manufacturing process, operating temperature, and operating voltage. The results in performance of an electronic component that may change as a result of (including, but not limited to): clock to output time, output driver edge rate, output drive current, and input drive current. Discussion of the "slow" corner would mean having a component operating at its slowest, weakest drive strength performance. Similar discussion of the "fast" corner would mean having a component operating at its fastest, strongest drive strength performance. Operation or simulation of a component at its slow corner and fast corner is expected to bound the

extremes between slowest, weakest performance and fastest, strongest performance.

Crosstalk — the reception on a victim network of a signal imposed by aggressor network(s) through inductive and capacitive coupling between the networks.

- **Backward Crosstalk** — coupling which creates a signal in a victim network that travels in the opposite direction as the aggressor's signal.
- **Forward Crosstalk** — coupling which creates a signal in a victim network that travels in the same direction as the aggressor's signal.
- **Even Mode Crosstalk** — coupling from multiple aggressors when all the aggressors switch in the same direction that the victim is switching.
- **Odd Mode Crosstalk** — coupling from multiple aggressors when all the aggressors switch in the opposite direction that the victim is switching.

Edge Finger — The cartridge electrical contact which interfaces to the SC330 connector.

Flight Time — is a term in the timing equation that includes the signal propagation delay, any effects the system has on the T_{CO} of the driver, plus any adjustments to the signal at the receiver needed to guarantee the setup time of the receiver.

More precisely, *flight time* is defined to be:

- The time difference between a signal at the **input** pin of a receiving agent crossing V_{REF} (adjusted to meet the receiver manufacturer's conditions required for AC timing specifications; i.e., ringback, etc..), and the output pin of the driving agent crossing V_{REF} if **the driver was driving the *Test Load*** used to specify the driver's AC timings.

See Section 3.7.2. for details regarding flight time simulation and validation.

Figure 13 in Appendix A shows the *V_{REF} Guardband* boundaries where maximum and minimum flight time measurements are taken. The V_{REF} Guardband takes into account sources of noise that may affect the way an AGTL+ signal becomes valid at the receiver. See the definition of the V_{REF} Guardband.

- **Maximum and Minimum Flight Time** — Flight time variations can be caused by many different parameters. The more obvious causes include variation of the board dielectric constant, changes in load condition, crosstalk, V_{TT} noise, V_{REF} noise, variation in termination resistance and differences in I/O buffer performance as a function of temperature, voltage and manufacturing process. Some less obvious causes include effects of *Simultaneous Switching Output (SSO)* and packaging affects.
- The **Maximum Flight Time** is the largest flight time a network will experience under all variations of conditions. Maximum flight time is measured at the appropriate V_{REF} Guardband boundary.
- The **Minimum Flight Time** is the smallest flight time a network will experience under all variations of conditions. Minimum flight time is measured at the appropriate V_{REF} Guardband boundary.

For more information on flight time and the V_{REF} Guardband, see Appendix A of this guideline and the *Pentium® II Processor Developer's Manual.*

GTL+ is the bus technology used by the Pentium Pro processor. This is an incident wave switching, open drain bus with pull-up resistors which provide both the high logic level and termination. It is an enhancement to the GTL (Gunning Transceiver Logic) technology. See the *Pentium® II Processor Developer's Manual* for more details of GTL+.

Network — the trace of a Printed Circuit Board (PCB) that completes an electrical connection between two or more components.

Network Length — the distance between extreme bus agents on the network and does not include the distance connecting the end bus agents to the termination resistors.

Overdrive Region — is the voltage range, at a receiver, located above and below V_{REF} for signal integrity analysis. See the *Pentium® II Processor Developer's Manual* for more details.

Overshoot — Maximum voltage allowed for a signal at the processor core pad. See the *Pentium® III Xeon™ Processor at 500 and 550 MHz* datasheet and *Intel® 450NX PCIset* for overshoot specifications.

Pad — a feature of a semiconductor die contained within an internal logic package on the cartridge substrate used to connect the die to the package bond wires. A pad is only observable in simulation.

Pin — a feature of a logic package contained within the S.E.C. cartridge used to connect the package to an internal substrate trace.

Ringback — is the voltage that a signal rings back to after achieving its maximum absolute value. Ringback may be due to reflections, driver oscillations, etc. See the *Pentium® III Xeon™ Processor at 500 and 550 MHz* datasheet and *Intel® 450NX PCIset* for ringback specifications.

Settling Limit — defines the maximum amount of ringing at the receiving pin that a signal must reach before its next transition. See the *Pentium® III Xeon™ Processor at 500 and 550 MHz* datasheet and *Intel® 450NX PCIset* for settling limit specifications.

Setup Window — is the time between the beginning of Setup to Clock (T_{SU_MIN}) and the arrival of a valid clock edge. This window may be different for each type of bus agent in the system.

Simultaneous Switching Output (SSO) Effects — refers to the difference in electrical timing parameters and degradation in signal quality caused by multiple signal outputs simultaneously switching voltage levels (i.e., high-to-low) in the opposite direction from a single signal (e.g., low-to-high) or in the same direction (i.e., high-to-low). These are respectively called odd-mode switching and even-mode switching. This simultaneous switching of multiple outputs creates higher current swings that may cause additional propagation delay (or "pushout"), or a decrease in propagation delay (or "pull-in"). These SSO effects may impact the setup and/or hold times and are not always taken into account by simulations. System timing budgets should include margin for SSO effects.

Stub — the branch from the trunk terminating at the pad of an agent.

Test Load — Intel uses a 25 Ω test load for specifying its components.

Trunk — the main connection, excluding interconnect branches, terminating at agent pads.

Undershoot — Maximum voltage allowed for a signal to extend below V_{SS} at the processor core pad. See the *Pentium® III Xeon™ Processor at 500 amd 550 MHz* datasheet and *Intel® 450NX PCIset* for undershoot specifications.

Victim — a network that receives a coupled crosstalk signal from another network is called the victim network.

V_{REF} Guardband — A guardband (ΔV_{REF}) defined above and below V_{REF} to provide a more realistic model accounting for noise such as crosstalk, V_{TT} noise, and V_{REF} noise.

3.0. AGTL+ DESIGN GUIDELINE

The following step-by-step guideline was developed for systems based on four Pentium III Xeon processor loads, one MIOC load, and one optional cluster controller.

The guideline recommended in this section is based on experience developed at Intel while developing many different Pentium Pro processor family and Pentium II Xeon processor based systems. Begin with component selection, an initial timing analysis, and topology definition. Perform pre-layout analog simulations for a detailed picture of a working "solution space" for the design. These pre-layout simulations help define routing rules prior to placement and routing. After routing, extract the interconnect database and perform post-layout simulations to refine the timing and signal integrity analysis. Validate the analog simulations when actual systems become available. The validation section describes a method for determining the flight time in the actual system.

Guideline Methodology:

- Determine Components
- Initial Timing Analysis
- Determine General Topology, Layout, and Routing
- Pre-Layout Simulation (Sensitivity sweep)
- Place and Route Board
  - — Estimate Component to Component Spacing for AGTL+ Signals
  - — Layout and Route Board
- Post-Layout Simulation
  - — Interconnect Extraction
  - — Intersymbol Interference (ISI), Crosstalk, and Monte Carlo Analysis
- Validation
  - — Measurements
  - — Determining Flight Time

3.1. Determine Components

Determine whether a cluster controller will be used and whether it will reside directly on the PCB or occupy a fifth Single Edge Contact (SEC) connector slot.

3.2. Initial Timing Analysis

Perform an initial timing analysis of the system using Equation 1 and Equation 2 shown below. These equations are the basis for the timing analysis. To complete the initial timing analysis, values for clock skew and clock jitter are needed, along with the component specifications. These equations contain a multi-bit adjustment factor, M_{ADJ}, to account for multi-bit switching effects such as SSO pushout or pull-in that are often hard to simulate. These equations **do not** take into consideration all signal integrity factors that affect timing. Additional timing margin should be budgeted to allow for these sources of noise.

Equation 1. Setup Time

$$T_{CO\_MAX} + T_{SU\_MIN} + CLK_{SKEW} + CLK_{JITTER} + T_{FLT\_MAX} + M_{ADJ} \leq \text{Clock Period}$$

Equation 2. Hold Time

$$T_{CO\_MIN} + T_{FLT\_MIN} - M_{ADJ} \geq T_{HOLD} + CLK_{SKEW}$$

Symbols used in Equation 1 and Equation 2:

- T_{CO_MAX} is the maximum clock to output specification.[1]
- T_{SU_MIN} is the minimum required time specified to setup before the clock.[1]
- CLK_{JITTER} is the maximum clock edge-to-edge variation.
- CLK_{SKEW} is the maximum variation between components receiving the same clock edge.
- T_{FLT_MAX} is the maximum flight time as defined in Section 2.3.
- T_{FLT_MIN} is the minimum flight time as defined in Section 2.3.
- M_{ADJ} is the multi-bit adjustment factor to account for SSO pushout or pull-in.
- T_{CO_MIN} is the minimum clock to output specification.[1]
- T_{HOLD} is the minimum specified input hold time.

NOTE

1. The Clock to Output (T_{CO}) and Setup to Clock (T_{SU}) timings are both measured from the signals last crossing of V_{REF}, with the requirement that the signal does not violate the ringback or edge rate limits. See the *Pentium® III Xeon™ Processor at 500 and 550 MHz* datasheet and the *Pentium® II Processor Developer's Manual* for more details.

Solving these equation for T_{FLT} results in the following equations:

Equation 3. Maximum Flight Time

$$T_{FLT\_MAX} \leq \text{Clock Period} - T_{CO\_MAX} - T_{SU\_MIN} - CLK_{SKEW} - CLK_{JITTER} - M_{ADJ}$$

Equation 4. Minimum Flight Time

$$T_{FLT\_MIN} \geq T_{HOLD} + CLK_{SKEW} - T_{CO\_MIN} + M_{ADJ}$$

There are multiple cases to consider. Note that while the same trace connects two components, component A and component B, the minimum and maximum flight time requirements for component A driving component B as well as component B driving component A must be met. The cases to be considered are:

- Pentium® III Xeon™ processor driving Pentium III Xeon processor
- Pentium III Xeon processor driving MIOC
- MIOC driving a Pentium III Xeon processor
- Pentium III Xeon processor driving cluster controller
- Cluster controller driving Pentium III Xeon processor
- MIOC driving cluster controller
- Cluster controller driving MIOC

A designer using components other than those listed above must evaluate any additional combinations of driver and receiver.

Table 1 lists the AGTL+ component timings of the Pentium III Xeon processor and MIOC defined at the pins. **These timings are for reference only; obtain component specifications from the *Pentium® III Xeon™ Processor at 500 and 550 MHz* datasheet and *Intel® 450NX PCIset.***

Table 1. Pentium® III Xeon™ Processor and MIOC AGTL+ Parameters[1, 2]

IC Parameters	Pentium® III Xeon™ Processor at 90 MHz or 100 MHz Bus	MIOC
Clock to Output maximum (T_{CO_MAX})	2.70	2.65
Clock to Output minimum (T_{CO_MIN})	-0.07	-0.15
Setup time (T_{SU_MIN})	1.75	1.58
Hold time (T_{HOLD})	0.62	0.63

NOTES

1. All times in nanoseconds.
2. Numbers in table are for reference only. These timing parameters are subject to change. Please check the appropriate component datasheets for valid timing parameter values.

Table 2 gives an example AGTL+ initial maximum flight time calculation for a 100 MHz, 4-way Pentium III Xeon processor/Intel 450NX PCIset design that does not include a cluster controller. Note that assumed values for clock skew and clock jitter were used. Clock skew and clock jitter values are dependent on the clock components and distribution method chosen for a particular design and must be budgeted into the initial timing equations as appropriate for each design.

Intel highly recommends adding margin as shown in the "M_{ADJ}" column to offset the degradation caused by SSO pushout and other multi-bit switching effects. The "Recommended T_{FLT_MAX}" column contains the recommended maximum flight time after incorporating the M_{ADJ} value. If edge rate, ringback, and monotonicity requirements are not met, flight time correction must be performed as documented in the *Pentium® II Processor Developer's Manual* with the additional requirements noted in Appendix A. The commonly used "textbook" equations used to calculate the expected signal propagation rate of a board are included in Section 5.1.

Simulation and control of baseboard design parameters can ensure that signal quality and maximum and minimum flight times are met. Baseboard propagation speed is highly dependent on transmission line geometry configuration (stripline vs. microstrip), dielectric constant, and loading. This layout guideline includes high-speed baseboard design practices that may improve the amount of timing and signal quality margin. The magnitude of M_{ADJ} is highly dependent on baseboard design implementation (stackup, decoupling, layout, routing, reference planes, etc.) and needs to be characterized and budgeted appropriately for each design.

Table 2. Example T_{FLT_MAX} Calculations for 100 MHz Bus[1]

Driver	Receiver	Clk Period (ns)	T_{CO_MAX} (ns)	T_{SU_MIN} (ns)	Clk_{SKEW} (ns)	Clk_{JITTER} (ns)	M_{ADJ} (ns)	Recommended T_{FLT_MAX}[2] (ns)
Pentium® III Xeon™ processor	Pentium III Xeon processor	10.00	2.70	1.75	0.15	0.15	0.80	4.45
Pentium III Xeon processor	MIOC	10.00	2.70	1.58	0.15	0.15	0.80	4.63
MIOC	Pentium III Xeon processor	10.00	2.65	1.75	0.15	0.15	0.50	4.80

NOTES:

1. All times in nanoseconds.
2. The flight times in this column include margin to account for the following phenomena which Intel has observed when multiple bits are switching simultaneously. These multi-bit effects can adversely affect flight time and signal quality and are sometimes not accounted for in simulation. Accordingly, maximum flight times depend on the baseboard design and additional adjustment factors or margins are recommended.

   - SSO pushout or pull-in
   - Rising or falling edge rate degradation at the receiver caused by inductance in the current return path, requiring extrapolation that causes additional delay.
   - Crosstalk on the PCB and internal to the package can cause variation in the signals.

   There are additional effects that **may not** necessarily be covered by the multi-bit adjustment factor and should be budgeted as appropriate to the baseboard design. Examples include:

   - The effective board propagation constant (S_{EFF}), which is a function of:
     - — Dielectric constant (ε_r) of the PCB material
     - — The type of trace connecting the components (stripline or microstrip)
     - — The length of the trace and the load of the components on the trace. (Note that the board propagation constant multiplied by the trace length is a **component** of the flight time but **not necessarily equal to** the flight time.)

Table 3. Example T_{FLT_MAX} Calculations for 90 MHz Bus (Cluster Controller Design)[1]

Driver	Receiver	Clk Period	T_{CO_MAX}	T_{SU_MIN}	Clk_{SKEW}	Clk_{JITTER}	M_{ADJ}	Recommended T_{FLT_MAX}[2]
Pentium® III Xeon™ processor	Pentium III Xeon processor	11.11	2.70	1.75	0.15	0.15	0.80	5.56
Pentium III Xeon processor	MIOC	11.11	2.70	1.58	0.15	0.15	0.80	5.73
MIOC	Pentium III Xeon processor	11.11	2.65	1.75	0.15	0.15	0.50	5.91

NOTES:

1. All times in nanoseconds.
2. The flight times in this column include margin to account for the following phenomena which Intel has observed when multiple bits are switching simultaneously. These multi-bit effects can adversely affect flight time and signal quality and are sometimes not accounted for in simulation. Accordingly, maximum flight times depend on the baseboard design and additional adjustment factors or margins are recommended.

   - SSO pushout or pull-in
   - Rising or falling edge rate degradation at the receiver caused by inductance in the current return path, requiring extrapolation that causes additional delay.
   - Crosstalk on the PCB and internal to the package can cause variation in the signals.

   There are additional effects that **may not** necessarily be covered by the multi-bit adjustment factor and should be budgeted as appropriate to the baseboard design. Examples include:

   - The effective board propagation constant (S_{EFF}), which is a function of:
     - — Dielectric constant (ε_r) of the PCB material
     - — The type of trace connecting the components (stripline or microstrip)
     - — The length of the trace and the load of the components on the trace. (Note that the board propagation constant multiplied by the trace length is a **component** of the flight time but **not necessarily equal to** the flight time.)

Table 4. Example T_{FLT_MIN} Calculations (Frequency Independent)[1]

Driver	Receiver	T_{HOLD}	Clk_{SKEW}	T_{CO_MIN}	M_{ADJ}	Recommended T_{FLT_MIN}[2]
Pentium® III Xeon™ processor	Pentium III Xeon processor	0.62	0.15	-0.07	0.36	1.20
Pentium III Xeon processor	MIOC	0.63	0.15	-0.07	0.35	1.20
MIOC	Pentium III Xeon processor	0.62	0.15	-0.15	0.28	1.20

NOTE:

1. All times in nanoseconds.
2. The flight times in this column include margin to account for the following phenomena which Intel has observed when multiple bits are switching simultaneously. These multi-bit effects can adversely affect flight time and signal quality and are sometimes not accounted for in simulation. Accordingly, maximum flight times depend on the baseboard design and additional adjustment factors or margins are recommended.

   - SSO pushout or pull-in
   - Rising or falling edge rate degradation at the receiver caused by inductance in the current return path, requiring extrapolation that causes additional delay.
   - Crosstalk on the PCB and internal to the package can cause variation in the signals.

   There are additional effects that **may not** necessarily be covered by the multi-bit adjustment factor and should be budgeted as appropriate to the baseboard design. Examples include:

   - The effective board propagation constant (S_{EFF}), which is a function of:
     - — Dielectric constant (ε_r) of the PCB material
     - — The type of trace connecting the components (stripline or microstrip)
     - — The length of the trace and the load of the components on the trace. (Note that the board propagation constant multiplied by the trace length is a **component** of the flight time but **not necessarily equal to** the flight time.)

Table 3 gives a similar example maximum flight time calculation for a 90 MHz AGTL+ 4-way Pentium III Xeon processor / Intel 450NX PCIset design that includes a cluster controller.

Table 4 is an example calculation for minimum flight time that is frequency independent. Intel highly recommends adding margin as shown in the "M_{ADJ}" column to offset the degradation caused by SSO pull-in and other multi-bit switching effects. The "Recommended T_{FLT_MIN}" column contains the recommended minimum flight time after incorporating the M_{ADJ} value.

Table 2, Table 3, and Table 4 are derived assuming:

- CLK_{SKEW} = 0.15 ns (PCB skew only – assumes zero driver skew by tying clock driver outputs)
- CLK_{JITTER} = 0.15 ns

3.3. Determine General Topology, Layout, and Routing Desired

After the selecting the processor bus components and calculating the timing budget, determine the approximate location of the devices on the baseboard. Estimate the printed circuit board parameters from the placement and other information. Locate the processors, Intel 450NX PCIset, and cluster controller as required to meet timing. The "Double Star" and "Crow's Foot" Topologies illustrated in Figure 1 have been shown in simulation to be successful for 6-load, 90 MHz bus operation. 100 MHz bus operation requires that the cluster controller be removed and replaced with AGTL+ termination or a terminator card.

Figure 2 shows a "Double Star" and "Crow's Foot" topology modified to provide more noise and timing margin for 5-load, 100 MHz operation. The modification involves complete removal of the 6th stub and AGTL+ termination that should improve worst case flight time margin on falling edge transition. Doing so will reduce the bus load, and should provide more noise and timing margin. Perform simulations on the entire system bus to ensure that ideal termination resistance at the MIOC is chosen for any given design. Analysis at Intel suggests that for the 5-load Crow's Foot topology in Figure 2, the AGTL+ termination at the MIOC may be decreased to 75 Ω to maintain an effective AGTL+ termination of 25 Ω. For the 5-load Double Star topology, keeping the MIOC AGTL+ termination at 150 Ω should provide better signal quality and timing margin.

Six-load, cluster-capable systems may gain timing and signal quality margin by using faster dielectric material (e.g., Getek) and better ground plane referencing for AGTL+ signals.

- Place termination resistors at the MIOC and cluster controller, which should be located at opposite ends of the AGTL+ network. Minimize the inductance between the V_{TT} distribution and the termination resistors.
- The placement of the Pentium® III Xeon™ processor, MIOC and/or custom ASIC(s) on the Pentium III Xeon processor system bus must be carefully chosen.

Using a custom ASIC (with different timings than Pentium III Xeon processor or MIOC) on the system bus will require additional analog simulations to determine the optimum location of each agent along the bus.

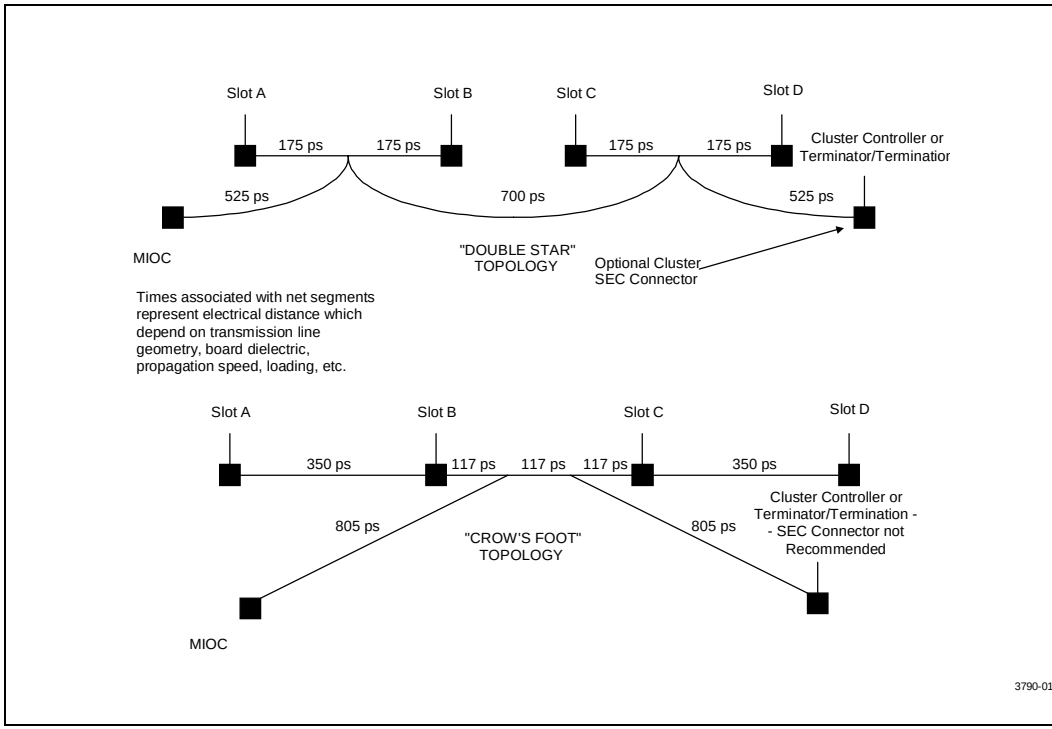

Figure 1. Example 6-load and 5-load Plus 6th Stub Termination Network Topology

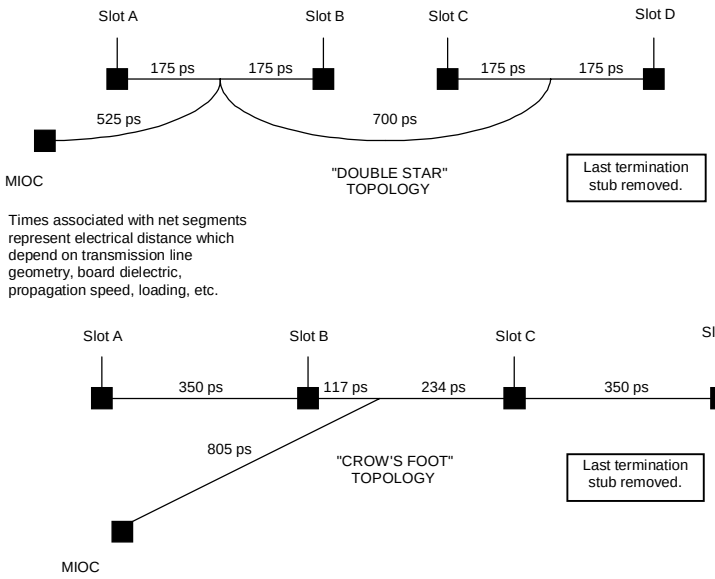

Figure 2. Example 5-load Network Topology Optimized for 100 MHz Bus

3.4. Pre-Layout Simulation

3.4.1. METHODOLOGY

Pentium III Xeon processor designs require analog simulations. Start simulations prior to layout. Pre-layout simulations provide a detailed picture of the working "solution space" that meets flight time and signal quality requirements. The layout recommendations in the previous sections are based on pre-layout simulations conducted at Intel. By basing board layout guidelines on the solution space, the iterations between layout and post-layout simulation can be reduced.

Intel specifies signal quality **at the device pads** and therefore recommends running simulations at the device pads for signal quality. However, the core timings are specified at the **device pins**, so simulation results at the device pins may be used to correlate simulation performance against actual system measurements

Pre-layout analysis includes a sensitivity analysis using parametric sweeps. Parametric sweep analysis involves varying one or two system parameters while all others such as driver strength, package, Z_0, and S_0 are constant. This way, the sensitivity of the proposed bus topology to varying parameters can be analyzed systematically. Sensitivity of the bus to minimum flight time, maximum flight time, and signal quality should be covered. Suggested sweep parameters include trace lengths, termination resistor values, and any other factors that may affect flight time, signal quality, and feasibility of layout. Minimum flight time and worst signal quality are typically analyzed using fast I/O buffers and interconnect. Maximum flight time is typically analyzed using slow I/O buffers and slow interconnect.

Outputs from each sweep should be analyzed to determine which regions meet timing and signal quality specifications. To establish the working solution space, find the common space across all the sweeps that result in passing timing and signal quality. The solution space should allow enough design flexibility for a feasible, cost-effective layout.

3.4.2. SIMULATION CRITERIA

Accurate simulations require that the actual range of parameters be used in the simulations. Intel has consistently measured the cross-sectional resistivity of PCB copper to be in the order of 1 ohm*mil^2/inch, not the 0.662 ohm*mil^2/inch value for annealed copper that is published in reference material. Using the 1 ohm*mil^2/inch value may increase the accuracy of lossy simulations.

Positioning drivers with faster edges closer to the middle of the network typically results in more noise than positioning them towards the ends. We have also shown that the worst-case noise margin can be generated by drivers located in all positions (given appropriate variations in the other network parameters). Therefore, we recommend simulating the networks from all driver locations, and analyzing each receiver for each possible driver.

Analysis has shown that **both fast and slow corner models** must be run for both rising and falling edge transitions. The fast corner is needed because the fast edge rate creates the most noise. The slow corner is needed because the buffer's drive capability will be a minimum, causing the V_{OL} to shift up, which may cause the noise from the slower edge to exceed the available budget. Slow corner models may produce minimum flight time violations on rising edges if the transition starts from a higher V_{OL}. So, Intel **highly recommends** checking for minimum and maximum flight time violations with both the fast and slow corner models.

The transmission line package models must be inserted between the output of the buffer and the net it is driving. Likewise, the package model must also be placed between a net and the input of a receiver model. This is generally done by editing your simulator's net description or topology file.

Intel has found wide variation in noise margins when varying the stub impedance and the PCB's Z_0 and S_0. Intel therefore recommends that PCB parameters be controlled as tightly as possible, with a sampling of the allowable Z_0 and S_0 simulated. The recommended effective line impedance (Z_{EFF}) is 65 Ω +/-10%. Intel recommends running uncoupled simulations using the Z_0 of the package stubs; and performing fully coupled simulations if increased accuracy is needed or desired. Accounting for crosstalk within the device package by varying the stub impedance was investigated and was not found to be sufficiently accurate. This lead to the development of full package models for the component packages.

3.5. Place and Route Board

3.5.1. ESTIMATE COMPONENT TO COMPONENT SPACING FOR AGTL+ SIGNALS

Estimate the number of layers that will be required. Then determine the expected interconnect distances between each of the components on the AGTL+ bus. Be sure to consider the guidelines in Section 3.3. Using the estimated interconnect distances, verify that the placement can support the system timing requirements.

The maximum network length between the bus agents is determined by the required bus frequency and the maximum flight time propagation delay on the PCB. The minimum network length is independent of the required bus frequency. Table 2, Table 3, and Table 4 assume values for CLK_{SKEW} and CLK_{JITTER}, parameters that are controlled by the system designer. In order to reduce system clock skew to a minimum, clock buffers which allow their outputs to be tied together are recommended. Intel suggests running analog simulations to ensure that each design has adequate noise and timing margin.

3.5.2. LAYOUT AND ROUTE BOARD

Route the board satisfying the estimated space and timing requirements. Stay within the solution space set from the pre-layout sweeps. Estimate the printed circuit board parameters from the placement and other information including the following general guidelines:

- Distribute V_{TT} with a power plane or a partial power plane. If this cannot be accomplished, use as wide a trace as possible and route the V_{TT} trace with the same design rules as the AGTL+ traces.
- Keep the overall length of the bus as short as possible (but don't forget minimum component to component distances to meet hold times).
- Plan to minimize crosstalk by:
  - — Routing the same type of AGTL+ I/O signals in isolated signal groups. I.e., route the data signals in one group, the arbitration signals in another group. Keep at least a 5:1 spacing to trace width ratio between each group.
  - — Keeping at least a 25 mil space between AGTL+ signals and non-AGTL+ signals (and at least a 5:1 spacing to line width ratio).
  - — Keeping at least a 4:1 spacing to trace width ratio between AGTL+ signals in the same group.

— Using a trace pitch to trace height ratio of 3:1 between AGTL+ signals in the same group.

— Using a trace pitch to trace height ratio of 4:1 between AGTL+ signals in different groups.

— Using a trace pitch to trace height ratio of 5:1 between AGTL+ signals and non-AGTL+ signals

— Minimizing the dielectric process variation used in the PCB fab.

— Eliminating parallel traces between layers not separated by a power or ground plane.

The spacing between the various bus agents causes variations in trunk impedance and stub locations. These variations cause reflections which can cause constructive or destructive interference at the receivers. A reduction of noise may be obtained by a minimum spacing between the agents. Unfortunately, a tighter spacing results in reduced component placement options and lower hold margins. Therefore adjusting the inter-agent spacing may be one way to change the network's noise margin, but mechanical constraints often limit the usefulness of this technique. Always be sure to validate signal quality after making any changes in agent locations or changes to inter-agent spacing.

There are six AGTL+ signals that can be driven by more than one agent simultaneously. These signals may require more attention during the layout and validation portions of the design. When a signal is asserted (driven low) by two or more agents on the same clock edge, the two falling edge wave fronts will meet at some point on the bus and can sum to form a negative voltage. The ringback from this negative voltage can easily cross into the overdrive region. The signals are AERR#, BERR#, BINIT#, BNR#, HIT#, and HITM#.

This document addresses AGTL+ layout. Chassis requirements for cooling, connector location, memory location, etc., may constrain the system topology and component placement location, therefore constraining the board routing. These issues are not directly addressed in this document.

3.6. Post-Layout Simulation

Following layout, extract the interconnect information for the board from the CAD layout tools. Run post-layout simulations to verify that the layout meets timing and noise requirements. A small amount of "tuning" may be required; experience at Intel has shown that sensitivity analysis dramatically reduces the amount of tuning required. The post layout simulations should take into account the expected variation for all interconnect parameters.

Intel specifies signal quality **at the device pads** and therefore recommends running simulations at the device pads for signal quality. However, Intel specifies core timings **at the device pins,** so simulation results at the device pins should be used later to correlate simulation performance against actual system measurements.

3.6.1. INTERSYMBOL INTERFERENCE

Intersymbol Interference (ISI) refers to the distortion or change in the waveform shape caused by the voltage and transient energy on the network when the driver begins its next transition.

For example, ISI may occur when the line is driven high, low, and then high in consecutive cycles (the opposite case is also valid). When the driver drives high on the first cycle and low on the second cycle, the signal may not settle to the minimum V_{OL} before the next rising edge is driven. This results in improved flight times in the third cycle. ISI simulations for the topology given in this section were performed by comparing flight times for the first and third cycle. ISI effects do not necessarily span only 3 cycles so it may be necessary to simulate beyond 3 cycles for certain designs. After simulating and quantifying ISI effects, adjust the timing budget accordingly to take these conditions into consideration.

3.6.2. CROSSTALK ANALYSIS

AGTL+ crosstalk simulations can consider the Pentium III Xeon processor core package, MIOC package, and SC330 connectors as non-coupled. Treat the traces on the Pentium III Xeon processor cartridge and baseboard as fully coupled for maximum crosstalk conditions. Simulate the traces as lossless for worst case crosstalk, and lossy where more accuracy is needed. Evaluate both odd and even mode crosstalk conditions.

AGTL+ crosstalk simulation involves the following cases:

- Intra-group AGTL+ crosstalk
- Inter-group AGTL+ crosstalk
- CMOS to AGTL+ crosstalk

3.6.3. MONTE CARLO ANALYSIS

Perform a Monte Carlo analysis to refine the passing solution space region. A Monte Carlo analysis involves

randomly varying parameters (independent of one another) over their tolerance range. This analysis is intended to ensure that no regions of failing flight time and signal quality exist between the extreme corner cases run in pre-layout simulations. For the example topology, vary the following parameters during Monte Carlo simulations:

- Trace lengths on baseboard
- AGTL+ termination resistance R_{TT} on Pentium® III Xeon™ processor cartridges
- AGTL+ termination resistance R_{TT} at the MIOC
- AGTL+ termination resistance R_{TT} at the Cluster Controller (if present)
- AGTL+ termination resistance R_{TT} on termination cards or 6th stub (if present).
- Z_0 of traces on Pentium III Xeon processor cartridges
- S_0 of traces on Pentium III Xeon processor cartridges
- Z_0 of traces on baseboard
- S_0 of traces on baseboard
- Fast and slow corner Pentium III Xeon processor I/O buffer models
- Fast and slow Pentium III Xeon processor package models
- Fast and slow corner Intel® 450NX PCIset I/O buffer models
- Fast and slow Intel 450NX PCIset package models

Refer to the *Pentium® III Xeon™ Processor at 500 and 550 MHz* datasheet and electronic I/O buffer models for the parameter ranges of the Pentium III Xeon processor and Intel 450NX PCIset.

3.7. Validation

Build systems and validate the design and simulation assumptions.

3.7.1. MEASUREMENTS

Note that the AGTL+ specification for signal quality is at the **pad** of the component. The expected method of determining the signal quality is to run analog simulations for the pin and the pad. Then correlate the simulations at the pin against actual system measurements at the pin. Good correlation at the pin leads to confidence that the simulation at the pad is accurate. Controlling the temperature and voltage to correspond to the I/O buffer model extremes should enhance the correlation between simulations and the actual system.

3.7.2. FLIGHT TIME SIMULATION

As defined earlier in Section 2.3., flight time is the time difference between a signal crossing V_{REF} at the input pin of the receiver, and the output pin of the driver crossing V_{REF} were it driving a ***test load***. The timings in the tables and topologies discussed in this guideline assume the actual system load is 25 Ω and is **equal to the *test load***. This may not be the case in a particular design and this section describes how to correlate the design load to the test load in simulation and in validation.

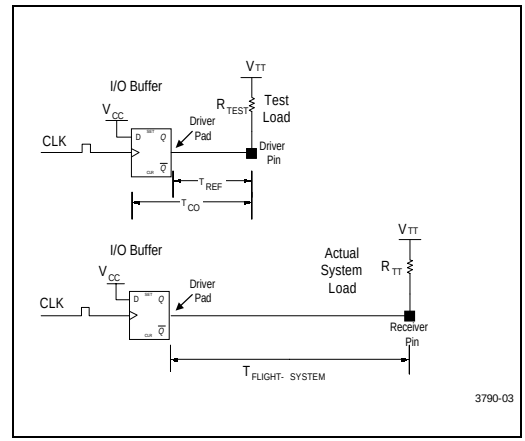

Figure 3. Test Load vs. Actual System Load

Figure 3 above shows the different configurations for T_{CO} testing and flight time simulation. The flip-flop represents the logic input and driver stage of a typical AGTL+ I/O buffer. T_{CO} timings are specified at the driver pin output. $T_{FLIGHT\text{-}SYSTEM}$ is usually reported by a simulation tool as the time from the driver pad starting its transition to the time when the receiver's input pin sees a valid data input. Since both timing numbers (T_{CO} and $T_{FLIGHT\text{-}SYSTEM}$) will include propagation time from the pad to the pin, it is necessary to subtract this time (T_{REF}) from the reported flight time to avoid double counting. T_{REF} is defined as the time that it takes for the driver output pin to reach the measurement voltage, V_{REF}, starting from the beginning of the driver transition at the pad. T_{REF} must be generated using the same test load for T_{CO}. Intel provides this timing value in the AGTL+ I/O buffer models.

In this manner, the following *valid delay* equation is satisfied:

Valid Delay Equation

$$\text{Valid Delay} = T_{CO} + T_{FLIGHT\text{-}SYS} - T_{REF} = T_{CO\text{-}MEASURED} + T_{FLIGHT\text{-}MEASURED}$$

This valid delay equation is the total time from when the driver sees a valid clock pulse to the time when the receiver sees a valid data input.

3.7.3. FLIGHT TIME HARDWARE VALIDATION

When a measurement is made on the actual system, T_{CO} and flight time do not need correction since these are the actual numbers. These measurements include all of the effects pertaining to the driver-system interface and the same is true for the T_{CO}. Therefore the addition of the measured T_{CO} and the flight time must be equal to the valid delay calculated above.

4.0. THEORY

4.1. AGTL+

AGTL+ is the electrical bus technology used for the Pentium III Xeon processor bus. This is an incident wave switching, open-drain bus with external pull-up resistors that provide both the high logic level and termination at each load. The Pentium III Xeon processor AGTL+ drivers contain a full cycle active pull-up device to improve system timings. The AGTL+ specification defines:

- Termination voltage (V_{TT}).
- Receiver reference voltage (V_{REF}) as a function of termination voltage (V_{TT}).
- Pentium® III Xeon™ processor termination resistance (R_{TT}).
- Input low voltage (V_{IL}).
- Input high voltage (V_{IH}).
- NMOS on resistance (R_{ONN}).
- Edge rate specifications.
- Ringback specifications.
- Overshoot/Undershoot specifications.
- Settling Limit.

The complete AGTL+ specification can be found in the *Pentium® III Xeon™ Processor at 500 and 550 MHz* datasheet. Layout recommendations for the AGTL+ bus can be found in Section 3.0. of this document.

4.2. Timing Requirements

The system timing for AGTL+ is dependent on many things. Each of the following elements combine to determine the maximum and minimum frequency the AGTL+ bus can support:

- The range of timings for each of the agents in the system.
  - Clock to output [T_{CO}]. (Note that the system load is likely to be different from the “specification” load therefore the T_{CO} observed in the system may not be the same as the T_{CO} from the specification.)
  - The minimum required setup time to clock [T_{SU_MIN}] for each receiving agent.
- The range of flight time between each component. This includes:
  - The velocity of propagation for the loaded printed circuit board [S_{EFF}].
  - The board loading impact on the effective T_{CO} in the system.
- The amount of skew and jitter in the system clock generation and distribution.
- Changes in flight time due to crosstalk, noise, and other effects.

4.3. Noise Margin

The goal of this section is to describe the total amount of noise that can be tolerated in a system (the noise budget), identify the sources of noise in the system, and recommend methods to analyze and control the noise so that the allowed noise budget is not exceeded.

There are several sources of noise which must be accounted for in the system noise budget, including:

- V_{REF} variation
- V_{CCCORE} variation
- Variation in V_{TT}
- Crosstalk

- Ringback due to impedance variation along the network, termination mismatch, and/or stubs on the network
- Simultaneous Switching Output Effects

The total noise budget is calculated by taking the difference in the worst case specified input level and the worst case driven output level.

Sections 4.3.1. and 4.3.2. discuss calculating noise margin. These sections do not discuss ringback tolerant receivers which can increase the effective noise margin. See the appropriate component datasheets for information about ringback tolerance.

4.3.1. FALLING EDGE OR LOW LEVEL NOISE MARGIN

Equation 5 below shows a method for calculating falling edge noise margin when the Pentium III Xeon processor is driving. An example calculation follows.

Equation 5. Low Level Noise Margin

$$\text{Noise Margin}_{\text{LOW LEVEL}} = V_{ILMAX} - V_{OLMAX}$$

$$\Rightarrow V_{REF\_MIN} - 100\text{ mV} - V_{OLMAX}$$

$$\Rightarrow [\,[\,2/3\,(\,V_{TT\_MIN})\,]\,-\,1\%\,]\,-\,100\text{ mV}\,-\,V_{OLMAX}$$

Symbols for Equation 5 are:

- V_{ILMAX} is the maximum specified valid input low level from the component specification. For this example, 100 mV below the reference voltage is assumed.
- V_{OLMAX} is the maximum output low level the component will drive. This V_{OLMAX} maximum condition corresponds to the slow corner components and models.
- V_{REF_MIN} is the minimum valid voltage reference used for the threshold reference.
- V_{TT_MIN} is the minimum termination voltage.

For the following example calculations for low level and high level noise margin, an R_{ON_MAX} equal to 12.5 Ω is assumed, along with V_{REF} and V_{TT} tolerance assumptions. These specs should be obtained from the *Pentium® III Xeon™ Processor at 500 and 550 MHz* datasheet.

Solving for V_{REF_MIN} with 1% V_{REF} uncertainty:

$$V_{REF\_MIN} = [\,2/3\,(\,V_{TT\_MIN})\,]\,-\,1\%$$
$$= [\,2/3\,(1.5\text{ V} - 9\%)\,]\,-\,1\%$$
$$= [\,2/3\,(1.37\text{ V})\,]\,-\,1\%$$
$$= 901\text{ mV}$$

The output low current in the case of V_{TT_MIN}, can be calculated as shown below:

$$I = V/R = 1.37/(25\ \Omega + 12.5\ \Omega) = 36.5\text{ mA}$$

Then the V_{OLMAX} for V_{REF_MIN} is (36.5 mA * 12.5 Ω) = 456 mV

Then,

$$\text{Noise Margin}_{\text{LOW LEVEL}}$$
$$= (V_{REF\_MIN} - 100\text{ mV}) - V_{OLMAX}$$
$$= (901\text{ mV} - 100\text{ mV}) - 456\text{ mV}$$
$$= 345\text{ mV}$$

These example calculations are for an effective termination resistance of 25 Ω. These calculations **do not** include any resistive drop along the trace.

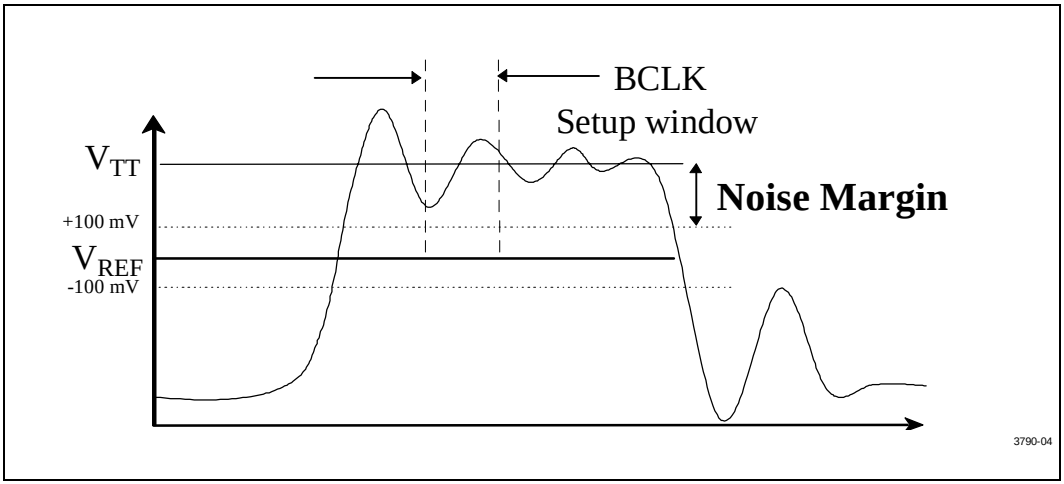

Figure 4. Rising Edge Noise Margin

4.3.2. RISING EDGE OR HIGH LEVEL NOISE MARGIN

Equation 6 below shows a method for calculating rising edge noise margin when the Pentium III Xeon processor is driving. An example calculation follows.

Equation 6. High Level Noise Margin

$$\text{Noise Margin}_{\text{HIGH LEVEL}} = V_{OH\_MIN} - V_{IH\_MIN} \Rightarrow$$
$$V_{TT\_MIN} - (V_{REF\_MAX} + 100\ \text{mV})$$

Symbols for Equation 6 are:

- V_{IH_MIN} is the minimum specified valid input high level from the component specification. For this example, 100 mV above the reference voltage is assumed.
- V_{OH_MIN} is the minimum output high level the component will drive.
- V_{TT_MIN} is the minimum termination voltage. This is assumed to be 1.5 V - 9%, or 1.37 V.
- V_{REF_MAX} is the maximum valid voltage reference used for the threshold reference. Since V_{REF} is defined as a function of $V_{TT,}$ the maximum V_{REF} with V_{TT_MIN} is 2/3 * (1.37 V) + 1% = 922 mV.
- V_{OH_MIN} for AGTL+ signals is $V_{TT\_MIN.}$
- Then $\text{Noise Margin}_{\text{HIGH LEVEL}}$
  $= V_{TT\_MIN} - (V_{REF\_MAX} + 100\ \text{mV})$
  $= 1.37\ \text{V} - 922\ \text{mV} - 100\ \text{mV}$
  $= 348\ \text{mV}$

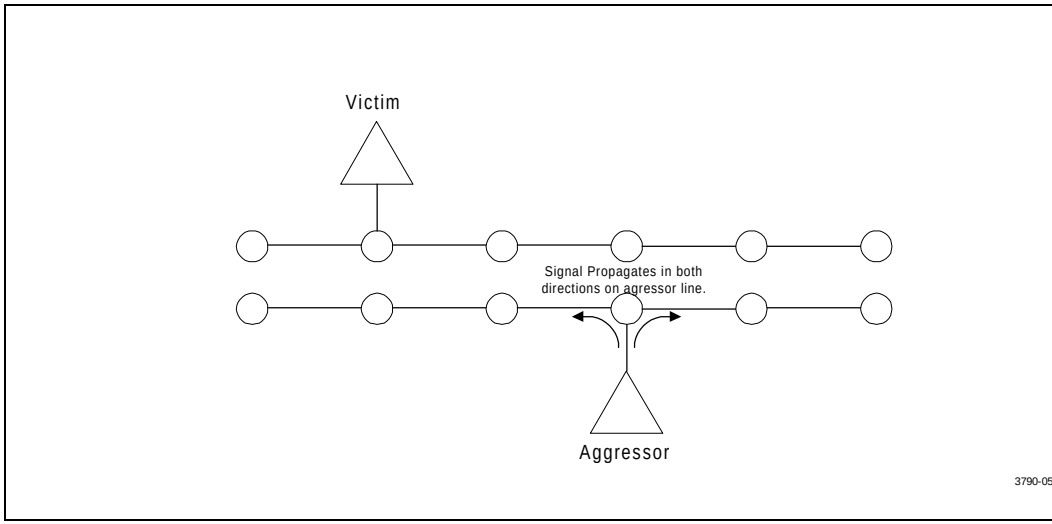

Figure 5. Propagation on Aggressor Network

4.4. Crosstalk Theory

AGTL+ signals swing across a smaller voltage range and have a correspondingly smaller noise margins than technologies that have traditionally been used in personal computer designs. This requires that designers using AGTL+ be more aware of crosstalk than they may have been in past designs.

Crosstalk is caused through capacitive and inductive coupling between networks. Crosstalk appears as both backward crosstalk and as forward crosstalk. Backward crosstalk creates an induced signal on a victim network that travels in a direction opposite that of the aggressor's signal. Forward crosstalk creates a signal that travels in the same direction as the aggressor's signal. On the AGTL+ bus, a driver on the aggressor network is not necessarily at the end of the network, therefore it sends signals in both directions on the aggressor's network. The signal propagating in each direction causes crosstalk on the victim network. Figure 5 shows a driver on the aggressor network and a receiver on the victim network which illustrates this effect. Figure 6 shows two aggressors on each side of the victim. Additional aggressors are possible in the z-direction, if adjacent signal layers are not routed in mutually perpendicular directions. Because crosstalk coupling coefficients decrease rapidly with increasing separation, it is rarely necessary to consider aggressors which are at least five line widths separated from the victim. The maximum crosstalk occurs when all the aggressors are switching in the same direction at the same time. There is crosstalk internal to the IC packages, which can also affect the signal quality.

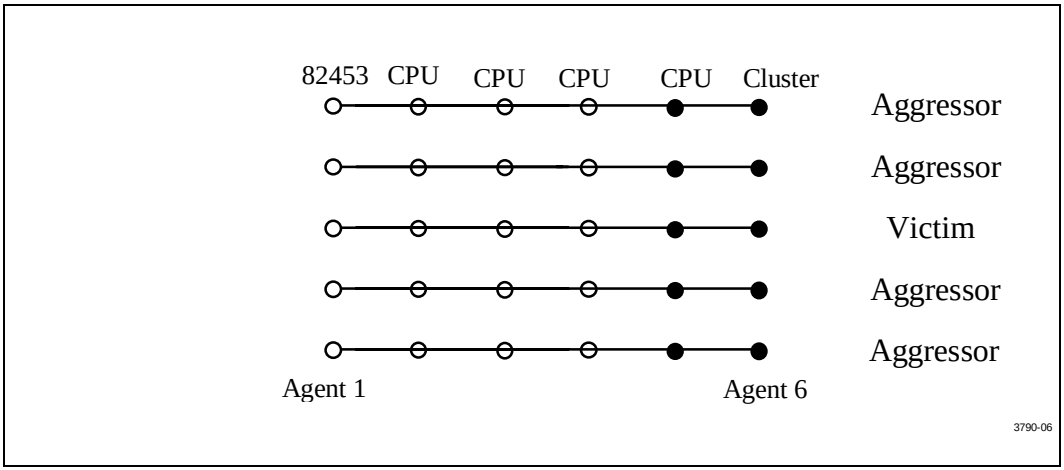

Figure 6. Aggressor and Victim Networks

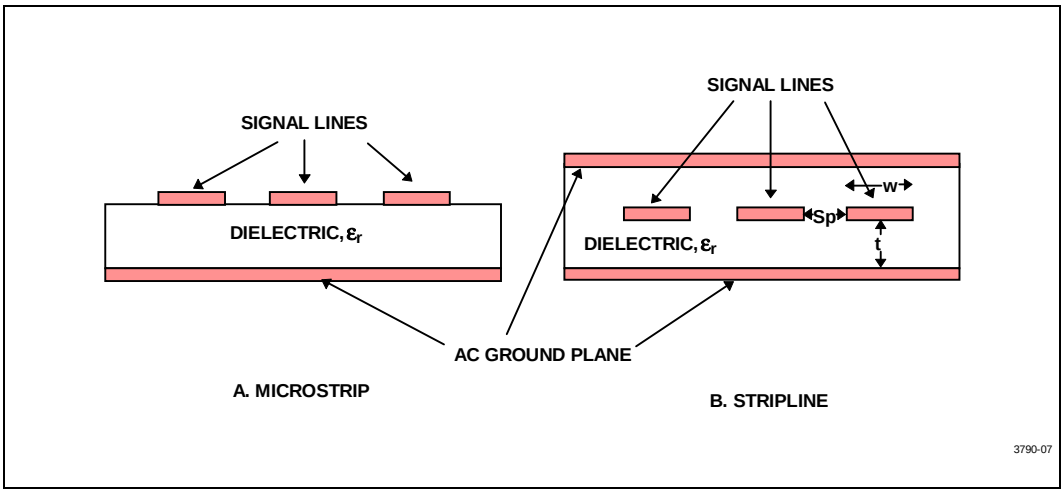

Figure 7. Transmission Line Geometry: (A) Microstrip (B) Stripline

Backward crosstalk is present in both stripline and microstrip geometry (see Figure 7). (A way to remember which geometry is stripline and which is microstrip is that a stripline geometry requires **stripping** a layer away to see the signal lines.) The backward coupled amplitude is proportional to the backward crosstalk coefficient, the aggressor's signal amplitude, and the coupled length of the network up to a maximum which is dependent on the rise/fall time of the aggressor's signal. Backward crosstalk reaches a maximum (and remains constant) when the propagation time on the coupled network length exceeds one half of the rise time of the aggressor's signal. Assuming the ideal ramp on the aggressor from 0% to 100% voltage swing, and the rise time on an unloaded coupled network, then:

$$Length\ for\ Max\ Backward\ Crosstalk = \frac{½ \times Rise\ Time}{Board\ Delay\ Per\ Unit\ Length}$$

An example calculation if fast corner fall time is 1.5 V/ns and board delay is 175 ps/inch (2.1 ns/foot) follows:

Table 5. Example Backward Crosstalk Coupling Factors with ε_r = 4.5, V_{OH_MAX} = 1.5 V, and Z_0 = 65 Ω

Space:Width:Thickness	Coupling Factor	Maximum Crosstalk
24:4:8	0.65%	9.8 mV
20:4:8	1.3%	19. 5 mV
16:4:8	1.75%	26.2 mV
14:4:8	2.5%	37.5 mV
12:4:8	3.4%	51.0 mV
8:4:8	6.55%	98.2 mV
4:4:8	13.5%	202.5 mV

Fall time = 1.5 V/1.5 V/ns = 1 ns

Length of maximum backward crosstalk =

½ * 1 ns * 1000 ps/ns /175 ps/in = 2.86 inches

Agents on the AGTL+ bus drive signals in each direction on the network. This will cause backward crosstalk from segments on two sides of a driver. The pulses from the backward crosstalk travel toward each other and will meet and **add** at certain moments and positions on the bus. This can cause the voltage (noise) from crosstalk to double.

Table 5 provides example coupling factors for various stripline space to width to dielectric thickness ratios (see Figure 6) with dielectric constant ε_r = 4.5, V_{OH_MAX} = 1.5 V, and Z_0 = 65 Ω. Note that the fast edge rates of falling edges place limits on the maximum coupled length allowable, and Table 5 illustrates the potential consequences of maximum coupled lengths. Also, it should be noted that multiple parallel coupled lines will increase the impact on the noise budget.

Forward crosstalk is absent in stripline topologies, but present in microstrip. (This is for the ideal case with a **uniform** dielectric constant. In actual boards, forward crosstalk is **nearly** absent in stripline topologies, but **abundant** in microstrip.) The forward coupled amplitude is proportional to the forward crosstalk coefficient, the aggressor's signal edge rate (*dv/dt*), and the coupled network's electrical length. The forward crosstalk coefficient is also a function of the geometry. Unlike backward crosstalk, forward crosstalk can grow with coupled section length, and may transition in a direction similar to or opposite to that of the aggressor's edge. Unlike backward crosstalk, forward crosstalk on the victim signal will continue to grow as it passes through more coupled length before the aggressor's wave front is absorbed by the termination.

4.4.1. CROSSTALK MANAGEMENT

To minimize crosstalk (and the "cost" of crosstalk) in terms of noise margin budget:

- Route adjacent trace layers in different directions (orthogonal preferred) to minimize the forward and backward crosstalk that can occur from parallel traces on adjacent layers. This reduces the source of crosstalk.

- Maximize the spacing between traces. Where traces have to be close and parallel to each other, minimize the distance that they are close together, and maximize the distance between sections that have close spacing. Routing close together could occur where multiple signals have to route between a pair of pins. When this happens the signals should be spread apart where possible. Also note that routing multiple layers in the same direction between reference planes can result in parallel traces that are close enough to each other to have significant crosstalk.
- Minimize the variation in board impedance (Z_0). For the example topologies covered in this guideline, 65 Ω +/- 10% was assumed.
- Minimize the nominal board impedance within the AGTL+ specification while maintaining the same trace width/spacing ratio. For a given dielectric constant, this reduces the trace width/trace height ratio, which reduces the backward and forward crosstalk coefficients. Having reduced crosstalk coefficients reduces the magnitude of the crosstalk.
- Minimize the dielectric constant used in the PCB fabrication. As above, all else being equal, this puts the traces closer to their reference planes and reduces the magnitude of the crosstalk.

- Watch out for voltage doubling at a receiving agent, caused by the adding of the backward crosstalk on either side of a driver. Minimize the total network length of signals that have coupled sections. If there has to be closely spaced/coupled lines, place them near the center of the net. This will cause the point in time that voltage doubling occurs to be before the setup window.
- Route synchronous signals that could be driven by different components in separate groups to minimize crosstalk between these groups. The Pentium® III Xeon™ processor uses a split transaction bus with six independent sub buses (arbitration, request, error, snoop, response, and data). This implies that in a given clock cycle, each sub bus could be driven by a different agent. If these two agents are at the opposite process corner (one fast and one slow), then separating the bus types will reduce the impact of crosstalk.

Simulation shows that space to line to dielectric ratios of less than 3:1:2 can produce excessive crosstalk between networks on the Pentium III Xeon processor bus. This is due to the lower voltage swing of AGTL+, high frequencies (even with the controlled edge rate buffers) and likely long parallel traces. Also, while rising edge rates are controlled, falling edge rates are not as well controlled.

4.4.2. POTENTIAL TERMINATION CROSSTALK PROBLEMS

The use of conventional “pull-up” resistor networks for Intel 450NX PCIset and cluster controller termination may not be suitable. These networks have a common power or ground pin at the extreme end of the package, shared by 13 to 19 resistors (for 14- and 20-pin components). These packages generally have too much inductance to maintain the voltage/current needed at each resistive load. Intel recommends using discrete resistors, resistor networks that have separate power/ground pins for each resistor, or working with a resistor network vendor to obtain resistor networks that have acceptable characteristics.

5.0. MORE DETAILS AND INSIGHTS

5.1. Textbook Timing Equations

The textbook equations used to calculate the propagation rate of a PCB are the basis for spreadsheet calculations for timing margin based on the component parameters. These equations are:

Equation 7. Intrinsic Impedance

$$Z_0 = \sqrt{\frac{L_0}{C_0}}\ (\Omega)$$

Equation 8. Stripline Intrinsic Propagation Speed

$$S_{0\_STRIPLINE} = 1.017 * \sqrt{\varepsilon_r}\ (\text{ns/ft})$$

Equation 9. Microstrip Intrinsic Propagation Speed

$$S_{0\_MICROSTRIP} = 1.017 * \sqrt{0.475 * \varepsilon_r + 0.67}\ (\text{ns/ft})$$

Equation 10. Effective Propagation Speed

$$S_{EFF} = S_0 * \sqrt{1 + \frac{C_D}{C_0}}\ (\text{ns/ft})$$

Equation 11. Effective Impedance

$$Z_{EFF} = \frac{Z_0}{\sqrt{1 + \frac{C_D}{C_0}}}\ (\Omega)$$

Equation 12. Distributed Trace Capacitance

$$C_0 = \frac{S_0}{Z_0}\ (\text{pF/ft})$$

Equation 13. Distributed Trace Inductance

$$L_0 = 12 * Z_0 * S_0\ (\text{nH/ft})$$

Symbols for Equation 7 through Equation 13 are:

- S_0 is the speed of the signal on an unloaded PCB in ns/ft. This is referred to as the board propagation constant.
- $S_{0\_MICROSTRIP}$ and $S_{0\_STRIPLINE}$ refer to the speed of the signal on an unloaded microstrip or stripline trace on the PCB in ns/ft.
- Z_0 is the intrinsic impedance of the line in Ω and is a function of the dielectric constant (ε_r), the line width, line height and line space from the plane(s). The equations for Z_0 are not included in this document. See the *MECL System Design Handbook* by William R. Blood, Jr. for these equations.
- C_0 is the distributed trace capacitance per unit length of the network in pF/ft.
- L_0 is the distributed trace inductance per unit length of the network in nH/ft.
- C_D is the sum of the capacitance of all devices and stubs divided by the length of the network's trunk, not including the portion connecting the end agents to the termination resistors in pF/ft.
- S_{EFF} and Z_{EFF} are the effective propagation constant and impedance of the PCB when the board is "loaded" with the components.

5.2. Effective Impedance and Tolerance/Variation

The impedance of the PCB needs to be controlled when the PCB is fabricated. The method of specifying control of the impedance needs to be determined to best suit each situation. Using stripline transmission lines (where the trace is between two reference planes) is likely to give better results than microstrip (where the trace is on an external layer using an adjacent plane for reference with solder mask and air on the other side of the trace). This is in part due to the difficulty of precise control of the dielectric constant of the solder mask, and the difficulty in limiting the plated thickness of microstrip conductors, which can substantially increase crosstalk.

- The effective line impedance (Z_{EFF}) is recommended to be 65 Ω +/-10%, where Z_{EFF} is defined by Equation 11.

5.3. Power/Reference Planes, PCB Stackup, and High Frequency Decoupling

5.3.1. POWER DISTRIBUTION

Designs using the Pentium III Xeon processor require several different voltages. The following paragraphs describe some of the impact of three common methods used to distribute the required voltages. Refer to the *Pentium® III Xeon™ Processor Power Distribution Guidelines* for more information on power distribution.

The most conservative method of distributing these voltages is for each of them to have a dedicated plane. If any of these planes are used as an "AC ground" reference, then the plane needs to be AC coupled to the system ground plane. This method may require more total layers in the PCB than other methods. 1 ounce/ft^2 thick copper is recommended for all power and reference planes.

A second method of power distribution is to use partial planes in the immediate area needing the power, and to place these planes on a routing layer on an as-needed basis. These planes still need to be decoupled to ground to ensure stable voltages for the components being supplied. This method has the disadvantage of reducing area that can be used to route traces. These partial planes may also change the impedance of adjacent trace layers. (For instance, the impedance calculations may have been done for a microstrip geometry, and adding a partial plane on the other side of the trace layer may turn the microstrip into a stripline.)

5.3.2. REFERENCE PLANES AND PCB STACKUP

The type and number of layers for the PCB need to be chosen to balance many requirements. Many of these requirements include:

The maximum trace resistance for AGTL+ signal paths should not exceed 2 ohms. Depending on the trace width chosen and PCB vendor's process tolerance, this may require 1 ounce/ft^2 thick copper instead of 1/2 ounce/ft^2 thickness. A higher trace resistance increases the voltage drop along the trace, which reduces the falling edge noise margin.

- Providing enough routing channels to support the minimum and maximum timing requirements of the components.

- Providing stable voltage distribution for each of the components.
- Providing uniform impedance for the Pentium® III Xeon™ processor bus and other signals as needed.
- Provide a ground plane under the principal component side of the baseboard. Preferably under both sides if active components are mounted on both sides.
- Minimizing coupling/crosstalk between the networks.
- Minimizing RF emissions.
- Maximizing PCB yield.
- Minimizing PCB cost.
- Minimizing cost to assemble PCB.

The following baseboard stackup recommendations should help reduce the amount of Simultaneous Switching Output (SSO) effects experienced.

It is recommended that baseboard stackup be arranged such that AGTL+ signal routes do not traverse multiple signal layers, as this can create discontinuities in the signal's return path. It is also recommended that each AGTL+ signal have a single reference plane for the entire route. Figure 8 shows the ideal case where a particular signal is routed entirely within the same signal layer, with a ground layer as the single reference plane.

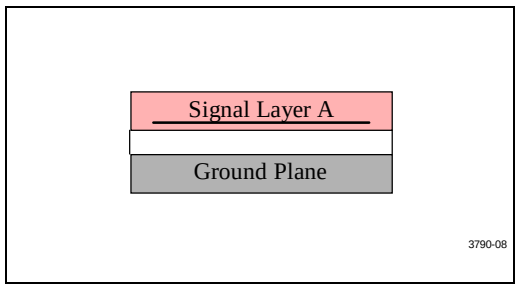

Figure 8. One Signal Layer and One Reference Plane

When it is not possible to route the entire AGTL+ signal on a single layer, there are methods to reduce the effects of layer switches whereby the signal still references the same plane (see Figure 9). Figure 10 shows another method of minimizing layer switch discontinuities, but may be less effective than Figure 9. In this case, the signal still references the same type of reference plane (ground). In such a case, it is good practice to stitch (i.e., connect) the two ground planes together with vias in the vicinity of the signal transition via.

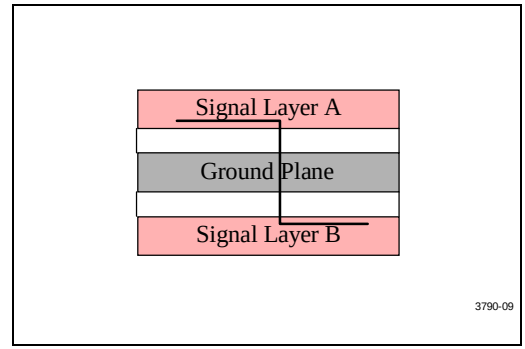

Figure 9. Layer Switch with One Reference Plane

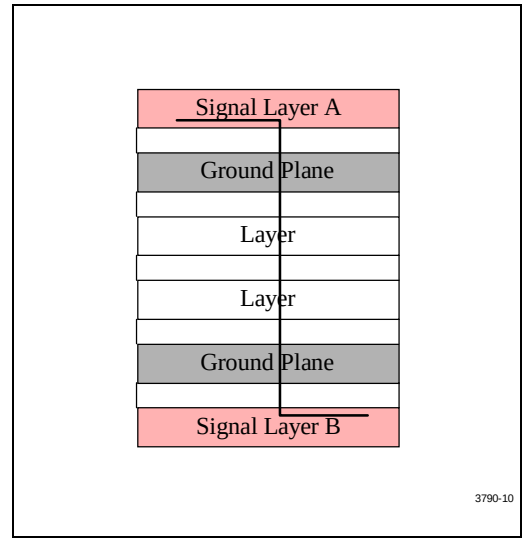

Figure 10. Layer Switch with Multiple Reference Planes (same type)

When routing and stackup constraints require that an AGTL+ signal reference multiple planes, one method of minimizing adverse effects is to add high-frequency decoupling wherever the transitions occur, as shown in Figure 11 and Figure 12. Such decoupling should, again, be in the vicinity of the signal transition via and use capacitors with minimal effective series resistance (ESR) and effective series inductance (ESL). Dual vias for these caps are recommended since via inductance may sometimes be higher than the actual capacitor inductance.

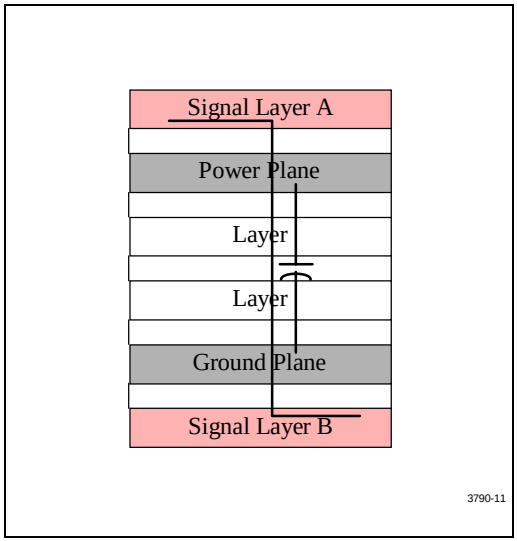

Figure 11. Layer Switch with Multiple Reference Planes

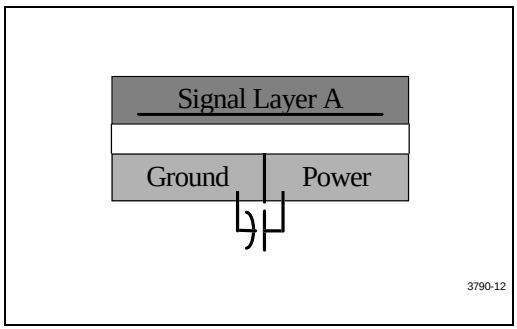

Figure 12. One Layer with Multiple Reference Planes

5.3.3. HIGH FREQUENCY DECOUPLING

This section contains several high frequency decoupling recommendations that will improve the return path for an AGTL+ signal. These design recommendations will very likely reduce the amount of SSO effects.

Just as layer switching and multiple reference planes can create discontinuities in an AGTL+ signal return path, discontinuities may also occur when a signal transitions between the baseboard and cartridge. Therefore, providing adequate high-frequency decoupling across VCC_{CORE} and ground at the SC330 connector interface on the baseboard will minimize the discontinuity in the signal's reference plane at this junction. Please note that these additional high-frequency decoupling capacitors are in addition to the high-frequency decoupling already on the processor.

Transmission line geometry also influences the return path of the reference plane. The following are decoupling recommendations that take this into consideration:

- A signal that transitions from a stripline to another stripline should have close proximity decoupling between all four reference planes.
- A signal that transitions from a stripline to a microstrip (or vice versa) should have close proximity decoupling between the three reference planes.
- A signal that transitions from a stripline or microstrip through vias or pins to a component (chipset, etc.) should have close proximity decoupling across all involved reference planes to ground for the device.

5.3.4. SC330 CONNECTOR

Internal studies indicate that the use of thermal reliefs on the connector pin layout pattern (especially ground pins) should be minimized. Such reliefs (cartwheels or wagon-wheels) increase the net ground inductance and reduce the integrity of the ground plane to which many signals are referenced. Increased ground inductance has been shown to aggravate SSO effects. Also, the anti-pad diameter (clearance holes in the planes) for the signal pins should be minimized since large anti-pads also reduce the integrity of the ground plane and increase inductance.

Some additional layout and EMI-reduction guidelines regarding the SC330 connector follow:

- Extend power/ground planes up to the SC330 connector pins.
- Extend the reference planes for AGTL+ and other controlled-impedance signals up to the SC330 connector pins.
- Minimize or remove thermal reliefs on power/ground pins.
- Route V_{TT} power with the widest signal trace or mini-plane as possible. Place decoupling caps across V_{TT} and ground in the vicinity of the connector pins.

- Use a ground plane under the principal component side of the baseboard (and secondary side if it contains active components).
- Distribute decoupling capacitors across power and ground pins evenly around the connector (less than 0.5 inch spacing) on the primary and secondary sides.
- Minimize serpentine traces on outer layers.

5.4. Clock Routing

Analog simulations are required to ensure clock net signal quality and skew is acceptable. The clock skew in Pentium III Xeon processor based systems must be kept to a minimum (the calculations and simulations for the example topologies in this document have a total clock skew of 150 ps and 150 ps of clock jitter). For a given design, the clock distribution system, including the clock components, must be evaluated to ensure these same values are valid assumptions. The *Pentium® III Xeon™ Processor at 500 and 550 MHz* datasheet specifies clock signal quality requirements for Pentium III Xeon processor systems. To help meet these specifications, follow these general guidelines:

- Tie clock driver outputs.
- Have equal electrical length and type of traces on the PCB (microstrip and stripline may have different propagation velocities).
- Maintain consistent impedance for the clock traces.
  - — Minimize the number of vias in each trace.
  - — Minimize the number of different trace layers used to route the clocks.
  - — Keep other traces away from clock traces.
- Lump the loads **at the end** of the trace if multiple components are to be supported by a single clock output.
- Have equal loads at the end of each network.

The **ideal** way to route each clock trace is on the same single inner layer, next to a ground plane, isolated from other traces, with the same total trace length, to the same type of single load, with an equal length ground trace parallel to it, and driven by a zero skew clock driver. When deviations from ideal are required, going from a single layer to a pair of layers adjacent to power/ground planes would be a good compromise. The fewer number of layers the clocks are routed on, the smaller the impedance difference between each trace is likely to be. Maintaining an equal length and parallel ground trace for the **total length of each** clock ensures a low inductance ground return and produces the minimum current path loop area. (The parallel ground trace will have lower inductance than the ground plane because of the mutual inductance of the current in the clock trace.)

5.5. Conclusion

AGTL+ routing requires a significant amount of effort. Planning ahead and leaving the necessary time available for correctly designing a board layout will provide the designer with the best chance of avoiding the more difficult task of debugging inconsistent failures caused by poor signal integrity. Intel recommends planning a layout schedule that allows time for each of the tasks outlined in this document.

APPENDIX A
DEFINITION OF V_{REF} GUARDBAND, AND FLIGHT TIME MEASUREMENTS/CORRECTIONS

Acceptable signal quality must be maintained over all operating conditions to ensure reliable operation. Signal Quality is defined by four parameters: Overshoot, Undershoot, Settling Limit, and Ringback which are specified in the *Pentium® III Xeon™ Processor at 500 and 550 MHz* datasheet and *Intel® 450NX PCIset*. Timings are measured at the pins of the driver and receiver, while signal integrity is observed at the receiver chip pad. When signal integrity at the pad violates the following guidelines and adjustments need to be made to flight time, the adjusted flight time obtained at the chip pad can be assumed to have been observed at the package pin, usually with a small timing error penalty.

6.0. V_{REF} GUARDBAND

To account for noise sources that may affect the way an AGTL+ signal becomes valid at a receiver, V_{REF} is shifted by ΔV_{REF} for measuring minimum and maximum flight times. The V_{REF} Guardband region is bounded by V_{REF}-ΔV_{REF} and V_{REF}+ΔV_{REF}. ΔV_{REF} has a value of 100 mV, which accounts for the following noise sources:

- 50 mV for motherboard coupling
- 35 mV for V_{TT} noise
- 15 mV for V_{REF} noise

The example topology covered in this guideline assumes ringback tolerance within 20 mV of V_{REF}. Since V_{REF} is guardbanded by 100 mV, this places the absolute ringback limits at:

- 1.12 V for rising edge ringback
- 0.88 V for falling edge ringback

A violation of these ringback limits requires flight time correction as documented in the *Pentium® II Processor Developer's Manual*.

7.0. OVERDRIVE REGION

The overdrive region is the voltage range, at a receiver, from V_{REF} to V_{REF} + 200 mV for a low-to-high going signal and V_{REF} to V_{REF} - 200 mV for a high-to-low going signal. The overdrive regions encompass the V_{REF} Guardband. So, when V_{REF} is shifted by ΔV_{REF} for timing measurements, the overdrive region **does not** shift by ΔV_{REF}. Figure 13 below depicts this relationship. Corrections for edge rate and ringback are documented in the *Pentium® II Processor Developer's Manual*. However, there is an exception to the documented correction method. The *Pentium® II Processor Developer's Manual* states that extrapolations should be made from the last crossing of the overdrive region back to V_{REF}. Simulations performed on this topology should extrapolate back to the appropriate V_{REF} Guardband boundary, and not V_{REF}. So, for maximum rising edge correction, extrapolate back to V_{REF} + ΔV_{REF}. For maximum falling edge corrections, extrapolate back to V_{REF} - ΔV_{REF}.

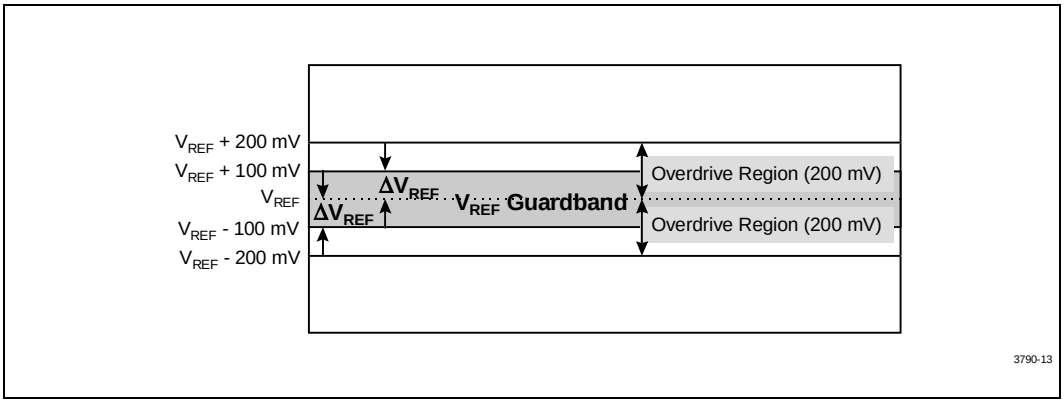

Figure 13. Overdrive Region and V_{REF} Guardband

8.0. FLIGHT TIME DEFINITION AND MEASUREMENT

The minimum flight time for a **rising edge** is measured from the time the driver crosses V_{REF} when terminated to a test load, to the time when the signal first crosses V_{REF}-ΔV_{REF} at the receiver (see Figure 14). Maximum flight time is measured to the point where the signal first crosses V_{REF}+ΔV_{REF}, assuming that ringback, edge rate, and monotonicity criteria are met. Similarly, minimum flight time measurements for a falling edge are taken at the V_{REF}+ΔV_{REF} crossing; maximum flight time is taken at the V_{REF}-ΔV_{REF} crossing.

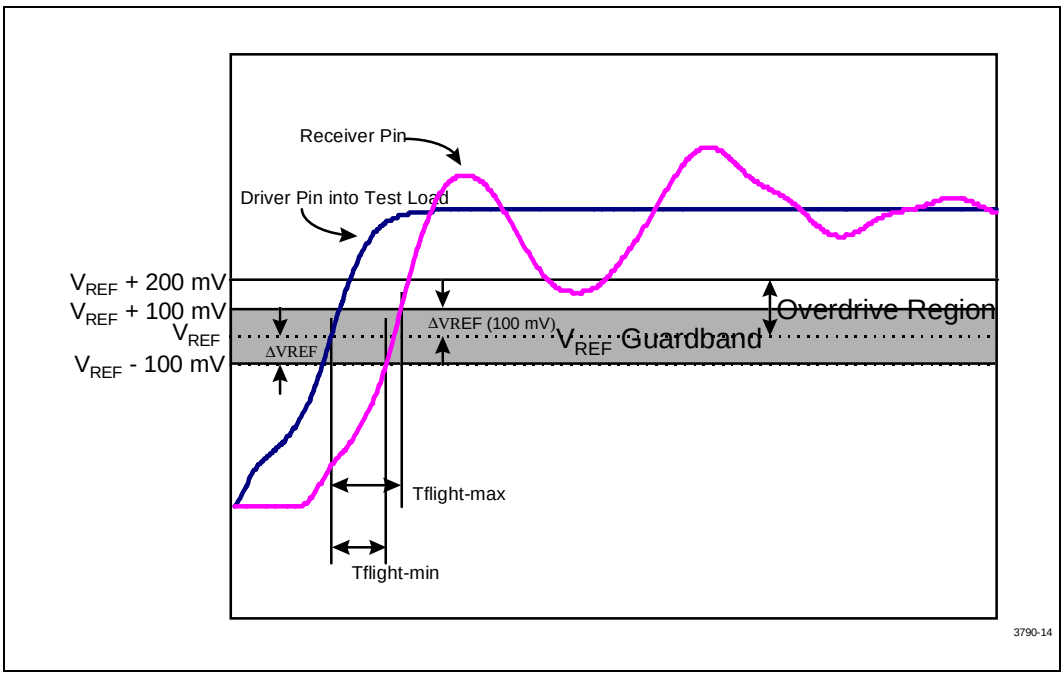

Figure 14. Rising Edge Flight Time Definition