



AP-912

**APPLICATION
NOTE**

**Pentium® III Xeon™
Processor Power
Distribution Guidelines**

March 1999

Order Number: 245095-001

Information in this document is provided in connection with Intel products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications.

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

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

The Pentium® III Xeon™ processor may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

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

Copies of documents which have an ordering number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725 or by visiting Intel's website at <http://www.intel.com>

Copyright © Intel Corporation 1999.

* Third-party brands and names are the property of their respective owners.

CONTENTS

PAGE	PAGE
1. 0.INTRODUCTION	5
1.1. Terminology	5
1.2. References.....	5
2. 0.TYPICAL POWER DISTRIBUTION	6
3. 0.SC330 PROCESSOR POWER REQUIREMENTS	8
3.1. Voltage Tolerance.....	9
3.2. Multiple Voltages.....	9
3.3. Voltage Sequencing.....	9
3.3.1. NON-AGTL+ SIGNALS (2.5 V)	11
3.3.2. BACK SIDE BUS.....	11
3.3.3. SM BUS SIGNALS	11
3.3.4. AGTL+ SIGNALS	11
3.3.5. MEMORY SIDE SIGNALS	11
3.3.6. PCI SIDE SIGNALS	12
3.3.7. CLOCK INPUT	12
3.3.8. CLOCK RATIO INPUTS.....	12
3.3.8.1. Mixed Frequency Processors.....	14
4. 0.MEETING THE SC330 PROCESSOR POWER REQUIREMENTS	14
4.1. Voltage Budgeting.....	14
4.2. Supplying Power	14
4.2.1. LOCAL DC-TO-DC CONVERTERS VS. CENTRALIZED POWER SUPPLY	15
4.2.2. AC VS. DC INPUT VOLTAGE	15
4.2.3. LINEAR REGULATORS VS. SWITCHING REGULATORS.....	16
4.3. Decoupling Technologies and Transient Response.....	16
4.3.1. BULK CAPACITANCE	19
4.3.2. HIGH FREQUENCY DECOUPLING.....	21
4.4. Power Planes.....	22
4.4.1. LOCATION OF HIGH FREQUENCY DECOUPLING.....	22
4.4.2. LOCATION OF BULK DECOUPLING.....	24
4.4.3. IMPEDANCE AND EMISSION EFFECTS OF POWER ISLANDS.....	24
5. 0.THE AGTL+ BUS POWER REQUIREMENTS	24
5.1. Tolerance	25
5.2. Reference Voltage	25
6.0. MEETING THE AGTL+ POWER REQUIREMENTS	25
6.1. Generating V_{TT}	25
6.2. Generating V_{REF}	25
7. 0.RECOMMENDATIONS	26
7.1. VCC_CORE and VCC_L2.....	26
7.2. VOLTAGE REGULATOR MODULE 8.3 (VRM 8.3).....	28
7.2.1. VRM 8.3 DESIGN.....	28
7.2.2. RECOMMENDED SYSTEM DESIGN FOR VRM 8.3	29
7.2.3. OPTIONS FOR MEETING VCC TOLERANCES	30
7.3. The Main Power Supply	32
7.4. V_{TT32}	
7.4.1. TERMINATION RESISTORS.....	32
7.5. V_{REF}	33
7.6. Component Models	33
8. 0.MEASURING TRANSIENTS	34
9.0. SC330 PROCESSOR POWER DISTRIBUTION NETWORK MODELING	34
FIGURES	
Figure 1. Ideal SC330 Processor Power Supply Scheme.....	6
Figure 2. Power Distribution of SC330 Processors	7
Figure 3. Detailed Power Distribution Model for System with One SC330 Processor	7
Figure 4. AGTL+ Bus Termination Layout	8
Figure 5. SC330 Processor Substrate and L2 Cache.....	9
Figure 6. Voltage Sequencing Example.....	10
Figure 7. Non-AGTL+ ESD Diodes.....	10

Figure 8. AGTL+ ESD Diodes 11

Figure 9. Timing Diagram of Compatibility Pins 12

Figure 10. Schematic of Pin Sharing..... 13

Figure 11. System Design Model 15

Figure 12. Remote Sense..... 15

Figure 13. Location of Capacitance in a Power Model with a DC-to-DC Converter 17

Figure 14. Effect of Transients on a Power Supply 18

Figure 15. ESR Cylindrical Capacitor..... 19

Figure 16. Capacitor Model 19

Figure 17. ESR Required for Various Current Demands 20

Figure 18. Capacitance Required vs. ESR at 8.5A, 60 mV and 30 μ s Δ t 20

Figure 19. 12 Layer PCB Stackup..... 23

Figure 20. 1206 Capacitor Pad and Via Layouts..... 23

Figure 21. VREF..... 25

Figure 22. Local Regulation..... 26

Figure 23. Connections to Via Patterns..... 29

Figure 24. VRM Layout Option 1 30

Figure 25. VRM Layout Option 2 30

Figure 26. VRM Layout Option 3 31

Figure 27. VRM 8.3 Placement in a 4 Processor System..... 31

Figure 28. VRM 8.2 Placement in a 4 Processor System 32

Figure 29. VCC_CORE Power Delivery Model for AC Transient Response..... 35

Figure 30. VCC_L2 Power Delivery Model for AC Transient Response..... 35

Figure 31. VTT Power Delivery Model for AC Transient Response..... 36

TABLES

Table 1. SC330 Power Delivery Models..... 8

Table 2. DC Voltage Tolerance 9

Table 3. Timing Parameters of Compatibility Pins..... 13

Table 4. DC and AC Voltage Tolerance at 400 MHz 14

Table 5. Efficiency of a Linear Regulator..... 16

Table 6. Estimating VTT Current..... 25

Table 7. Core and L2 Voltage Identification Code 27

Table 8. VRM 8.3 Pinout..... 28

Table 9. Voltage Specifications at SC330 Pins ... 29

Table 10. Various Component Models Used at Intel (Not Vendor Specifications)..... 34

1.0. INTRODUCTION

As computer performance demands increase, new, higher speed logic with increased density is developed to fulfill these needs. To reduce their overall power dissipation, modern microprocessors are being designed with lower voltage implementations. This in turn requires power supplies to provide lower voltages with higher current capability. Because of this, processor power is now becoming a significant portion of the system design, and demands special attention. Now more than ever, power distribution requires careful design practices. Pentium® III Xeon™ processors have unique requirements for voltages supplied to them. Their bus implementation, called AGTL+ (Assisted Gunning Transceiver Logic+), requires a voltage supply of its own.

The 100 MHz SC330 processor system bus operates using GTL+ (Gunning Transceiver Logic +) signaling levels with a new type of buffer utilizing active negation and multiple terminations. This new bus logic is called **Assisted Gunning Transistor Logic**, or **AGTL+**.

For older personal computer designs, a power plane with a mix of high frequency and bulk decoupling capacitors spread evenly across the system board is a low cost way to ensure sufficient power distribution. As the switching rate of current increases, the cost of the power distribution system becomes significant enough to merit careful calculation. Centralized distribution of power, for example, may no longer be the most cost effective solution to power distribution.

Another side effect of lowering voltages of some components is the existence of multiple voltages within the system. On a basic Pentium III Xeon processor-based system board there will be 1.5 V for AGTL+ termination, 1.8 V to 2.8 V for the processor and L2 cache, 2.5 V for CMOS non-AGTL+ signals, 3.3 V for the chipset and SM Bus, 5 V and 12 V for other components. The sequencing possibilities of all these voltages must be taken into account. This is discussed in Section 3.3.

The reader should be familiar with basic electrical engineering theory, as the first few sections of this document explain in detail the issues involved in designing a system with proper power distribution. The

last few sections discuss the recommendations and SC330 processor power distribution network modeling.

1.1. Terminology

“Power-Good” or “PWRGOOD” (an active high signal) indicates that all of the supplies and clocks within the system are stable. PWRGOOD should go active a constant time after system voltages are stable and should go inactive as soon as any of these voltages fail their specifications. The time constant should be set such that, in a working system, all clocks and other supply levels have reached a stable condition before PWRGOOD goes active.

VCC_CORE, VCC_L2 and AGTL+ refer to the processor core’s V_{CC}, SC330 processor’s cache supply voltage and Assisted Gunning Transceiver Logic + supply voltage, respectively. “AGTL+” is the bus between the SC330 processor and it’s chipset. The AGTL+ bus and the front side bus are therefore synonymous. Further, in this document, the terminology SC330 processor and Pentium III Xeon processor has been used to refer to the same processor.

1.2. References

This document contains references to the

- AP-587, *Slot 1 Processor Power Distribution Guidelines* (Order Number 243332-002)
- AP-830, *Pentium® II Xeon™ Processor/Intel® 450NX PCIsset AGTL+ Layout Guidelines* (Order Number 243790-001)
- *Pentium® II Xeon™ Processor at 400 MHz datasheet* (Order Number 243770-001)
- *Pentium® II Xeon™ Processor Bus Termination Design Guidelines* (Order Number 243774-001)
- *Pentium® III Xeon™ Processor at 500 and 550 MHz datasheet* (Order Number 245094-001)
- *VRM 8.2 DC-DC Converter Design Guidelines* (Order Number 243773-002)
- *VRM 8.3 DC-DC Converter Design Guidelines* (Order Number 243870-002)
- *SC330 Processor Bus Terminator Design Guidelines*

2.0. TYPICAL POWER DISTRIBUTION

Power distribution is generally thought of as **getting power to the parts that need it**. Most digital designers typically begin by assuming that an ideal supply will be provided, and plan their schematics with little thought to power distribution until the end. The printed circuit board (PCB) designers attempt to create the ideal supply with two power planes in the PCB or by using large width traces to distribute power. High frequency noise created when logic gates switch is controlled with high frequency ceramic capacitors, which are subsequently recharged from lower frequency bulk capacitors (such as tantalum capacitors). Various **rule of thumb** methods

exist for determining the amount of each type of capacitance that is required. For SC330 processor designs, the system designer needs to design beyond the rule of thumb and architect a power distribution system that meets SC330 processor specs.

Figure 1 shows the ideal power model. However, a more realistic power distribution scheme appears as in Figure 2.

Note that, in a SC330 system, there is a capability to use up to four processors. In Figure 2, one can see the recommended solution involving local voltage regulator modules.

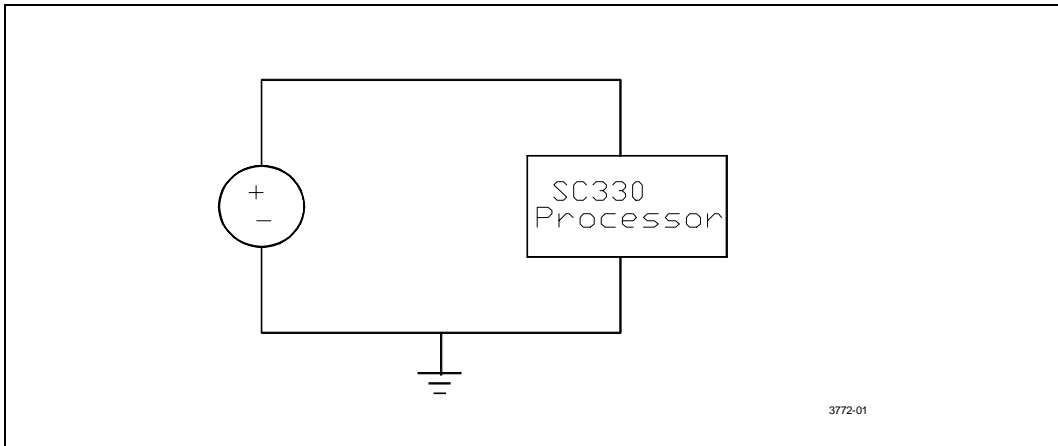


Figure 1. Ideal SC330 Processor Power Supply Scheme

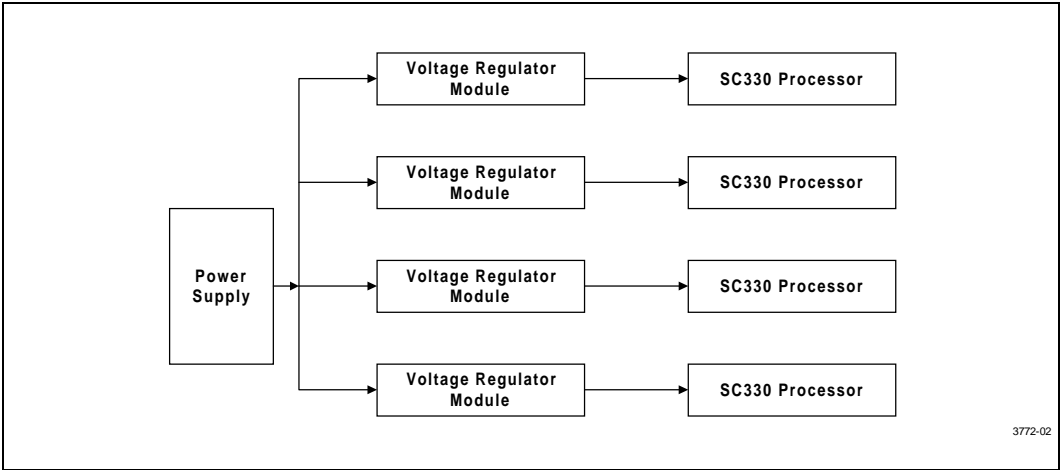


Figure 2. Power Distribution of SC330 Processors

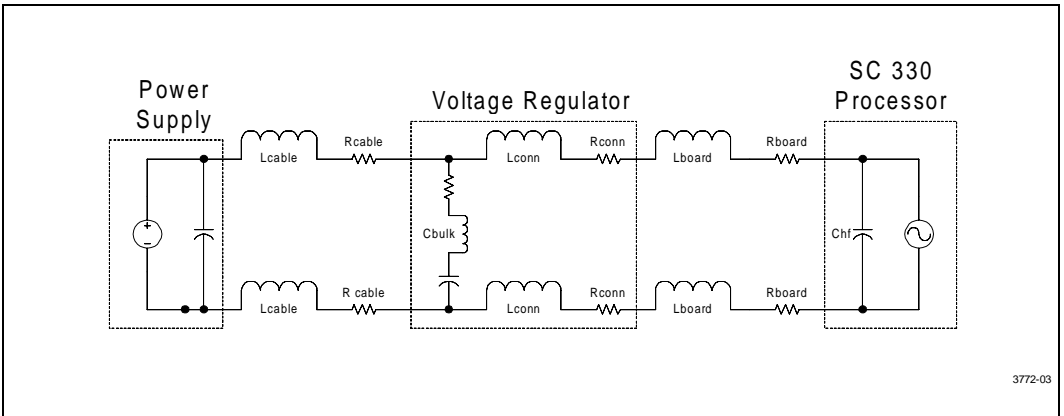


Figure 3. Detailed Power Distribution Model for System with One SC330 Processor

To completely model this system, one must include the inductance and resistance which exists in the cables, connectors, PCB, the pins and body of components (such as resistors and capacitors), and the edge fingers and contacts of the processor and voltage regulator. A more detailed model showing these effects is shown in Figure 3.

In the past, voltage drops due to inductance ($V = Ldi/dt$) and resistance ($V = IR$) have been nearly negligible relative to the tolerance of components in most systems.

This has caused the creation of simple rules for decoupling. For example, with the current at 1 amp, and a $\pm 5\%$ tolerance on 5 V (± 250 mV), one could easily ignore the effects of 25 m Ω of resistance in the distribution path, since this amounts to only 25 mV of drop. However, at 10 amps, this IR drop is equal to the 250 mV tolerance of the supply. Similarly, 3nH of inductance can typically be ignored in a power distribution system, unless current transients of 30 amp/ μ s exist, as seen in SC330 processors. The Ldi/dt drop in this case is equal to 90 mV.

Table 1. SC330 Power Delivery Models

Supply	Regulator Capacitance	Regulator Inductance	Regulator Resistance	Motherboard Inductance	Motherboard Resistance
VCC_CORE	8000 μ F	2.0 nH	6.6 m Ω	0.2 nH (HOT) + 0.2 nH (RTN)	0.6 m Ω (HOT) + 0.1 m Ω (RTN)
VCC_L2	8000 μ F	2.0 nH	6.6 m Ω	0.2 nH (HOT) + 0.2 nH (RTN)	1.5 m Ω (HOT) + 0.5 m Ω (RTN)
VTT	8000 μ F	2.0 nH	25 m Ω	10 nH (HOT) + 5.0 nH (RTN)	10 m Ω (HOT) + 1.0 m Ω (RTN)

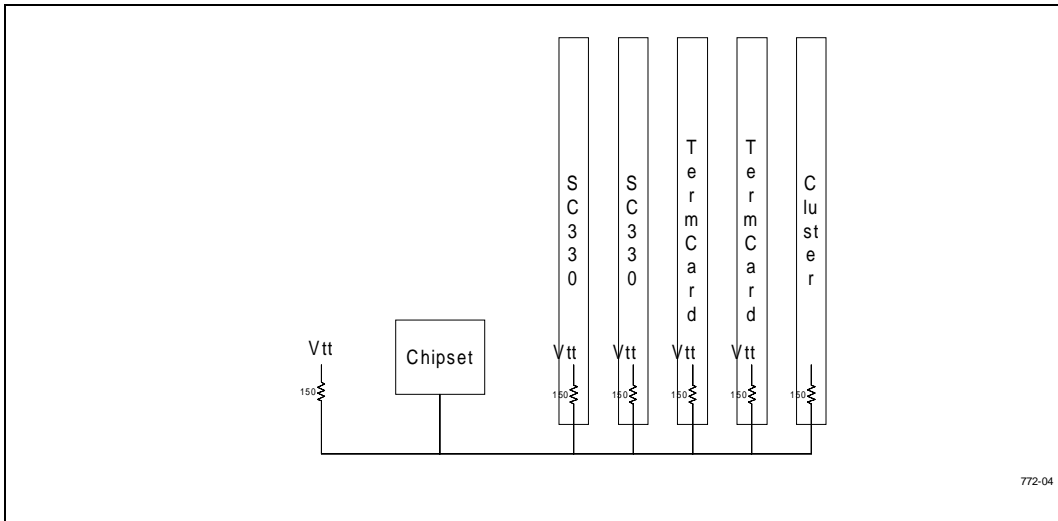


Figure 4. AGTL+ Bus Termination Layout

The high I_{CC} and di/dt requirements of a SC330 processor must both be taken into account for a successful design. Section 3 describes the requirements of the SC330 processor. Section 4 discusses meeting these requirements.

3.0. SC330 PROCESSOR POWER REQUIREMENTS

This section describes the issues related to supplying power to a SC330 processor. For detailed electrical specifications, please refer to *Pentium® III Xeon™ Processor at 500 and 550 MHz* datasheet.

The maximum power of a SC330 processor is specified by the maximum power of the substrate, not by the maximum current specification of each voltage source. This is due to the fact that all components can not possibly be running at maximum power simultaneously.

Pentium III Xeon processor typically operate on 2.0 V and depending on the frequency of operation will draw currents up to 14 A. However, Flexible Mother Board (FMB) designers should keep in mind that I_{ccCORE} requirements for FMB are as high as 16 A. The Pentium III Xeon processor allows the use of Auto HALT, Stop-Grant and Sleep states to reduce power consumption by stopping the clock to specific internal sections of the processor, depending on each particular state. There is no Deep Sleep state on the Pentium III Xeon processor.

A recommended termination for the AGTL+ bus is shown in Figure 4. All AGTL+ bus lines should be terminated to the V_{TT} supply, through 150 ohm resistors. Any unused SC330 slots must be terminated using Front Side Bus Termination cards. This bus implementation allows up to 6 loads and may be run at speeds up to 100 MHz. *Pentium® II Xeon™ Processor Bus Termination*

Design Guidelines document describes termination card design requirements.

3.1. Voltage Tolerance

The SC330 processor requires a ± 85 mV DC ‘steady-state’ tolerance and a ± 130 mV AC tolerance at the gold fingers.

Table 2. DC Voltage Tolerance

Voltage	DC Tolerance
VCC_L2	± 125 mV
V _{TT}	± 135 mV
VCC_TAP	± 125 mV
VCC_SMBUS	± 165 mV

Failure to meet these specifications on the low-end tolerance results in transistors slowing down and not meeting timing specifications. Not meeting the specifications on the high-end tolerance can induce electro-migration, causing damage or reducing the life of the processor.

3.2. Multiple Voltages

The VRM 8.2/8.3 module, which provides VCC_CORE supply to the Pentium III Xeon processor, has the capability of supplying voltages from 1.8 V to 2.8 V.

Typical VCC_CORE is 2 V. Similarly, the VRM 8.2/8.3 module, which provides VCC_L2 supply to the caches, has the capability of supplying voltages from 1.8 V to 2.8 V. Typical VCC_L2 is 2.5 V or 2.7V (depending on cache size). Reference Table 7 and/or VRM 8.2/8.3 DC-DC Converter *Design Guidelines* document, for available voltage details. Figure 5 shows a SC330 processor and 4 L2 caches, 512 KB each, on a substrate. The Pentium III Xeon processor will be available at introduction with 512K and 1 MB L2 cache options.

Typical multiple voltages required for a Pentium III Xeon processor system are, VCC_CORE = 2 V, VCC_L2 = 2.5 V or 2.7 V, V_{TT} = 1.5 V, VCC_TAP = 2.5 V and VCC_SM = 3.3 V.

3.3. Voltage Sequencing

When designing a system with multiple voltages, there is always the issue of ensuring that no damage occurs to the system during voltage sequencing. Voltage sequencing is the timing relationship between two or more voltages, such as AGTL+ signals and VCC_CORE. Sequencing applies when the user turns on or off the power supply, or the system enters a failure condition. Sequencing applies to the power voltage levels and the levels of certain other crucial signals.




Figure 5. SC330 Processor Substrate and L2 Cache

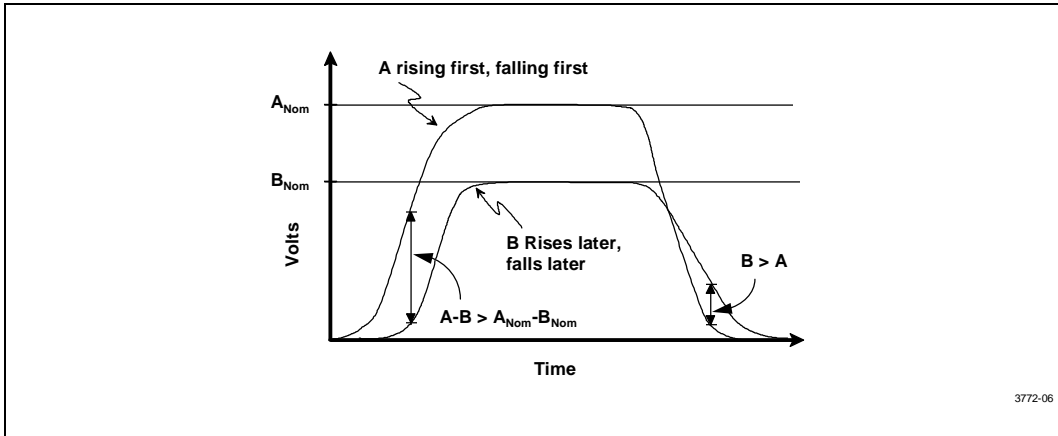


Figure 6. Voltage Sequencing Example

Figure 6 shows an example of power voltage sequencing. Here, voltage levels A and B are powered on and off in different sequence. On power-on, the A-B differential may be larger at any instant than the nominal levels. On power-off, the voltage B input may actually exceed voltage A at any instance in time. Intel designed SC330 processors, the AGTL+ bus, and Intel’s chipsets such that no additional circuitry is required in the power system to ensure the order of voltage sequencing. However, systems should be designed such that neither supply

stays on for extended time while the other is off. Excessive exposure to these conditions can compromise long term component reliability.

The following discussions describes the worst case situation where one voltage is on while the other is off. See Figure 7 and Figure 8 for highly simplified models of the buffers that show the ESD protection diodes. This model is provided for discussion purposes only and is not meant to imply any implementation scheme.

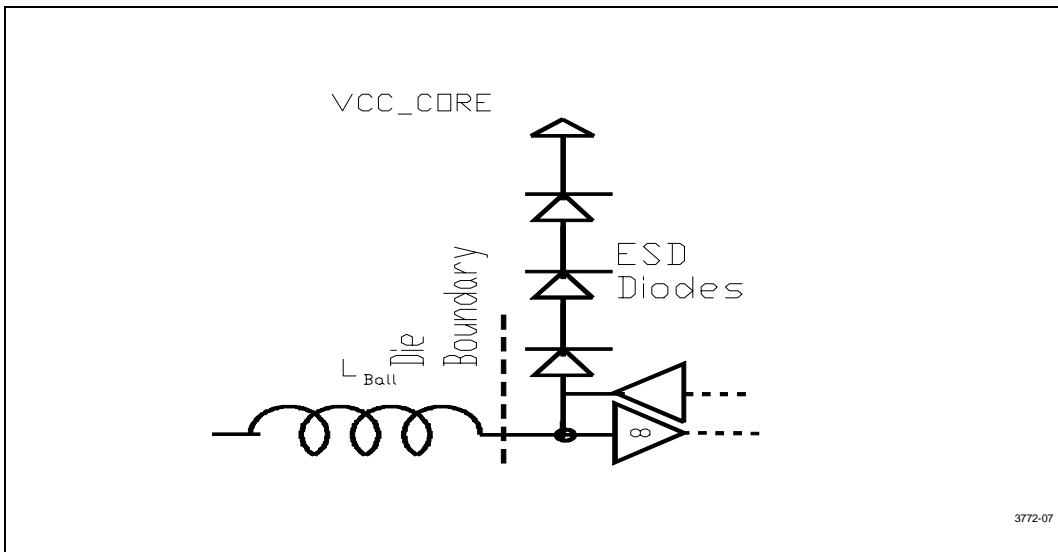
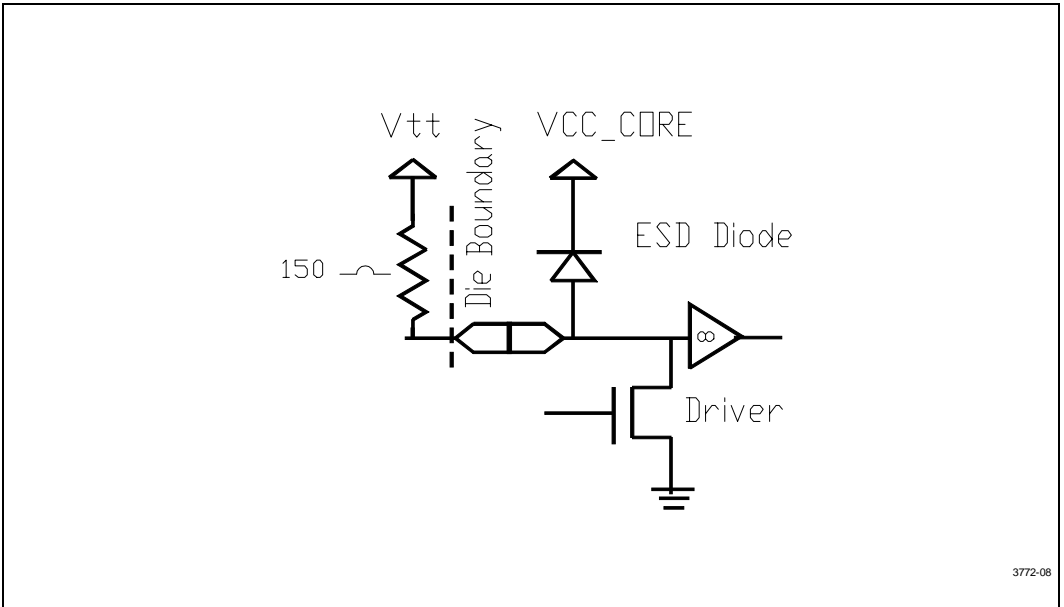


Figure 7. Non-AGTL+ ESD Diodes



3772-08

Figure 8. AGTL+ ESD Diodes

3.3.1. NON-AGTL+ SIGNALS (2.5 V)

When the VCC_CORE supply is on, and the non-AGTL+ supply (2.5 V) is off, the ESD protection diodes of the buffers are reverse biased and no power is supplied to the signal lines. As the processor sees RESET#, the outputs switch to the high or inactive state so bus contention after 2.5 V comes up is avoided. If the 2.5 V supply is on while the VCC_CORE supply is off, the 2.5 V supply delivers current to the SC330 processor core through the string of three ESD protection diodes connecting the pads to VCC_CORE . If a pull-up is used for the high level of the signals, then 100 ohms will allow a maximum of only 25 mA of current to be supplied to the core per pad cell. If the inputs are driven by a CMOS output, then the current from the output should be limited to 200 mA maximum output current per SC330 processor pin.

3.3.2. BACK SIDE BUS

No sequencing issues exist on the back side bus. See the absolute maximum ratings table in the *Pentium® III Xeon™ Processor at 500 and 550 MHz* datasheet.

3.3.3. SM BUS SIGNALS

There is no interaction with the L2 or CPU core.

3.3.4. AGTL+ SIGNALS

When the VCC_CORE supply is on and V_{TT} is off, all inputs appear low and there will be no current flowing on the AGTL+ bus. If V_{TT} is on and VCC_CORE is off, the AGTL+ bus attempts to power up the core through the ESD protection diode. The resulting VCC_CORE level will be low enough that no significant current will be consumed by the core.

NOTE

Every device on the bus must have power in order for the AGTL+ bus to operate properly.

3.3.5. MEMORY SIDE SIGNALS

3.3 V DRAM is used with 450NX chipset. There are no memory side sequencing issues, when 3.3 V DRAM is used.

By providing the MIOC with the PWRGOOD signal (as described in Section 1.1), so that it asserts MRESET# to RCGs and MUXs, it will drive the CAS lines of the DRAM inactive, and reset the data buffers as soon as it receives 3.3 V. This holds the DRAM outputs off and keeps the chipset buffer components in reset during a period of power supply stabilization. This includes a poor V_{TT} that would prevent the AGTL+ bus RESET# signal from being created correctly. This action protects these devices from producing bus contention between themselves

3.3.6. PCI SIDE SIGNALS

PCIRST# tells all PCI devices to remain in a tri-state condition. The PCI bus controller holds this signal active when the bus controller receives power and its PWR_GD signal is inactive. The PCI bus controller also tri-states its outputs during this time. In addition, the PCI inputs use a 5 V input for their ESD protection. This eliminates any issue with turning on its ESD diodes.

3.3.7. CLOCK INPUT

The clock input frequency must never exceed the intended final value while the PWRGOOD signal to the processor is active. (See terminology in Section 1.1.)

PWRGOOD should be inactive anytime that VCC_CORE or VCC_L2 are invalid. This can be accomplished by logically AND-ing ‘power good’ signals from both supplies, and connecting this output to the chipset and the SC330 processor’s PWRGOOD input for reset generation. (In this case, ‘power good’ is a signal from each supply that indicates stable voltage levels that are within tolerance.)

3.3.8. CLOCK RATIO INPUTS

The pins A20M#, IGNNE#, LINT1, and LINT0 are shared with the function for programming the PLL core clock multiplier ratio. These pins control the setting of the clock multiplier ratio during RESET# and until two clocks beyond the end of the RESET# pulse. At all other times their functionality is defined as the compatibility signals that the pins are named after. These signals have been made 2.5 V tolerant so that they may be driven by existing logic devices. This is important for both functions of the pins.

Figure 9 shows the timing relationship required for the clock ratio signals with respect to RESET# and BCLK. Table 3 shows the timing parameters. Note that the minimum setup time for these signals is 1 ms. This table also shows the timing relationship of the compatibility signals. Figure 9 also shows a signal called CRESET# (CMOS Reset), with the timing needed for controlling the multiplexing function required to share the pins. The chipset may generate CRESET#.

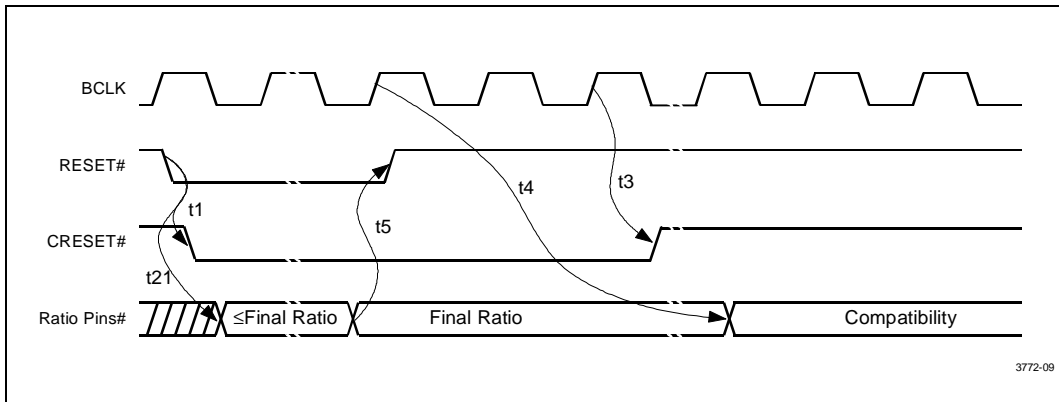


Figure 9. Timing Diagram of Compatibility Pins

Table 3. Timing Parameters of Compatibility Pins

t#	Parameter	Minimum	Maximum	Units
t1	RESET# active to CRESET#		10	ns
t2	RESET# active to Ratio Delay		5	BCLKs
t3	BCLK to CRESET# Inactive		10	ns
t4	BCLK to Compatibility		20	BCLKs
t5	Ratio Setup to RESET# rising	1		ms

Figure 10 illustrates one method for using CRESET# to select the clock multiplier ratio and the normal functionality of the compatibility signals. The pins of the processors are bussed together to allow either of them to be the compatibility processor. The outputs of the multiplexer must meet the requirements stated in the non-AGTL+ tolerant buffer specified above.

The compatibility inputs to the multiplexer must meet the input specifications of the multiplexer. This may require a level translation before the multiplexer inputs unless the inputs and the signals driving them are already compatible.

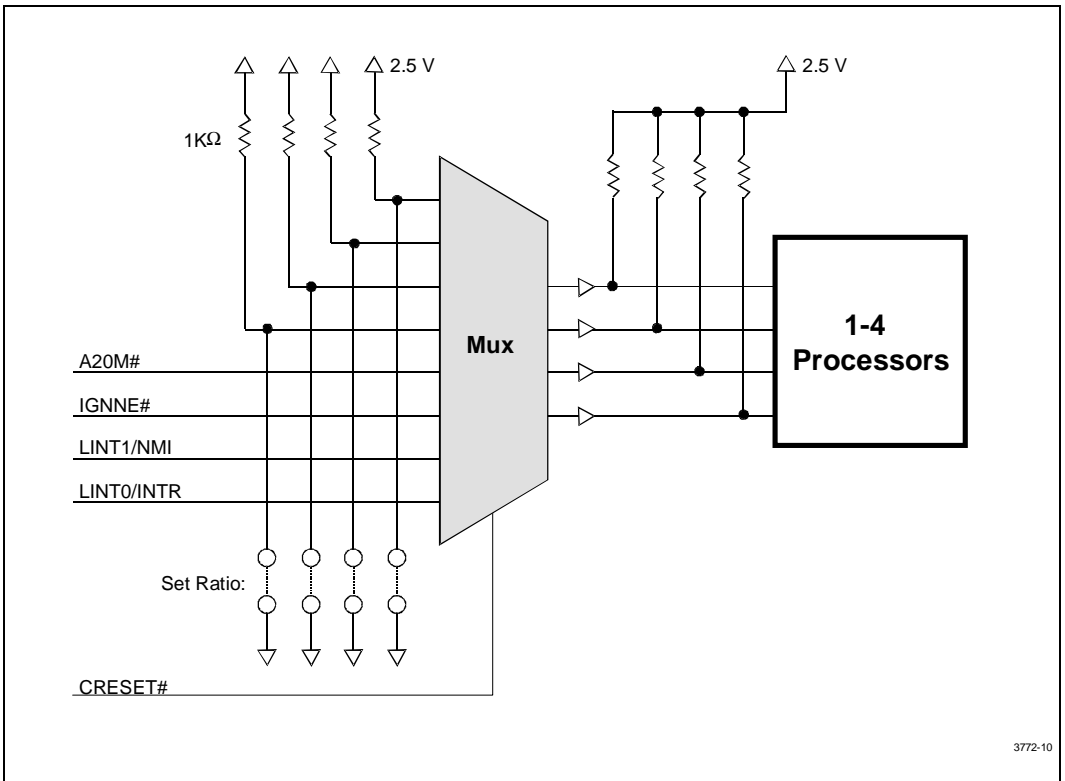


Figure 10. Schematic of Pin Sharing

3.3.8.1. Mixed Frequency Processors

Intel does not support nor validate mixing processors rated at different frequencies.

4.0. MEETING THE SC330 PROCESSOR POWER REQUIREMENTS

SC330 processor power supply design requires tradeoffs between power supply, distribution and decoupling technologies. This section discusses how to design a system using the more accurate power distribution model shown in Figure 11, one step at a time. Even though this chapter deals with different type of power supplies and different configurations, Intel recommends using local regulation (VRM 8.2/8.3 modules) for VCC_CORE and VCC_L2.

4.1. Voltage Budgeting

Before beginning the design of a power distribution system one must have an idea of how to budget the tolerance specifications at different points in the system. This provides a target for component and board layout and helps reduce iterations to reach a solution.

Table 4 provides an estimation of the effects of the factors that system designers need to consider when calculating DC and AC voltage tolerance budget. Table 4 also lists the tolerable voltage fluctuations/ losses at VRM out put and at gold fingers.

Table 4. DC and AC Voltage Tolerance at 400 MHz

Location	DC (mV)	AC (mV)
VRM out put (at 2 V)	±60	±100
At gold fingers	-85	±130

4.2. Supplying Power

The power distribution system starts with the source of power, or the power supply. A central power supply unit may create the required voltages. Another option, local regulation may create the voltages closer to the load. The section below discusses the tradeoffs involved.

Due to higher current requirements and in order to maintain power supply tolerance, the SC330 processor requires either local regulation or a power supply with remote sense capabilities. A DC loss occurs over the power distribution system due to the resistance of such things as cables, power planes, and connectors.

The formula $\Delta V = I \times R$ represents this loss. Where ΔV is the voltage loss, I is the current and R is the effective resistance of the distribution system. In a system with consistent current demand, setting the voltage slightly higher than the nominal value overcomes resistance of the distribution system. This ensures that the voltage at the farthest reaches of the system remains within specification. However, in systems where current can change significantly between a high and a low state (i.e., ΔI is high), ΔV changes significantly as well. The formula $\Delta V = \Delta I \times R$ represents this change in voltage. The tighter tolerance specification of SC330 processors make this loss significant.

Intel recommends local regulation (the use of a supply or regulator near the load) to create the voltage needed. For example, a local DC-to-DC converter, placed close to the load, converts a higher DC voltage to a lower level using either a linear or a switching regulator. Distributing lower current at a higher voltage to the converter minimizes unwanted losses ($I \times R$). (Power companies use this same method in high tension lines to distribute electricity from the generating source to local residential use.) More importantly however, a discrete regulator regulates the voltage locally which minimizes DC line losses by eliminating R_{CABLE} and reducing R_{BOARD} on the processor voltage.

Power supplies with remote sense may work if local regulation is not appropriate. A power supply typically regulates the voltage at its terminals before cabling to the board. Again, changing distribution losses based on the current demand make it difficult to hold a tight tolerance at the load. A remote sense, shown in Figure 12, may solve this problem by running a separate connection from near the load to the feedback loop of the power supply. The feedback loop has very low current draw (in the μ amp range) and does not suffer from the line losses described above. This allows the supply to regulate its output based on the voltage level at the load that is affected by the line losses. Remote-sense supplies suffer from added inductance due to cabling to a power supply and noise induced in the remote sense feedback signal. Section 4.3 explains this issue. The system designer must also deal with finding a representative load point that applies for all processors in a multi-processor system.




Figure 11. System Design Model

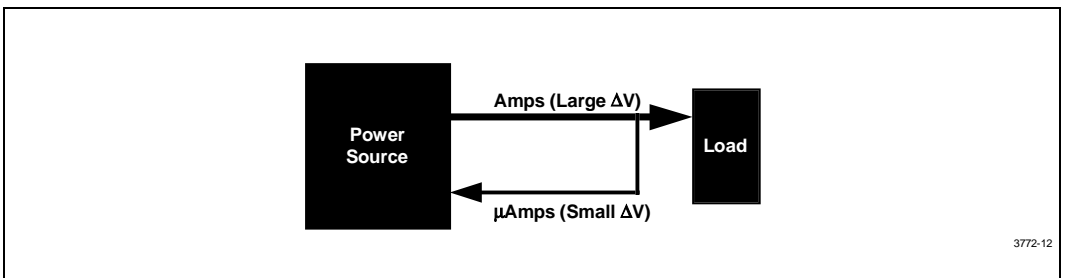


Figure 12. Remote Sense

Either method of regulation can easily maintain $\pm 2\%$ accuracy, plus a small ripple and noise budget, under a **stable** load. However, the further demands of a SC330 processor tax the abilities of a remote-sense supply. The large current transients of a SC330 processor means that the system designer must exercise extreme care to eliminate noise coupling and ringing when using remote sense feedback.

4.2.1. LOCAL DC-TO-DC CONVERTERS VS. CENTRALIZED POWER SUPPLY

Most desktop computers today utilize a self-contained multiple output power supply. This convenient and cost effective strategy isolates the issues of power generation from the system designer and allows the creation of a large reusable sub-system. However, lower operating voltages and increased transient response make long bus distribution schemes and self-contained supplies less suitable due to the resistance and inductance of the

distribution scheme. The use of distributed local DC-to-DC converters provides another alternative.

Distributed local DC-to-DC converters improve upgrade potential. Sockets allow these converters to be added or replaced as required. Furthermore, self-adjusting regulators meet the varying needs of the processor socket.

While the decision lies in the hands of the system designer, Intel recommends the use of local regulators. Refer to *VRM 8.2/8.3 DC-DC Converter Design Guidelines* for more details on such regulators.

4.2.2. AC VS. DC INPUT VOLTAGE

The new SC330 processor DC voltage can be created directly from the line voltage or from a low voltage AC or DC tap of the central power supply.

Creating a DC voltage from an AC voltage is generally **easier** than converting from one DC level to another. A DC-to-DC voltage converter must first **chop** a DC voltage in order to create an alternating voltage before the converter can step that voltage up or down. Typically however, PC power supplies today do not provide AC voltage taps to the system.

Creating the additional DC voltage from the line voltage requires the addition of an extra winding to the line transformer. This incurs additional costs and suffers from issues of distribution explained in the next sections. Changing the output voltage in this system requires changing the transformer, which makes the design less versatile.

Creating the additional DC voltage from an existing DC voltage requires a DC-to-DC converter. These converters work well in the PC market as they can work off of the existing 5 V or 12 V taps of typical PC power supplies, and can be manufactured in high volumes. System designers can place DC-to-DC converters very near the SC330 processor (thus reducing distribution loss) or design them into the existing power supply case. They can also design DC-to-DC converters to have selectable output voltages, as well.

4.2.3. LINEAR REGULATORS VS. SWITCHING REGULATORS

A linear regulator drops a variable voltage across itself in order to maintain an output voltage within tolerance regardless of load changes (within its specifications). Due to their simplicity, linear regulators respond to load changes fairly quickly (about 1µs response time). A linear regulator's efficiency drops off as the input/output voltage differential increases, as evidenced in Equation 1.

Equation 1. Loss Within a Linear Regulator

$$P_{LOSS} \approx (V_{IN} - V_{OUT}) \times I$$

The linear regulator also requires a minimum drop from the input to the output of about a diode drop (0.5 V to 1.0 V), making it impossible to have small changes from V_{IN} to V_{OUT} .

The formula **efficiency** = V_{OUT}/V_{IN} approximates the efficiency of a linear regulator. Table 5 illustrates the significant power loss and poor efficiency of a linear regulator for a V_{IN} of 5 V and a fixed output current of 10 amps.

Linear regulators tend to have faster reaction times than switching regulators. However, due to the high power loss of the linear regulator, designers should consider switching regulators for the high output current ratings required by SC330 processors. A 2.8 V switching regulator can achieve 80% efficiency at 10 amps.

Table 5. Efficiency of a Linear Regulator

V _{OUT}	Efficiency with V _{IN} of 5 V	Power Loss at 10 Amps
3.3	66%	17W
2.8	56%	22W
2.5	50%	25W
1.8	36%	32W

A switching regulator first **chops** the input voltage to make it **AC-like**. The faster it switches or chops, the faster the converter's reaction time. A faster reaction time reduces capacitance requirements. Low end switching regulators operate at a 100 kHz switching rate, while high end devices start at 1 MHz.

4.3. Decoupling Technologies and Transient Response

As shown earlier, inductance is also an issue in distribution of power. The inductance of the system due to cables and power planes further slows the power supply's ability to respond quickly to a current transient.

Decoupling a power plane can be broken into several independent parts. Figure 13 shows each of the locations where capacitance could theoretically be applied. The closer to the load the capacitor is placed, the more inductance that is bypassed. By bypassing the inductance of leads, power planes, etc., less capacitance is required. However, closer to the load there is less room for capacitance. Therefore tradeoffs must be made.

Typically a digital component causes switching transients. These sharp surges of current occur at each clock edge and taper off by the end of the cycle. Intel designed the SC330 processor such that it manages the highest frequency components of the current transients. Intel accomplished this by adding capacitance to the cartridge (C_{HF}) as well as directly on the die (C_{DIE}). To lower SC330 connector and substrate inductance (L_{CONN} and L_{SUB}) as well as the board inductance (L_{BOARD}), the

SC330 processor is designed with approximately 56 ground pins and 35 VCC_CORE pins. These processor design considerations reduce the current slew rate to the order of 20A/ μ s. SC330 processors require no external high frequency capacitance, since C_{HF} is sufficient to lower the di/dt to 20 A/ μ s. Larger bulk storage (C_{BULK}),

such as electrolytic capacitors, supply current during longer lasting changes in current demand by the component, such as coming out of an idle condition. Similarly, they act as a storage well for current when entering an idle condition from a running condition.

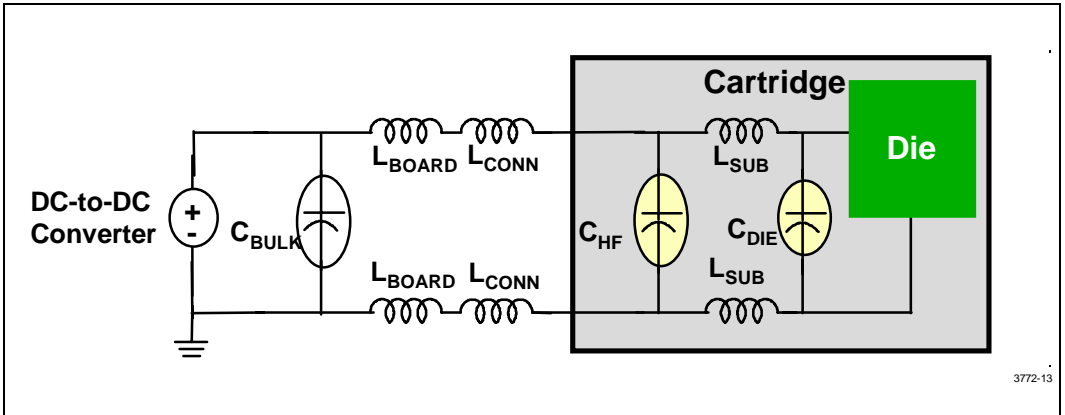


Figure 13. Location of Capacitance in a Power Model with a DC-to-DC Converter




Figure 14. Effect of Transients on a Power Supply

All of this power bypassing is required due to the relatively slow speed at which a power supply or DC-to-DC converter can react. A typical voltage converter has a reaction time on the order of 1 to 100 μ s while the processor's current transients are on the order of 1 to 20 ns. Bulk capacitance supplies energy from the time the high frequency decoupling capacitors are drained until the power supply can react to the demand. More correctly, the bulk capacitors in the system slow the transient requirement seen by the power source to a rate that it is able to supply, while the high frequency capacitors slow the transient requirement seen by the bulk capacitors to a rate that they can supply. Figure 14 shows a poorly controlled supply versus a well-controlled supply during an increase in current demand. Notice how the poorly controlled supply dips below the allowed tolerance specification. A similar situation exists as the current demand decreases.

A load-change transient occurs when coming out of or entering a low power mode. For the SC330 processor this load-change transient can be on the order of 13 amps. These are not only quick changes in current demand, but also long lasting average current requirements. This occurs when the STPCLK# pin is asserted or de-asserted and during AutoHALT. AutoHALT is a low power state that the processor enters when the HALT op-code is executed. Note that even during normal operation the current demand can still change by as much as 7 amps as activity levels change within the SC330 processor component.

Maintaining voltage tolerance, during these changes in current, requires high-density bulk capacitors with low Effective Series Resistance (ESR). Use thorough analysis when choosing these components.

4.3.1. BULK CAPACITANCE

To understand why just adding more capacitance is not always effective, one must consider the ESR of the capacitance being added. This is the inherent resistance of the capacitor plate material. One way to understand where ESR comes from, and how to recognize a low ESR capacitor, is to analyze a cylindrical capacitor. By

unrolling the metal of the capacitor it appears as a sheet. This sheet has some linear resistance in Ω /inch. A longer sheet (bigger diameter capacitor) increases ESR. A wider sheet (taller capacitor) decreases the ESR.




Figure 15. ESR Cylindrical Capacitor

Another effect is the fairly high inductance of the bulk capacitors. These elements can be modeled as shown in Figure 16.




Figure 16. Capacitor Model

Again, this was taken from the complex model of Figure 3. Overcoming ESR is discussed here while assuming for now that the inductance effect will be addressed by the high frequency decoupling capacitors discussed in Section 4.3.2.

Figure 17 shows the relationship between current delivered (with a 60 mV budget) and the ESR of the capacitors. Even with infinite capacitance, 6 m Ω of ESR at 10A drops the full budget of 60 mV as shown in Equation 2.

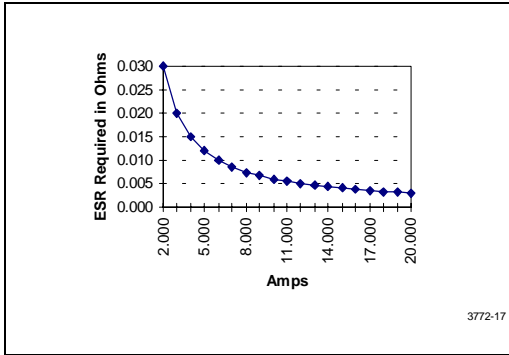


Figure 17. ESR Required for Various Current Demands

Equation 2. ESR Allowed for 60 mV Budget

$$R = 60mV / I$$

Another useful formula for estimating the amount of bulk capacitance required is shown in Equation 3. This ignores the ESR of the component but furnishes the amount of capacitance that would be required from an ideal component.

Equation 3. Capacitance for an Ideal Capacitor

$$C = \frac{\Delta I}{\Delta v / \Delta t}$$

ΔI represents the current that the bulk capacitance must be able to deliver or sink. This is equal to the difference between high and low current states since the power supply will initially continue to supply the same current that it had been prior to the load change. Δv is the allowable voltage change budgeted for bulk capacitive sag (discharge) over the period Δt . Δt is the reaction time of the power source.

Assuming some representative numbers for I , ΔV , and Δt , the capacitance required is shown by Equation 4.

Equation 4. Capacitance Needed if ESR is 0 Ohms

$$C = \frac{8.5A}{0.060V / 30 \times 10^{-6} s} = 4250 \mu F$$

Combining the above formulae to remove the resistive drop from the budget for the bulk capacitance gives Equation 5.

Equation 5. Capacitance vs. ESR

$$C = \frac{I \times \Delta t}{\Delta V - I \times ESR}$$

This equation leads to the capacitance vs. ESR graph shown in Figure 18, when ΔV is assumed to be 60 mV, I is assumed to be 8.5A, and the reaction time (Δt) of the power source is 30 μs . Any capacitance value falling in the area below the curve, in Figure 18, is sufficient for this application. Again this provides a figure that can be used to get a feel for the type of capacitors required. For example, to satisfy this equation one could use twelve 1000 μF capacitors if the ESR of each was 53 m Ω . The parallel resistance of 12 capacitors would be 4.4 m Ω and the parallel capacitance would be 12,000 μF , which falls in the area above the curve, in Figure 18.

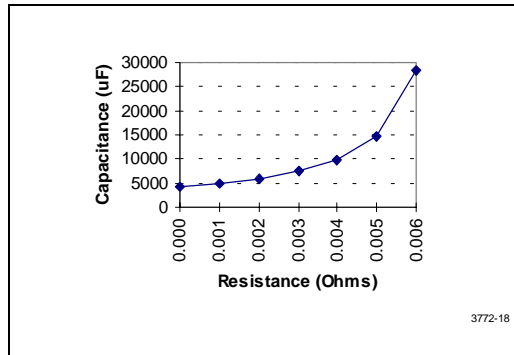


Figure 18. Capacitance Required vs. ESR at 8.5A, 60 mV and 30 μs Δt

This is a fairly conservative analysis. Using a reaction time for a power source assumes that the power source does not compensate at all for the change in current demand until Δt has passed, and then immediately is capable of delivering to that demand. Also, it is unnecessarily conservative to assume that the IR drop is the full drop the whole period in which the capacitor discharges as the current drops as the capacitor discharges. To analyze the power distribution system in more detail requires running a simulation from the power source model to the SC330 processor power model, including all board, cable, and capacitor effects. See Section 7.5 for more information on component models and Section 9 for the SC330 processor power model.

Note that the numbers used in various calculations, in this section (4.3.1), are assumed values and do not represent DS2P actual values.

4.3.2. HIGH FREQUENCY DECOUPLING

SC330 processors contain all of the capacitors necessary for high frequency decoupling of a properly designed system. This section discusses high frequency decoupling for background purposes only.

Since the bulk storage not only contains an effective series resistance, but also a fairly high inductance, these capacitors need to be assisted by other capacitors that have a lower inductance (but typically less capacitance). These **high frequency** capacitors control the switching transients and hold-over the power planes during an average load change until the higher inductance capacitors can react.

The 1206 surface mount package is a fairly low inductance package, and is actually lower than the inductance of an 0603 package due to the geometry of the board interconnects. For even lower inductance one can use a 0612 package since the board interconnect area gets even larger. A 0612 is the same size as the 1206 but has its pads along the long edge. However, due to the complexity of mass producing them, the cost of these is significantly higher. The 1206 package capacitors on the other hand are readily available and low cost.

One difficulty in simulating with high frequency capacitors however, is that vendors do not readily offer a specification for the inductance of their parts. In Section 7.5 are some measured values from capacitors that Intel has investigated which should be verified against the vendors' parts that will actually be used in any design. After calculating the number of capacitors required, one can look at the impact that averaging tolerances over many measured components has to the

design and pad the design appropriately with additional components.

Since the capacitor inductance is package related, choose the largest value available in the package that has been chosen. The highest capacitance obtainable will be the most beneficial for the design since the amount of capacitance behind this inductance is still critical.

This simple law of inductance is useful as an example for estimating the number of high frequency capacitors required:

Equation 6. Simple Law of Inductance

$$V = L di/dt$$

V is the voltage drop that will be seen due to the inductance. The di/dt value can be expressed in A/ μ s and L is the inductance of a series combination of via, trace, and all of the high frequency capacitors in parallel. See Section 4.4 for ideas on reducing via and trace inductance.

Once the allowable inductance for the budgeted voltage drop (due to high frequency transitions) is calculated, the number of capacitors (N) required can be estimated by:

Equation 7. Number of Capacitors Required

$$N = L_n / L$$

where L_n is the inductance of a single capacitor and L is the inductance required that was calculated above.

For example, to meet a 0.3 A/ns di/dt and not produce more than 60 mV of noise due to high frequency capacitor inductance (1.9 nH from Table 8) one would simply plug into Equation 6 and Equation 7.

Equation 8. Inductance Allowed

$$L = 0.060V \div 0.3A/ns = 0.2nH$$

Equation 9. Number of Capacitors for 0.2 nH

$$N = 1.9nH \div 0.2nH = 10\text{capacitors}$$

The above analysis can also include resistance of the high frequency capacitors.

While the above calculation provides a theoretical number of capacitors required to meet a di/dt requirement, high frequency noise may yet persist. More capacitors may be necessary to control noise from other sources. However, mixing additional values in the design to create higher resonance points should not be useful since the capacitors described (1206 package) have very high resonant frequencies already. This is shown by using the values from Table 8 in Equation 10.

Equation 10. Resonant Frequency

$$f = \frac{1}{2\pi\sqrt{LC}} \approx \frac{1}{2\pi\sqrt{(0.47 \times 10^{-9}) \times (1 \times 10^{-6})}} \approx 7.3\text{MHz}$$

Note that all 1206 capacitors will have basically the same inductance value and that smaller components actually have more inductance. Also, the inductance of the vias are the larger contributors and actually cause the resonance to be more like 3.6 MHz.

4.4. Power Planes

The imperfections of the power planes themselves have so far been ignored. These may also introduce unwanted resistance and inductance into the power distribution system. The complex model in Figure 3 refers to these imperfections as R_{BOARD} and L_{BOARD} .

Multilayer boards are approximately ten times superior to double sided boards, in terms of both emissions and susceptibility.

Figure 19 shows a configuration for 12 layer board, in which a separate plane is allocated to V_{TT} . However, if cost is the main consideration, one can opt to go for fewer layers by embedding V_{TT} plane in other planes. In that situation, a thorough analysis is recommended.

Power should definitely be distributed as a plane. This plane can be constructed as an **island** on a layer used for other signals, on a supply plane with other power islands, or as a dedicated layer of the PCB. Processor power should never be distributed by traces alone.

Due to the fact that the SC330 processor voltage is unique to most system designs, a voltage island, or islands, will probably be the most cost effective means of distributing power to the processors. This island from the source of power to the load should not have any breaks, so as to minimize inductance in the plane. It should also completely surround all of the pins of the source and all of the pins of the load.

4.4.1. LOCATION OF HIGH FREQUENCY DECOUPLING

The SC330 processor contains all of the high frequency decoupling required for a properly designed system.

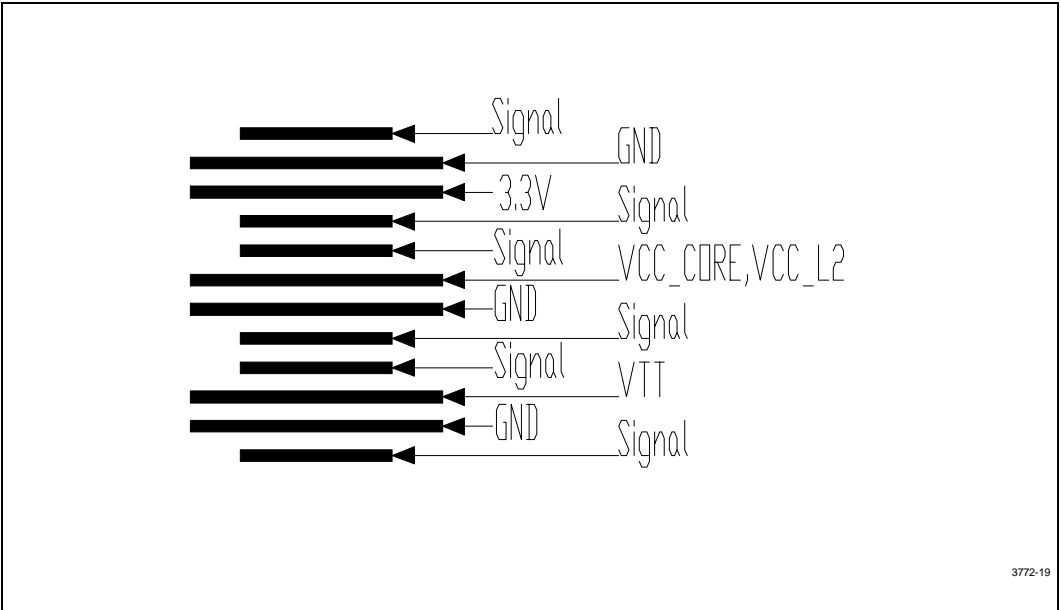


Figure 19. 12 Layer PCB Stackup




Figure 20. 1206 Capacitor Pad and Via Layouts

Where needed, high frequency decoupling should be placed as close to the power pins of the load as physically possible. Use both sides of the board if necessary for placing components in order to achieve the optimum proximity to the power pins. This is vital as the inductance of the board's metal plane layers could cancel the usefulness of these low inductance components.

Another method to lower the inductance that should be considered is to shorten the path from the capacitor pads to the pins that it is decoupling. If possible, place the vias connecting to the planes within the pad of the capacitor. If this is not possible, keep the traces as short as is feasible. Possibly one or both ends of the capacitor can be connected directly to the pin of the processor without the use of a via. Even if simulation results look good, these practical suggestions can be used to create an even better decoupling situation where they can be applied in layout. Figure 20 illustrates these concepts.

4.4.2. LOCATION OF BULK DECOUPLING

The location of bulk capacitance is not as critical since more inductance is already expected for these components. However, knowing their location and the inductance values involved will be useful for simulation. In this example the bulk capacitance is on the voltage converter module electrically **behind** the inductance of the converter pins. This is Intel's recommended solution.

4.4.3. IMPEDANCE AND EMISSION EFFECTS OF POWER ISLANDS

There are impedance consequences for signals that cross over or under the edges of the power island that exists on another layer. While neither of these may be necessary for most designs, there are two reasonable options to consider which can protect a system from these consequences.

The SC330 processor power islands can be isolated from signals by one of the solid power plane layers such as the ground layer. This forces a particular stack-up model.

Another option that helps, but does not completely eliminate radiation effects, is to decouple the edges of the processor power islands to ground on regular intervals of about 1" using good high frequency decoupling capacitors (1206 packages). This requires more

components but does not require any particular board stack-up.

In either event, for controlling emissions, all planes and islands should be well decoupled. The amount of decoupling required for controlling emission will be determined by the exact board layout, and the chassis design. One should plan ahead by allowing additional pads for capacitors to be added in case they are discovered necessary during initial EMI testing.

5.0. THE AGTL+ BUS POWER REQUIREMENTS

The AGTL+ bus is a multiple terminated, open-drain bus. Intel recommends terminating all un used SC330 slots with termination cards, as shown in Figure 4. V_{TT} (1.5 V) supplies current when output drivers turn on. There are approximately 141 AGTL+ lines in a SC330 processor system.

The AGTL+ bus power requirements present a different situation than creating power for the SC330 processors. While the AGTL+ bus requires less current than the processor, it still has a tight tolerance specification. Just as the processor can start and stop executing within a few clock cycles, the bus usage follows, which in turns causes load changes and transients on the V_{TT} power supply. V_{TT} must be available to the termination resistors at different locations, as shown in Figure 4. This can be best accomplished by having a separate plane allocated to V_{TT} , in a multilayer base board and having a dedicated 1.5 V Power Supply.

An AGTL+ buffer sinks a maximum of 55 mA. When considering the duty cycle of the signals, the 141 GTL signals draw a maximum of about 6.57 amps at 100% utilization of the bus. Table 6 illustrates AGTL+ current draw using relatively conservative duty cycles. Utilization of the bus, the value of the AGTL+ termination resistors, chipset functionality and motherboard design limit actual current draw. Power supply designers need to take these benefits into account as well. In a typical system, in addition to the front side bus, MIOC Memory bus and two F-16 buses exist. It is estimated that F-16 buses draw about 1.5A each and MIOC memory bus draws about 5A. Hence, a dedicated V_{TT} supply with around 15A rating will be sufficient.

Table 6. Estimating V_{TT} Current

Signal Group	Quantity of Signals	Max Duty Cycle	Average Current
Data + ECC	72	100	3.96
Address + Parity	35	67	1.29
Arbitration	7	100	0.39
Request	7	67	0.26
Error	5	20	0.06
Response	6	33	0.11
Other	9	100	0.50
Total	141		6.57

5.1. Tolerance

V_{TT} at the processor edge fingers must be held to $\pm 9\%$. It is recommended that a regulator that can maintain $\pm 3\%$ at low current draw be used in order to guarantee $\pm 9\%$ over all conditions. It is again important to note that this tolerance specification covers all voltage anomalies including power supply ripple, power supply tolerance, current transient response, and noise. Not meeting the specification on the low or high end will change the rise and fall time specifications. Failure to meet this specification on the low end will also result in reduced margins for the AGTL+ buffers thus making it more difficult to meet timing specifications.

5.2. Reference Voltage

The AGTL+ bus requires a Voltage Reference called V_{REF} as well. The SC330 processor generates its own copy of V_{REF} . Set V_{REF} to $2/3 V_{TT}$ for the PCIset. A simple voltage divider of two resistors can meet the V_{REF} current requirements, due to the very low current draw of this signal (at most $15 \mu A$ per device). Bear in mind that leakage current varies and may be significant when building the voltage divider.

6.0. MEETING THE AGTL+ POWER REQUIREMENTS

Due to the different nature of powering the AGTL+ bus versus powering a processor, meeting the V_{TT} requirements may be addressed in a different way.

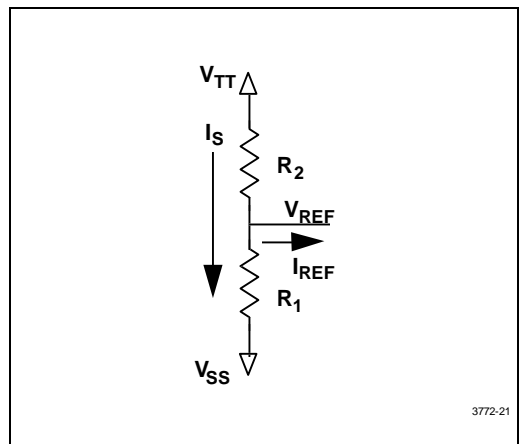
6.1. Generating V_{TT}

Intel recommends terminating AGTL+ bus as shown in Figure 4, and using a single V_{TT} regulator with a dedicated plane for its distribution. Any un used SC330 slots must be terminated by termination cards, as shown in Figure 4. For cost purposes, the mother board designers may choose not to have a dedicated V_{TT} plane. In such cases, board designers should conduct thorough analysis to optimize V_{TT} layout and distribution.

6.2. Generating V_{REF}

V_{REF} is a low current input (about $15 \mu A$ per device) to the differential receivers within each of the components on the AGTL+ bus. **Each SC330 processor generates its own V_{REF} .** A simple voltage divider can generate V_{REF} . Because V_{REF} is used only by the input buffers, it does not need to maintain a tight tolerance from component to component. It does however, need to meet the 2% specification at all V_{REF} inputs.

Equation 11 uses $R1 = 2 \times R2$ to generate a V_{REF} set at a nominal value of $2/3 V_{TT}$. Figure 21 illustrates using 1% resistors to generate the V_{REF} specification of $2/3 V_{TT} \pm 2\%$.


 Figure 21. V_{REF}

Equation 11. Creating V_{REF} of 2/3 V_{TT}

$$V_{REF} = V_{TT} \times \frac{R_1}{R_1 + R_2} = V_{TT} \times \frac{2 \times R_2}{2 \times R_2 + R_2} = \frac{2}{3} V_{TT}$$

R₁ and R₂ should be small enough values that the current drawn by the V_{REF} inputs (I_{REF}) is negligible versus the current caused by R₂ and R₁.

A complete analysis of this circuit’s currents into and out of the center node, as in Equation 12, will provide the final V_{REF} of the circuit. n is the number of I_{REF} inputs supplied by the divider.

Equation 12. Node Analysis

$$I(R_2) = I(R_1) + n \times I_{REF}$$

Plugging in for the currents and rearranging, gives:

Equation 13. Node Analysis in Terms of Voltage

$$\frac{V_{TT} - V_{REF}}{R_2} - \frac{V_{REF}}{R_1} = n \times I_{REF}$$

Which leads to:

Equation 14. Solving for V_{REF}

$$V_{REF} = \frac{V_{TT}/R_2 - n \times I_{REF}}{1/R_2 + 1/R_1}$$

The worst case V_{REF} should be analyzed with I_{REF} at the maximum and minimum values determined for the number of loads being provided voltage. If the number of loads can change from model to model or because of upgrades, this should be taken into account as well. Analyze Equation 14 with R₁ and R₂ at the extremes of their tolerance specifications.

7.0. RECOMMENDATIONS

Intel recommends using simulation to design and verify SC330 processor based systems. With the above estimates, a model of the power source, and the model of the SC330 processor provided in Section 9, system developers can begin analog modeling. Intel recommends the following as a starting point or benchmark.

7.1. VCC_CORE and VCC_L2

Intel recommends using a socketed local Voltage Regulator Module 8.2/8.3 (VRM 8.2/8.3) DC-to-DC converter, as shown in Figure 22, for both VCC_CORE and VCC_L2. This removes cable inductance from the distribution, reduces board inductance. These modules should be capable of accepting 5 signals, used to indicate the voltage required by the individual processor unit. Table 7 lists the Core and L2 Voltage Identification Codes. Refer to *VRM 8.2/8.3 DC-DC Converter Design Guidelines* document for actual specification.




Figure 22. Local Regulation

One of the benefits of using separate regulators per processor is the ability to vary processor types in the system, if allowed by the product line in the future.

Intel recommends placing the bulk decoupling on the DC-to-DC converter module. Since these capacitors tend to be large and not available in surface mount technology, it makes sense to isolate these to a smaller module that can be run in a different manufacturing environment than the typical system board designs.

The SC330 processor contains all of the high frequency decoupling required for a properly designed system.

Table 7. Core and L2 Voltage Identification Code^{1, 2}

Processor Pins							
VID4	VID3	VID2	VID1	VID0	V _{CC}	Core ³	L2 ³
00110b - 01111b					Reserved		
0	0	1	0	1	1.80	x	x
0	0	1	0	0	1.85	x	x
0	0	0	1	1	1.90	x	x
0	0	0	1	0	1.95	x	x
0	0	0	0	1	2.00	x	x
0	0	0	0	0	2.05	x	x
1	1	1	1	0	2.1	x	x
1	1	1	0	1	2.2		x
1	1	1	0	0	2.3		x
1	1	0	1	1	2.4		x
1	1	0	1	0	2.5		x
1	1	0	0	1	2.6		x
1	1	0	0	0	2.7		x
1	0	1	1	1	2.8		x
1	0	1	1	0	2.9		
1	0	1	0	1	3.0		
1	0	1	0	0	3.1		
1	0	0	1	1	3.2		
1	0	0	1	0	3.3		
1	0	0	0	1	3.4		
1	0	0	0	0	3.5		
1	1	1	1	1	no core		

NOTES:

- 0 = Processor pin connected to V_{SS}, 1 = Open on processor; may be pulled up to TTL V_{IH} on motherboard. See the *VRM 8.2/8.3 DC-DC Converter Design Guidelines*.
- VRM output should be disabled for V_{CCCORE} values less than 1.80 V.
- x** = Required

7.2. VOLTAGE REGULATOR MODULE 8.3 (VRM 8.3)

Intel has defined a VRM 8.2 derivative, called VRM 8.3. The VRM 8.3 definition will include a remote-sense input pin and will specify V_{CCCORE} at the SC330 input. VRM suppliers may elect to provide products that meet both the VRM 8.2 and VRM 8.3 requirements.

The following two sections cover the VRM 8.3 design and recommendations for designing systems with the VRM 8.3. The third section provides four layout options, first two using VRM 8.2, third using VRM 8.3 and the fourth using VRM 8.2/8.3 in a four processor configuration, for meeting V_{CC} tolerances.

7.2.1. VRM 8.3 DESIGN

VRM 8.3 design guidelines will include the expectation that, if a VRM 8.3 is installed into a VRM 8.2 header with pin B5 (remote sense input to VRM 8.3) not connected on the baseboard, the static output at the VRM pins will not exceed the V_{ID} -selected voltage + 60 mV. Table 8 shows the VRM 8.3 pinout:

Table 8. VRM 8.3 Pinout

Pin #	Row A	Row B
1	5 Vin	5 Vin
2	5 Vin	5 Vin
3	5 Vin	5 Vin
4	12 Vin	12 Vin
5	12 Vin	SENSE *
6	Ishare	OUTEN
7	VID0	VID1
8	VID2	VID3
9	VID4	PWRGOOD
10	V_{CCCORE}	V_{SS}
11	V_{SS}	V_{CCCORE}
12	V_{CCCORE}	V_{SS}
13	V_{SS}	V_{CCCORE}
14	V_{CCCORE}	V_{SS}
15	V_{SS}	V_{CCCORE}
16	V_{CCCORE}	V_{SS}
17	V_{SS}	V_{CCCORE}
18	V_{CCCORE}	V_{SS}
19	V_{SS}	V_{CCCORE}
20	V_{CCCORE}	V_{SS}

* Remote sense input to VRM 8.3 (Pin B5)

Table 9. Voltage Specifications at SC330 Pins

Symbol	Parameter	Processor Core Frequency (MHz)	Minimum	Typical	Maximum	Unit
$V_{CC_{CORE/L2}}$	V_{CC} for processor core/L2 cache			2.0		V
	$V_{CC_{CORE/L2}}$ static tolerance at VRM pins on system board	350	-0.060		0.060	V
		400	-0.060		0.060	
		450	-0.060		0.060	
	$V_{CC_{CORE/L2}}$ static tolerance at SC330 connector	350	-0.070		0.070	V
		400	-0.070		0.070	
		450	-0.070		0.070	
	$V_{CC_{CORE/L2}}$ transient tolerance at VRM pins on system board	350	-0.100		0.100	V
		400	-0.100		0.100	
		450	-0.100		0.100	

Table 9 depicts the output voltage specification at SC330 pins for frequencies up to 450 MHz (*VRM 8.3 Design Guidelines* will include frequencies above 450 MHz).


7.2.2. RECOMMENDED SYSTEM DESIGN FOR VRM 8.3

System boards can include a SENSE input trace for each VRM 8.3. The trace resistance should not be greater than 1.0 Ohm. The following information explains how to route the SENSE trace:

- Route SENSE trace from mid-point of SC330 $V_{CC_{CORE}}$ pins to core VRM pin B5.
- Route SENSE trace from mid-point of $V_{CC_{L2}}$ island (assuming two processor share a VRM) to L2 VRM pin B5. If each processor has its own L2 VRM, then route SENSE trace from the middle $V_{CC_{L2}}$ pin on the DS2P connector to L2 VRM pin B5.

To meet the transient response for resistances greater than 0.5m Ohm and inductances greater than 0.2n H, it

may be necessary to add capacitors at the connector — for example two 22uF polymer capacitors for $V_{CC_{CORE}}$ and one for the $V_{CC_{L2}}$. To avoid loss of the low ESL characteristic of the 22u F polymer capacitor, connections to via patterns should be as wide as the capacitor with multiple via holes per connection, as shown in Figure 23. In addition, some experiments have shown the need for low ESR bulk decoupling (300-600u F). However, mechanical interference with the cartridge and heatsink, in this situation, must be considered. Use these information as a guideline, as system requirements might vary.


Figure 23. Connections to Via Patterns




Figure 24. VRM Layout Option 1




Figure 25. VRM Layout Option 2

7.2.3. OPTIONS FOR MEETING V_{CC} TOLERANCES

V_{CC_{CORE}} and V_{CC_{L2}} static and transient tolerances of the Pentium III Xeon processor and the corresponding VRM tolerances assume power distribution paths with resistances no greater than 0.5m Ohm and inductances no greater than 0.2n H. Due to SC330 layout constraints, meeting these limits can be a challenge. This section introduces four layout options for system board developers to consider in order to meet the required V_{CC} tolerances; first two using VRM 8.2, third using VRM 8.3 and the fourth using VRM 8.2/8.3, in a four processor configuration.

1. It may be possible to meet the 0.5m Ohm and 0.2n H limits by placing each of the 8.2 VRMs at the side of the Pentium® III Xeon™ processor, with the VRM output pins near the SC330 V_{CC_{CORE}} or V_{CC_{L2}} input pins, as shown in Figure 24.
2. It may be possible to meet the 0.5m Ohm and 0.2n H limits by placing each of the 8.2 VRMs at the ends of the Pentium III Xeon processor closest to SC330 V_{CC_{CORE}} and V_{CC_{L2}} input pins as shown in Figure 25, using 1 plane of 2 oz copper, or 2 planes of 1 oz copper in parallel.

3. As a third option, sense the VRM output voltage on the system board at the Pentium III Xeon processor connector by using a VRM 8.3 DC-DC converter, as shown in Figure 26. Add on-board capacitors as needed.

4. Figure 27 and Figure 28 provide yet another layout option using VRM 8.3 and VRM 8.2, respectively, in a four processor configuration.

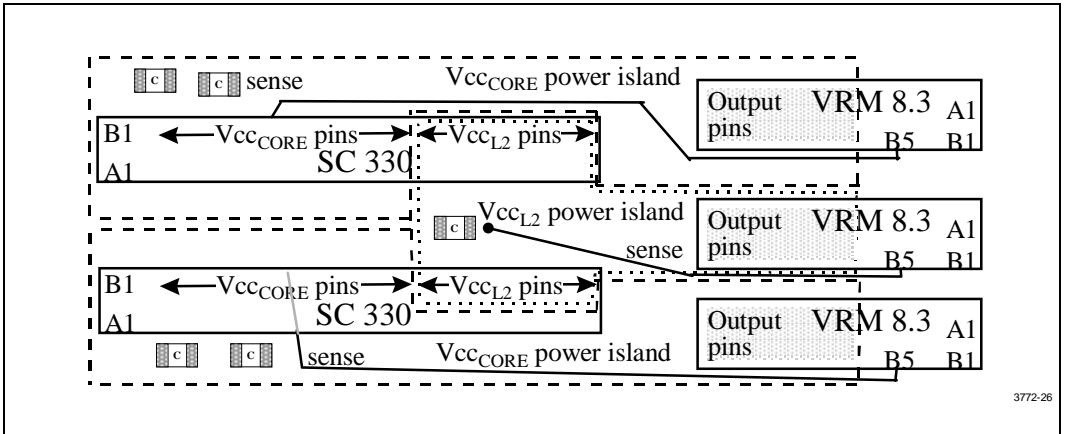


Figure 26. VRM Layout Option 3

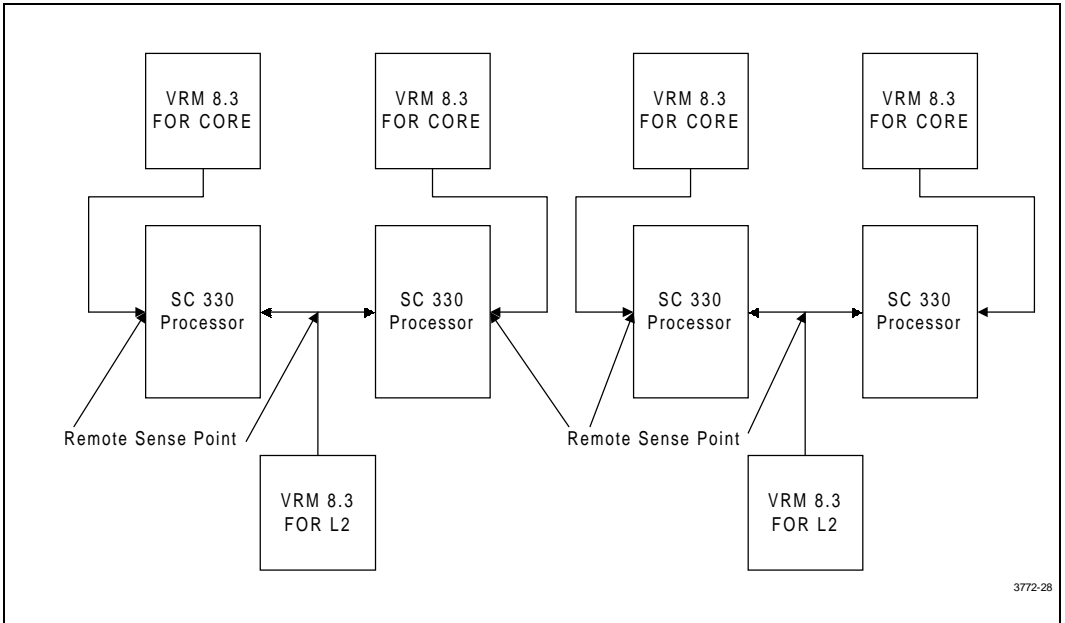


Figure 27. VRM 8.3 Placement in a 4 Processor System

3772-26

3772-28

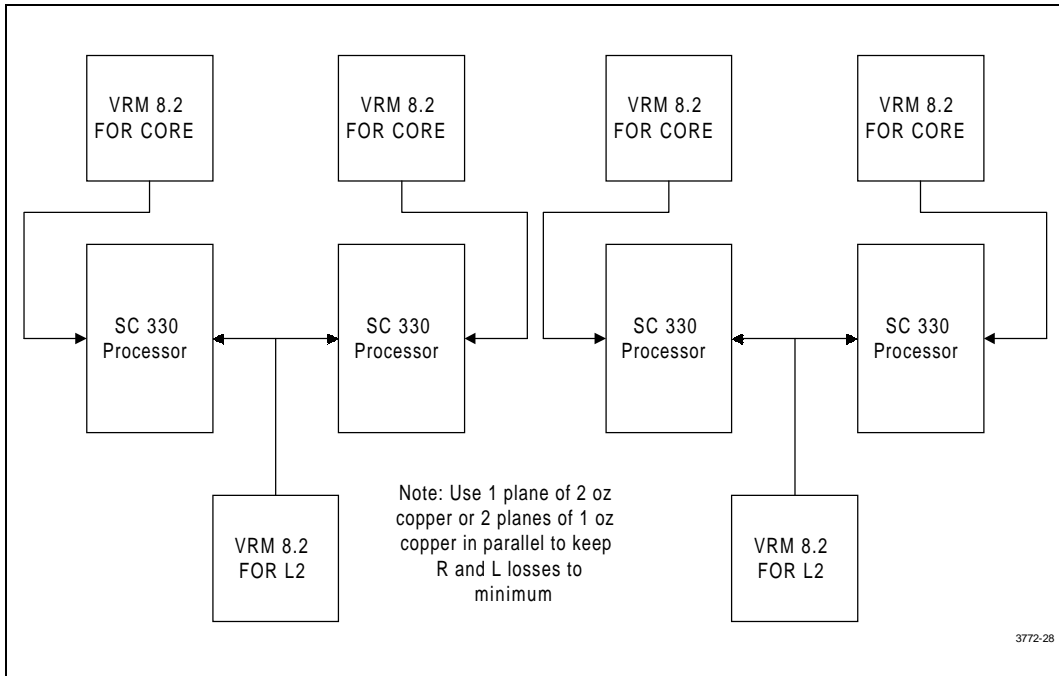


Figure 28. VRM 8.2 Placement in a 4 Processor System

7.3. The Main Power Supply

The main supply must provide power to the DC-to-DC converter as well as to the rest of the system. One should ensure that the input voltage to the converter meets the converter’s requirements, and that the DC-to-DC converter does not create a transient problem of its own on the 5 V or 12 V outputs of the main supply.

7.4. V_{TT}

Intel recommends using a single V_{TT} regulator and distributing V_{TT} on a single dedicated plane.

Intel also recommends one 0.1 μF capacitor for each termination resistor package for high frequency decoupling. The equivalent inductance values can further be lowered by using 0.1 μF and 0.001 μF , at each pin of the resistor package. However, this may be an over kill. In either case, use 0805 package. Place these capacitors as near to the termination resistors as possible.

7.4.1. TERMINATION RESISTORS

To avoid large quantities of resistors required in a system, Intel recommends using resistor network for termination resistors. The best resistor networks have a separate pin access to each side of every resistor in the package. This minimizes any inductance or crosstalk within the package.

NOTE

When using resistor networks with single corner pin V_{CC} connections for AGTL+ termination, beware of inductive packages. Intel has found that these packages can cause significant voltage drops due to the inductance in the 24 pin SOIC packages being used for this purpose. A better option is to use resistor networks in which both ends of each resistor are available as pins.

7.5. V_{REF}

Intel recommends one voltage divider at each component. The SC330 processor generates its own V_{REF} internally. Tie all of the V_{REF} inputs of each component together. Assume a maximum of 15μ amps of leakage current per load. Note that these leakage currents can be positive or negative.

The following discussion illustrates using a single voltage divider to support two V_{REF} Loads. Using 1% resistors for the voltage divider in Figure 21, make R₁ a 150Ω resistor, and use 75Ω for R₂. This creates a static usage of 7 mA (1.5 V/225Ω) per voltage divider. After looking at all combinations of R₁ and R₂ (above and below tolerance) and I_{REF} (±30 μA), the worst case solution for Equation 14 can be found with I_{REF} at 30μamps, R₁ at the low end of its tolerance specification (148.5Ω), and R₂ at the high end of its tolerance specification (75.75Ω). This yields:

Equation 15. Resistor Tolerance Analysis

$$V_{REF} = \frac{1.5/75.75 - .000030}{1/75.75 + 1/148.5} = 0.99V$$

Since the target of 2/3 of V_{TT} is 1.00 V, this setting is within 0.97% of the 2/3 point and satisfies the 2% specification. A spreadsheet program allows the reader to

easily verify the other corners. Varying over its tolerance range has minimal effect.

These values chosen for R₁ and R₂ have additional benefits: The parallel combination terminates the V_{REF} line to 50 ohms. This generally available resistance value reduces resistor cost. Further, the resistance value for R₁ (75Ω) may be generated by using two 150Ω resistors in parallel. This way, the probability of getting resistors (R₁ and R₂) with similar tolerance will be higher.

Decouple V_{REF} at each V_{REF} input and at the voltage divider with a 0.001 μF capacitor to V_{SS}. Decoupling V_{REF} to V_{TT} at the voltage dividers with a 0.001 μF capacitor may further enhance the ability for V_{REF} to track V_{TT}. The actual benefit of this decoupling is controversial.

When routing V_{REF} to the pins, use a 30-50 mil trace (the wider the better) and keep it as short as feasible. Also, keep all other signals at least 20 mils away from the V_{REF} trace. This provides a low impedance line without the cost of an additional plane or island.

7.6. Component Models

Acquire component models from their manufacturers. Intel can not guarantee the specifications of another manufacturer's components. This section contains some of the models developed by Intel for its simulations. The SC330 processor model can be found in Section 9.

Table 10. Various Component Models Used at Intel (Not Vendor Specifications)

Component of Simulation	ESR (Ω)	ESL (nH)	ESL+ Trace + Via (nH)
0.1 μ F Ceramic 0603 package	0.100	1.60	3.0
1.0 μ F Ceramic 1206 package	0.120	0.47	1.9
100 μ F MLC (2.05"x0.71")	0.005	0.30	1.7
47 μ F, 16 V Tantalum D Case	0.100	0.602	2.0
330 μ F, 16 V Aluminum Electrolytic	0.143	2.37	3.8
1000 μ F, 10 V Aluminum Electrolytic (20mm)	0.053	N/A	N/A
1000 μ F, 25 V Aluminum Electrolytic (25mm)	0.031	N/A	N/A
LBOARD. One used for V_{SS} , one for VCC_CORE. This estimate accommodates traces to vias, planes and the socket connections to the plane.	0.000	0.40	N/A

8.0. MEASURING TRANSIENTS

In order to measure transients on a voltage island, requires a clean connection. Achieve this by placing a coaxial connection directly into the power island during layout. An SMA type connector can be used and should be placed near the centrum of the voltage island.

Cable the signal directly into the oscilloscope and take the reading with the oscilloscope bandwidth limited to 20 MHz. This filters out the components of the V_{CC} noise that the processor also filter out. There is no need to decouple frequencies above this range since the SC330 processor filters them out.

9.0. SC330 PROCESSOR POWER DISTRIBUTION NETWORK MODELING

Intel provides the AC electrical models shown in Figure 23, Figure 24 and Figure 25 for use in the simulation of the AC transient response of the SC330 processor power delivery systems. Due to tool capability limitations, these models have been greatly simplified and are provided as a rough illustration of the SC330 power delivery systems.

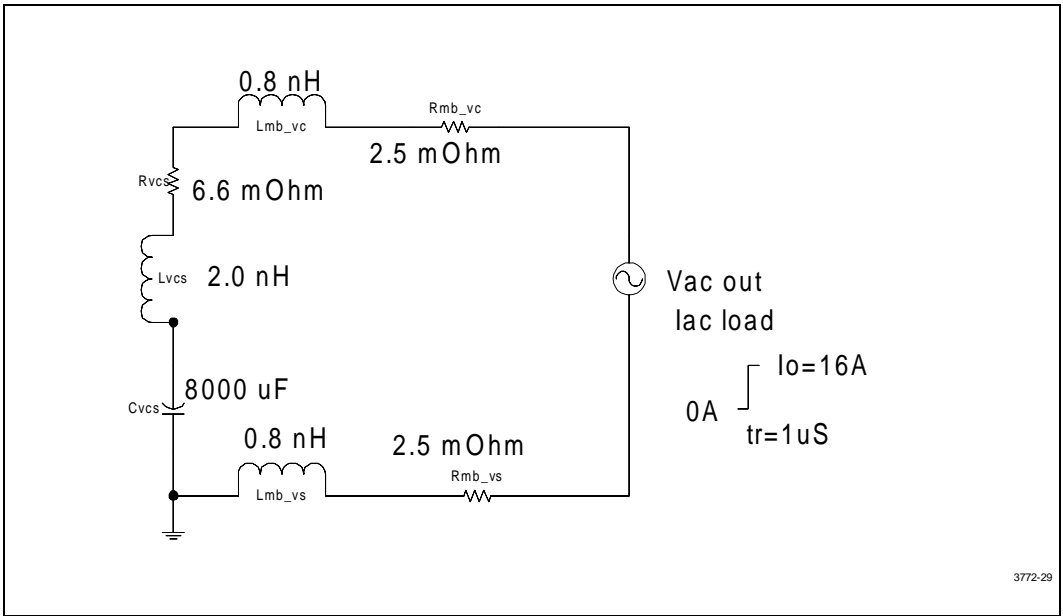


Figure 29. VCC_CORE Power Delivery Model for AC Transient Response

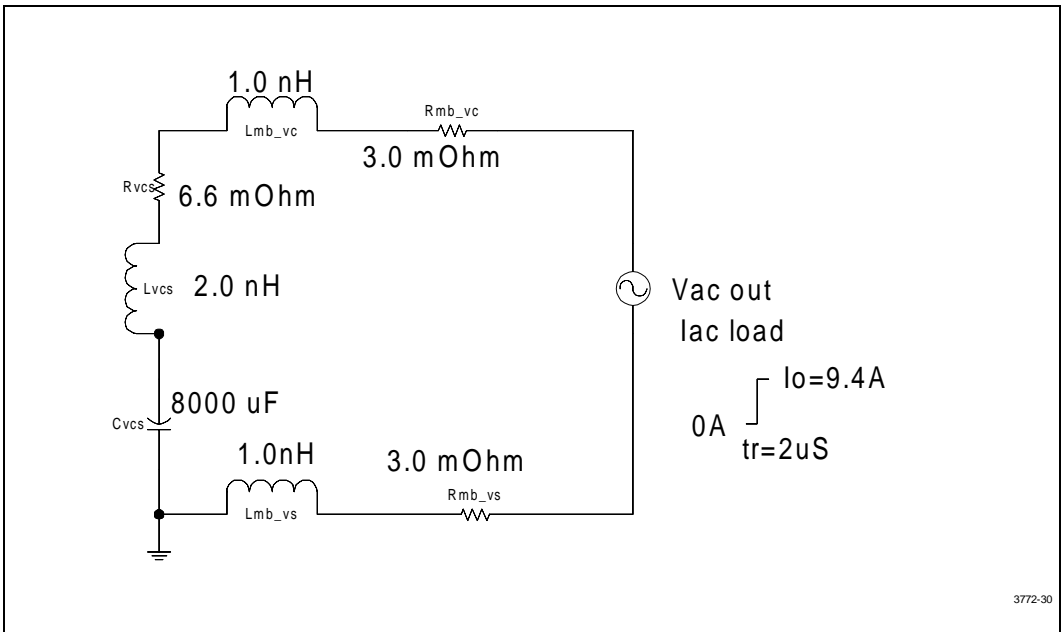


Figure 30. VCC_L2 Power Delivery Model for AC Transient Response

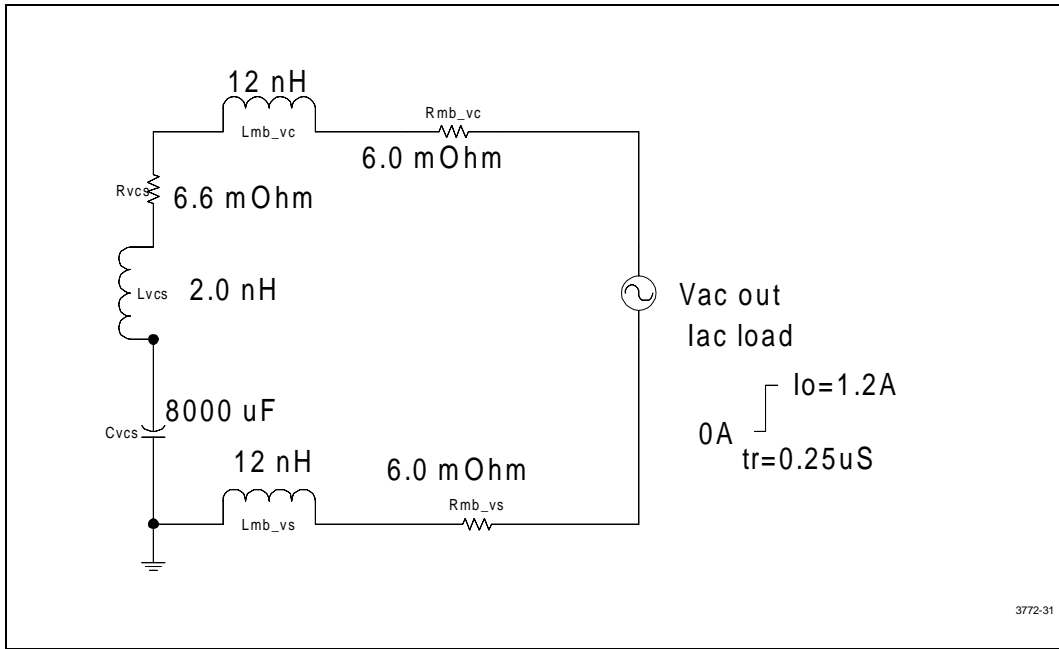


Figure 31. V_{TT} Power Delivery Model for AC Transient Response