AP-587

APPLICATION NOTE

Slot 1 Processor Power Distribution Guidelines

May 1997

Order Number: 243332-001

CONTENTS

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. Slot 1 processors have unique requirements for voltages supplied to them. Their bus implementation, called GTL+, requires a voltage supply of its own.

For most 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 current differences between the low power state and the high power state increase, 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 Slot 1 processor-based system board there will be 1.5V for GTL+ termination, 1.8V to 2.8V for the processor, 2.5V for CMOS non-GTL signals, 3.3V for the chipset and the L2 cache, and 5V for other components. The possibility that any of these voltages may come up before another 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 sections of this document explain in detail the issues involved in designing a system with proper power distribution. The last section includes a specification for a DC-to-DC converter module.

1.1. Terminology

"Power-Good" or "PWRGOOD" (an active high signal) indicates that all of the supplies and clocks within the system stabilize. PWRGOOD should go active some constant time after 5V, 3.3V and $V_{CC}P$ are stable and should go inactive any time 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" is the processor core's V_{CC}. The VCC_CORE voltage level varies for different Slot 1 processors. "VCC_L2", the Slot 1 processor's cache supply voltage, is always 3.3V.

"GTL+" is the technology used for the bus between the Slot 1 processor and its chipset. The GTL+ bus and the system bus are therefore synonymous.

1.2. References

This document contains numerous references to the *Pentium® II Processor at 233 MHz, 266 MHz and 300 MHz* datasheet (Order Number 243335).

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 in turn recharged from 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 Slot 1 processor designs, the system designer needs to reach beyond the rule of thumb and architect the power distribution system with the specifications of the Slot 1 processor in mind.

Figure 1 shows the ideal power model. However, in a real Slot 1 system, the power distribution scheme typically appears as inFigure 2. This system has physical components such as cables, connectors, the PCB, and the processor package. In this figure, one can see the recommended solution involving local voltage regulator modules.

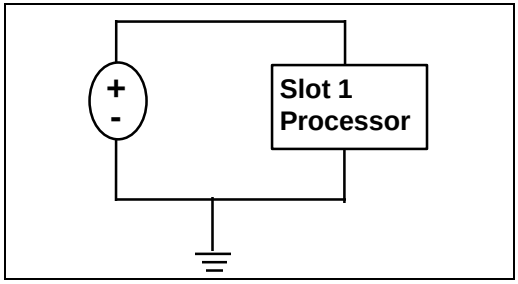

Figure 1. Ideal Slot 1 Processor Power Supply Scheme

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 the tolerance at 250 mV (5% of 5V), one could easily ignore the effects of 25 mΩ of resistance in the distribution path. However, at 10 amps, this IR drop is equal to the 250 mV tolerance. Similarly, 250 pH of inductance can typically be ignored in a power distribution system, unless current transients of 1 amp/ns exist, as they do when using Slot 1 processors. The Ldi/dt drop in this case is also equal to 250 mV.

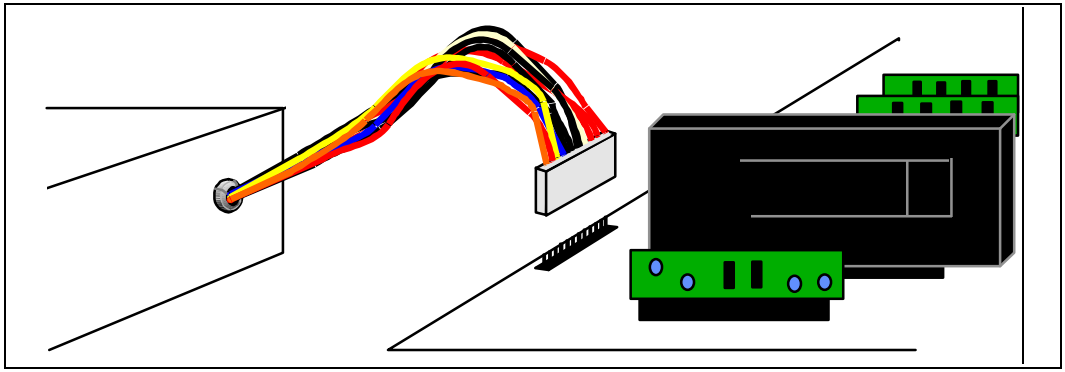

Figure 2. Physical Power Distribution of a Slot 1 System

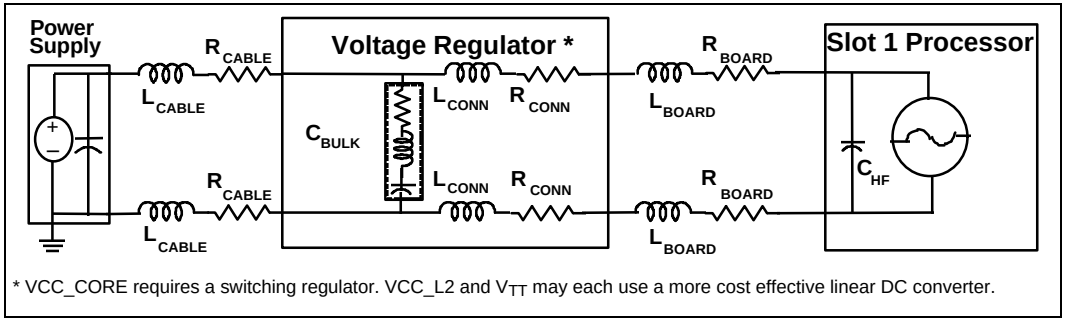

Figure 3. Detailed Power Distribution Model

Table 1. Slot 1 Power Delivery Models

Supply	Regulator Capacitance	Regulator Inductance	Regulator Resistance	Motherboard Inductance	Motherboard Resistance
VCC_CORE	6000 µF	3.0 nH	11 mΩ	0.35 nH	0.15 mΩ
Vcc3.3	200 µF	2.0 nH	25 mΩ	1.0 nH	1.0 mΩ
V_{TT}	3000 µF	2.0 nh	25 mΩ	1.0 nH	1.5 mΩ

The high value of a Slot 1 processor's current and the high rate of change of the current must both be taken into account for a successful design. Section 3 describes the requirements of the Slot 1 processor. Section 4 discusses meeting these requirements.

3.0. SLOT 1 PROCESSOR POWER REQUIREMENTS

This section describes the issues related to supplying power to a Slot 1 processor using approximate values for Pentium® II processor specifications. However, the *Pentium® II Processor at 233 MHz, 266 MHz and 300 MHz* datasheet lists the actual specifications.

The Pentium II processor operates at 2.8V, compared with 3.1V for a Pentium Pro, 3.3V for a Pentium processor, and 5V for previous Intel processors.

The Pentium II processor requires approximately three times the average current of the Pentium processor. In addition, the Pentium II processor shuts off unused units to conserve power, and includes features such as Stop Clock and AutoHALT, which create load-change transients as high as 30 amps per microsecond. In this document, a load-change transient is a change from one current requirement (averaged over many clocks) to another. Future Slot 1 processors may require higher current and different voltages for the processor die.

System developers must terminate the GTL+ bus at each end to a voltage source called V_{TT}. V_{TT} is nominally 1.5V, with a tolerance of three percent during steady state operations. This bus implementation allows up to eight loads, and may be run at speeds up to 66 MHz. Just as the processor can start and stop executing within a few clock cycles, the bus usage will follow.

The following sections discuss each of these concepts. Section 5 and Section 6 discuss the GTL+ power requirements.

3.1. Voltage Tolerance

To ease measurement of the VCC_CORE supply, Intel specifies tolerances on either side of the Slot 1 connector:

- The Slot 1 processor requires a +100/–70 mV DC 'steady-state' tolerance and a ±150 mV AC tolerance at the input (motherboard) side of the Slot 1 connector.

  or

- The Slot processor requires a +100/–90 mV DC tolerance and a ±195 mV AC tolerance at the processor's edge fingers.

The VCC_L2 voltage tolerance is ±165 mV.

The GTL+ voltage tolerance is ±135 mV.

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

3.2. Multiple Voltages

While the Pentium II processor core runs at 2.8V, the L2 cache runs at 3.3V. Future Slot 1 processors will run at different voltages. Reference Table 6 for details.

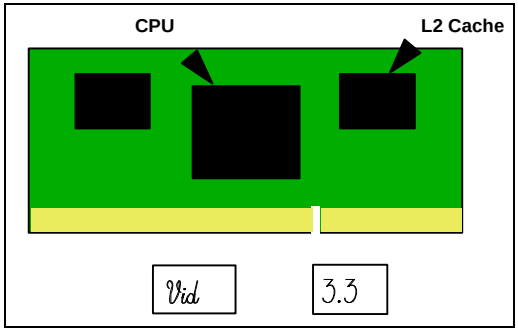

Figure 4. Slot 1 Processor Substrate

The Slot 1 pinout supports the L2 cache by including 3.3V *cache support* pins to the package. These pins are called VCC_L2 pins, while the primary voltage is supplied on the VCC_CORE pins. See Figure 5 for the location of these pins. A Slot 1 processor requires a well decoupled 3.3V supply be connected to the VCC_L2 pins. Intel expects future cache current requirements not to exceed a maximum average current (over many cycles) of 1.1A. The maximum power of a Slot 1 processor is specified as 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.

A system designer planning for upgrade potential should also be aware that future devices beyond the Pentium II processor will require a VCC_CORE other than 2.8V. This may range from 1.8V to 2.8V. To support this level of upgrade potential, the power source for the *main* processor supply should be designed with the ability to be easily configured to provide a voltage within the range of 1.8V to 2.8V. Intel recommends using the processor's Voltage ID pins with a resistor tree or a digital-to-analog converter (DAC). The voltage ID scheme is described fully in the *Pentium® II Processor at 233 MHz, 266 MHz and 300 MHz* datasheet. Intel has worked with power supply vendors to create replaceable voltage regulators which support voltage selection. See your local field applications engineer for assistance.

The GTL+ bus also requires another voltage called V_{TT} (1.5V). Section 5 discusses this voltage in further detail.

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 GTL+ 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.

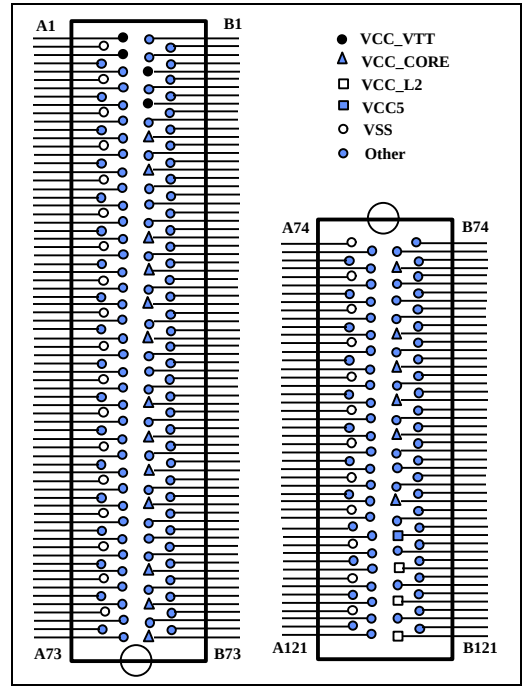

Figure 5. Slot 1 Power Pins, Top View (Through the SLot 1 Connector)

Figure 6 shows an example of power voltage sequencing. Here voltage levels A and B are shown to trade places with each other. On power-on, A-B may be larger at any point than they will be once they reach their nominal levels. On power-off, the voltage B input may actually be higher than the voltage A input for some period of time. Intel designed Slot 1 processors, the GTL+ 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 permanently while the other is off. Excessive exposure to these conditions can compromise long term component reliability.

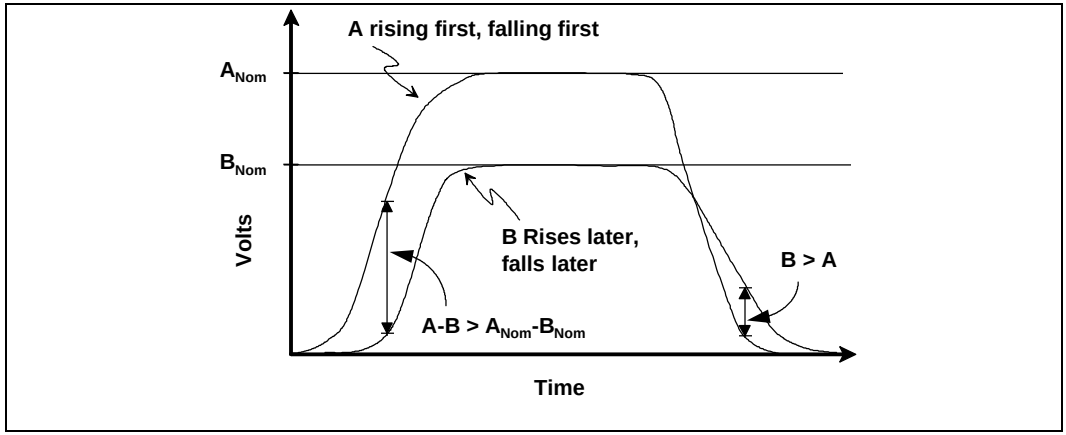

Figure 6. Voltage Sequencing Example

The discussion following is simplified by assuming the worst case which is 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.

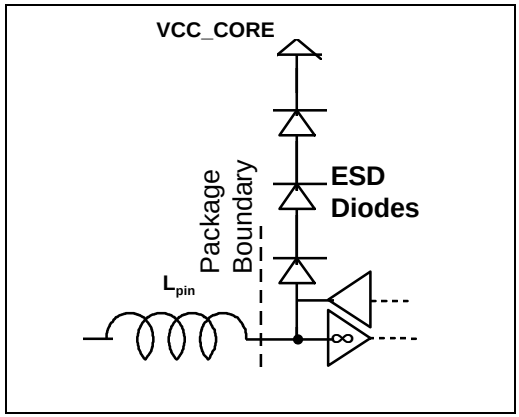

Figure 7. Non-GTL+ ESD Diodes

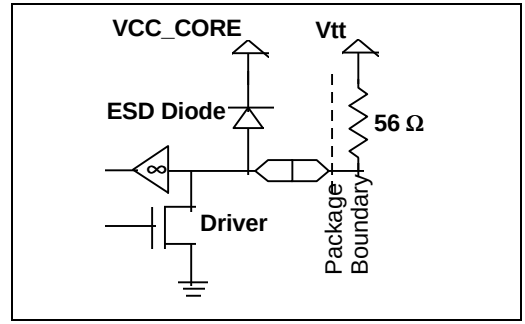

Figure 8. GTL+ ESD Diodes

3.3.1. NON-GTL+ SIGNALS

The non-GTL+ buffers are open drain. When the VCC_CORE supply is on, and the non-GTL+ supply (2.5V) 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 2.5V comes up is avoided.

If the 2.5V supply is on while the VCC_CORE supply is off, the 2.5V supply delivers current to the Slot 1 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 150 ohms will allow a maximum of only 5 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 Slot 1 processor pin.

3.3.2. GTL+ SIGNALS

The GTL+ outputs are also open drain. 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 GTL+ bus.

If V_{TT} is on and VCC_CORE is off, the GTL+ 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 GTL+ bus to operate properly.

3.3.3. MEMORY SIDE SIGNALS

Intel's chipset's 5V tolerant signals are internally buffered in a similar manner. When using 3.3V DRAM there are no memory side sequencing issues. When the 5V supply is on to 5V DRAM and the 3.3V memory controller supply is off, the CAS lines may float. This could cause the DRAM to drive 5V signals to a component that has no voltage applied. The system should provide weak pull-ups to 5V on the CAS lines to prevent the 5V DRAM devices from driving 3.3V inputs when there is no power to the memory controller.

By providing the memory controller with the PWRGOOD signal (as described in Section 1.1), it will drive the CAS lines of the DRAM inactive, and reset the data buffers as soon as it receives 3.3V. 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 GTL+ bus RESET# signal from being created correctly. This action protects these devices from producing bus contention between themselves.

3.3.4. 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 5V input for their ESD protection. This eliminates any issue with turning on its ESD diodes.

3.3.5. 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 OR-ing a 'power good' signal from both supplies, and connecting this output to the chipset and the Slot 1 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.)

3.3.6. 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 2.5V 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 2 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#.

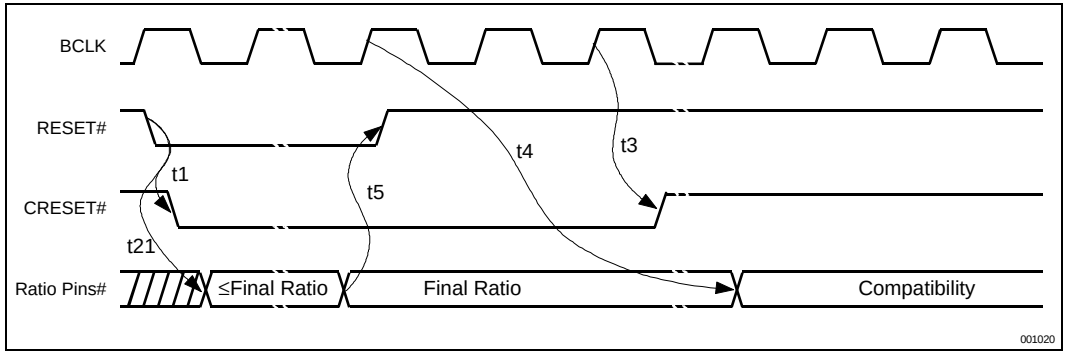

Figure 9. Timing Diagram of Compatibility Pins

Table 2. 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

The circuit below 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-GTL+ 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.

3.3.6.1. Mixed Frequency Processors

Intel does not support nor validate mixing processors rated at different frequencies. In order to support different frequency multipliers to each processor, the design requires two multiplexers. Before implementing this strategy, system designers must understand the impact of mixed processors on the operating systems running on the system.

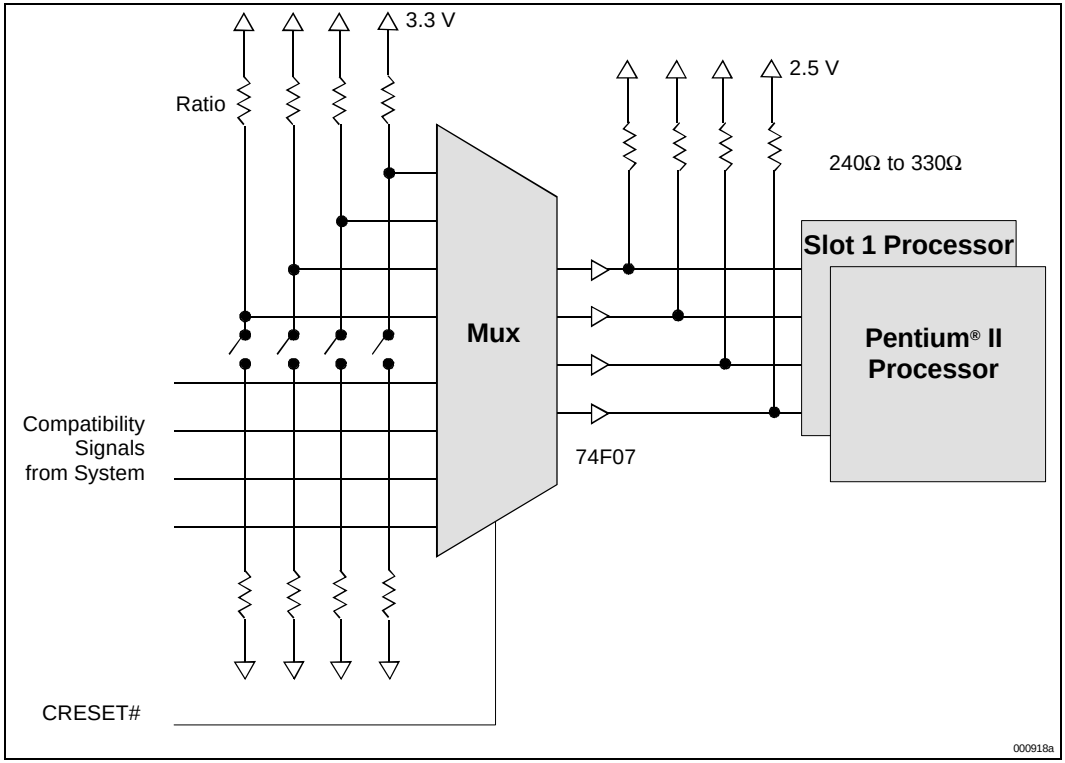

Figure 10. Schematic of Pin Sharing

4.0. MEETING THE SLOT 1 PROCESSOR POWER REQUIREMENTS

Slot 1 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, one step at a time.

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 for each of the components involved. This provides a target for each component and helps reduce iterations to reach a solution.

The high frequency decoupling of the Slot 1 processor eliminates the need to calculate a high frequency budget.

Table 3 provides an estimation of the effects of the factors that system designers need to consider when calculating a low frequency voltage budget.

Table 3. Sample Low Frequency Budget

Component	LF Budget (mV)
Regulator Set Point Tolerance	–25
Bulk Capacitance ESR/ Capacitance Sag	–60
Resistive Losses in Board	–10
Ripple, Noise	–10
Slot 1 Connector	–20
Total	–125

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 Slot 1 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 Slot 1 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 x 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 µamp 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.

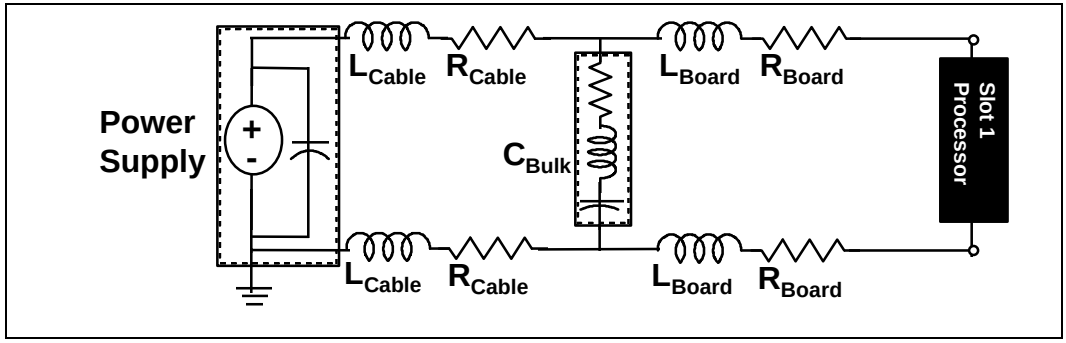

Figure 11. System Design Model

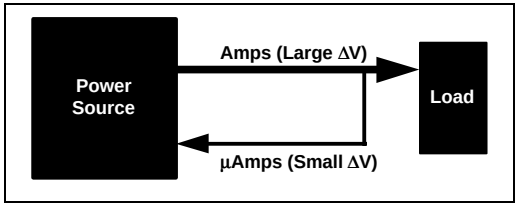

Figure 12. Remote Sense

Either method of regulation can easily maintain ±2% accuracy, plus a small ripple and noise budget, under a *stable* load. However, the further demands of a Slot 1 processor tax the abilities of a remote-sense supply. The large current transients of a Slot 1 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. Converter sockets meeting the upgrade specification in the *Pentium® II Processor at 233 MHz, 266 MHz and 300 MHz* datasheet can then be installed by each empty Slot 1 processor socket to provide an inexpensive upgrade strategy. A socketed regulator or a regulator with selected output levels can power Slot 1 processors.

4.2.2. AC VS. DC INPUT VOLTAGE

The new Slot 1 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 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 5V or 12V taps of typical PC power supplies, and can be manufactured in high volumes. System designers can place DC-to-DC converters very near the Slot 1 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 voltage and output voltage become farther separated as evidenced in Equation 1.

Equation 1. Loss Within a Linear Regulator

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

The linear regulator also requires a minimum drop from the input to the output of about a diode drop (0.5V to 1.0V), 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 4 illustrates the significant power loss and poor efficiency of a linear regulator for a V_{IN} of 5V and a fixed output current of 10 amps.

Table 4. Efficiency of a Linear Regulator

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

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 Slot 1 processors. A 2.8V switching regulator can achieve 80% efficiency at 10 amps.

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 Slot 1 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 Slot 1 connector and substrate inductance (L_{CONN} and L_{SUB}) as well as the board inductance (L_{BOARD}), the Slot 1 processor is designed with approximately 30 ground pins and 26 power pins. These processor design considerations reduce the current slew rate to the order of 30A/us. Slot 1 processors require no external high frequency capacitance, since C_{HF} is sufficient to lower the di/dt to 30 A/us. 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.

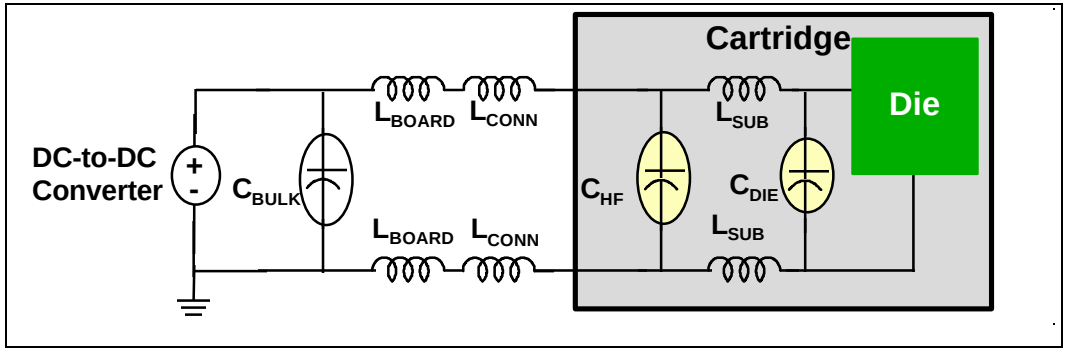

Figure 13. Location of Capacitance in a Power Model with a DC-to-DC Converter

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 µs 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 Slot 1 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 Slot 1 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.

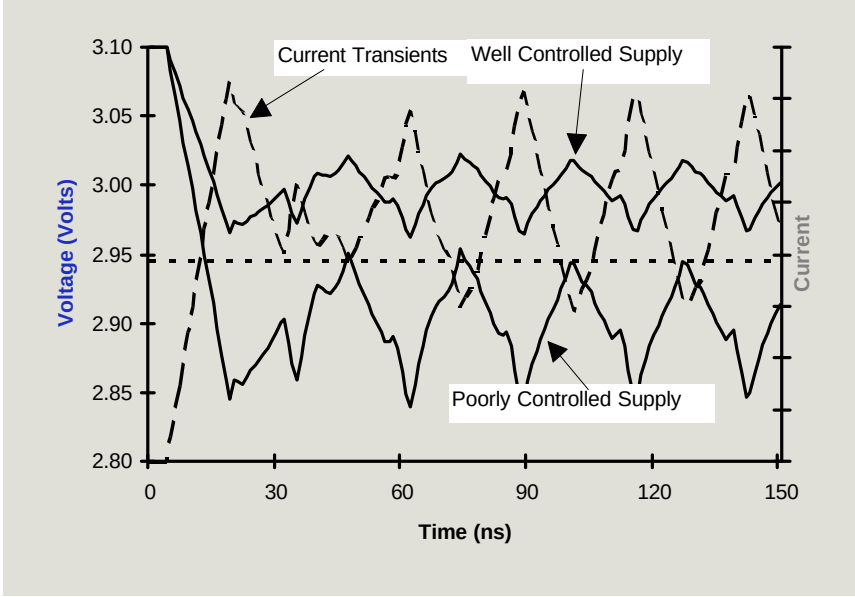

Figure 14. Effect of Transients on a Power Supply

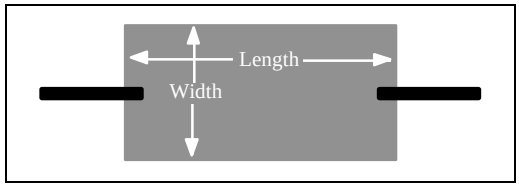

Figure 15. ESR Cylindrical Capacitor

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

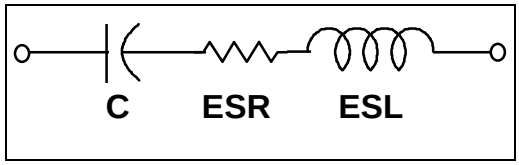

Figure 16. Capacitor Model

Again, this was taken from the complex model of. 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 mΩ of ESR at 10A drops the full budget of 60 mV as shown in Equation 2.

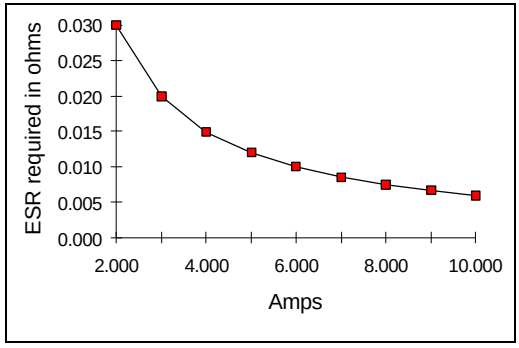

Figure 17. ESR Required for Various Current Demands

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

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

D*I* 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. D*v* is the allowable voltage change budgeted for bulk capacitive sag (discharge) over the period D*t*. D*t* is the reaction time of the power source.

Assuming some representative numbers for *I*, D*V*, and D*t*, the capacitance required is shown by Equation 4.

Equation 4. Capacitance Needed if ESR is 0 Ohms

$$C = \frac{8.5A}{0.060V / 30x10^{-6}\,s} = 4250mF$$

Combining the above formulas 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 D*V* is assumed to be 60 mV, *I* is assumed to be 8.5A, and the reaction time (D*t*) of the power source is 30 µs. The shaded area of the graph covers capacitance types that are insufficient 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 mΩ. The parallel resistance of 12 capacitors would be 4.4 mΩ and the parallel capacitance would be 12,000 µF, which falls in the white zone of the graph in Figure 18.

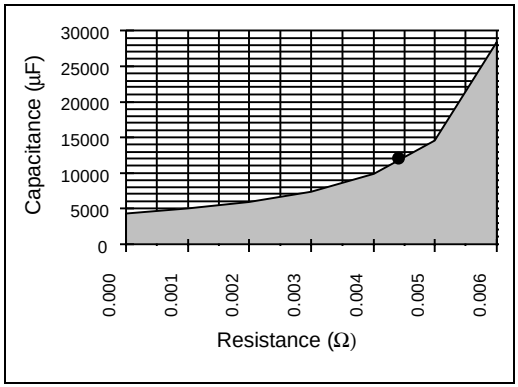

Figure 18. Capacitance Required vs. ESR at 8.5A, 60 mV ΔV, and 30 µs Δt

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 Slot 1 processor power model, including all board, cable, and capacitor effects. See Section 7.5 for more information on component models and Section 11 for the Slot 1 processor power model.

4.3.2. HIGH FREQUENCY DECOUPLING

Slot 1 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. An 0612 is the same size as the 1206 but has its pads along the long edge. The cost of these is significantly higher however due to the complexity of mass producing them. 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/µs 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 7) 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

$$N = 1.9nH \div 0.2nH = 10 capacitors$$ [1]

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 7 in Equation 10.

Equation 10. Resonant Frequency

$$f = \frac{1}{2\pi\sqrt{LC}} \approx \frac{1}{2\pi\sqrt{(0.47\times10^{-9})\times(1\times10^{-6})}} \approx 7.3MHz$$

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 or Islands

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}.

Figure 19 illustrates Intel's recommendation for the layers of a four layer Slot 1 processor baseboard. Route VCC_CORE and VCC_L2 on the V_{CC} layer.

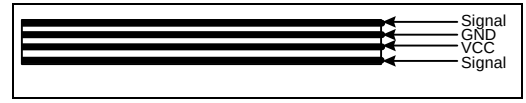

Figure 19. Baseboard Layer Definition

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. See Figure 20 and Figure 21 for examples of voltage islands.

Due to the fact that the Slot 1 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 should be continuous from the source of power to the load. It should also completely surround all of the pins of the source and all of the pins of the load.

Footnotes

[1] More capacitors will actually be required to achieve the necessary capacitance prior to the voltage regulator module due to the limited space within a 1206 package. The number of capacitors required for a Slot 1 processor is therefore "capacitance dependent".

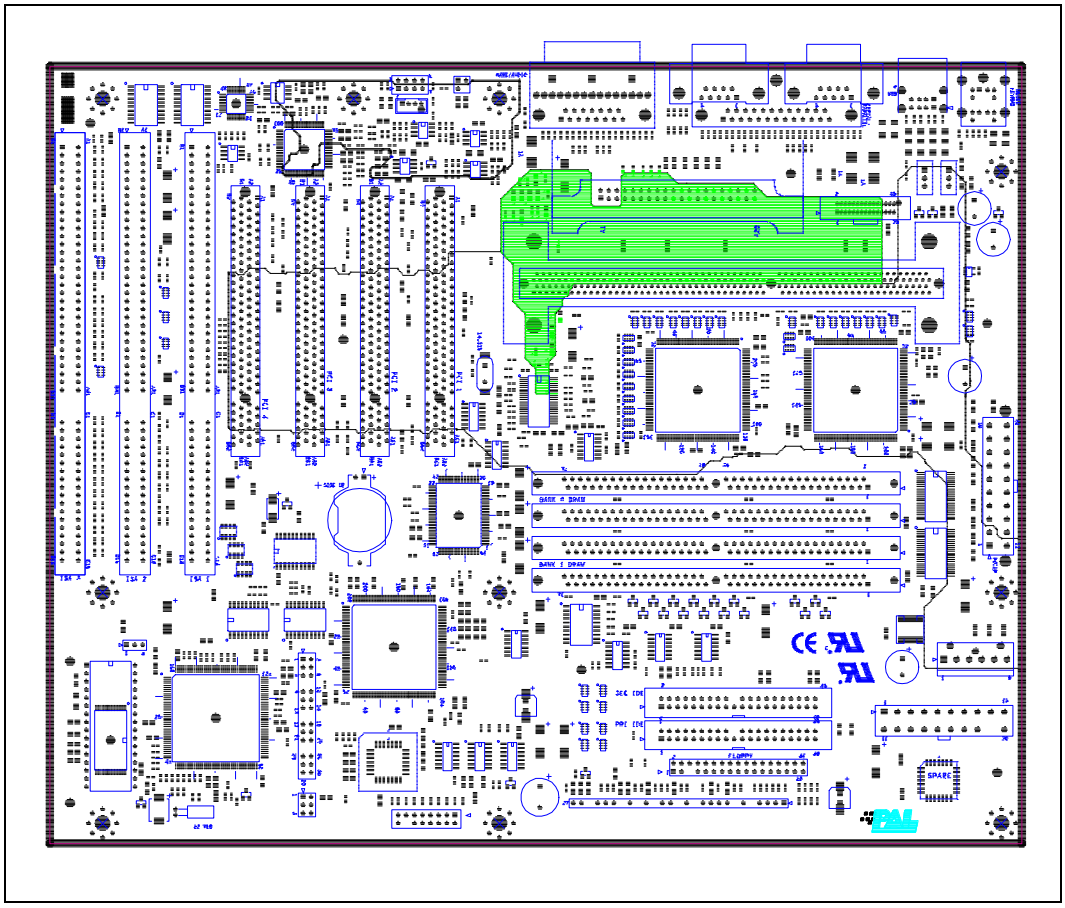

Figure 20. Slot 1 Processor VCC_CORE Voltage Island

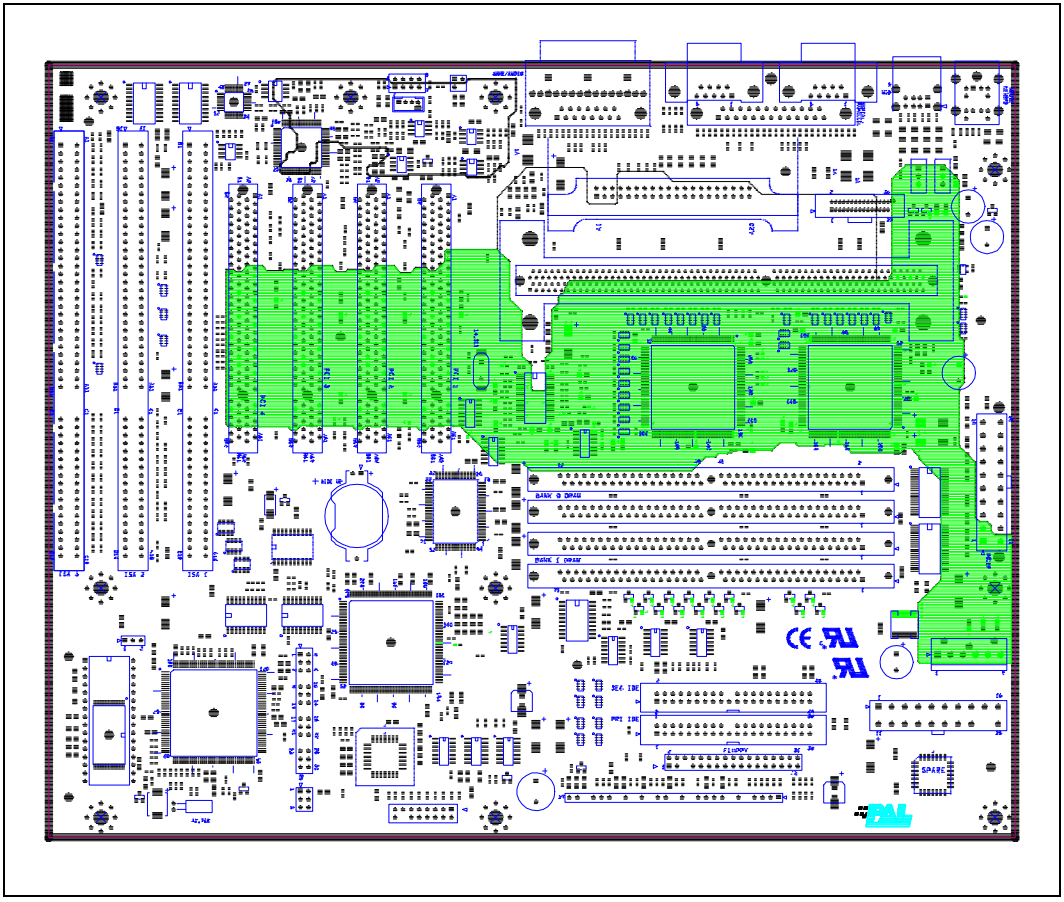

Figure 21. Slot 1 Processor VCC_L2 Voltage Island

4.4.1. LOCATION OF HIGH FREQUENCY DECOUPLING

The Slot 1 processor contains all of the high frequency decoupling required for a properly designed system.

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 22 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 Slot 1 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 GTL+ BUS POWER REQUIREMENTS

The GTL+ bus is an end terminated, open-drain bus. Intel recommends terminating both ends to a voltage level called V_{TT} (1.5V). V_{TT} supplies current when output drivers turn on. There are approximately 150 GTL+ lines in a Slot 1 processor system design.

The GTL+ bus power requirements present a different situation than creating power for the Slot 1 processors. While the GTL+ 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 both ends of the bus. This can be accomplished by having two sources of V_{TT} or by distributing V_{TT}.

A GTL+ buffer sinks a maximum of 45 mA. When considering the duty cycle of the signals, the 141 GTL signals draw a maximum of about 5.38 amps at 100% utilization of the bus. Table 5 illustrates GTL+ current draw using relatively conservative duty cycles. Utilization of the bus, the value of the GTL+ termination resistors, chipset functionality and motherboard design limit actual current draw. Power supply designers need to take these benefits into account as well.

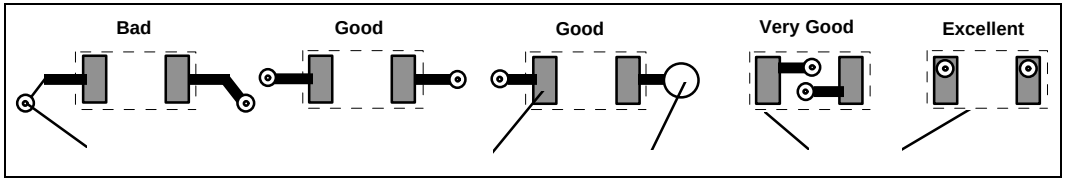

Figure 22. 1206 Capacitor Pad and Via Layouts

Table 5. Estimating V_{TT} Current

Signal Group	Quantity of Signals	Max Duty Cycle	Average Current
Data + ECC	72	100	3.24
Address + Parity	35	67	1.06
Arbitration	7	100	0.32
Request	7	67	0.21
Error	5	20	0.05
Response	6	33	0.09
Other	9	100	0.41
Total	141		5.38

5.1. Tolerance

V_{TT} at the processor edge fingers must be held to ±9%. It is recommended that a regulator that can maintain ±3% at low current draw be used in order to guarantee ±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 GTL+ buffers thus making it more difficult to meet timing specifications.

5.2. Reference Voltage

The GTL+ bus requires a Voltage Reference called V_{REF} as well. The Pentium II 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 µA per device). Bear in mind that leakage current varies and may be significant when building the voltage divider.

6.0. MEETING THE GTL+ POWER REQUIREMENTS

Due to the different nature of powering the GTL+ 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 both ends of the GTL+ bus. Since each Slot 1 processor contains the termination for one side of the bus, a dual-processor (DP) motherboard needs no V_{TT} terminating resistors. A uni-processor (UP) motherboard requires one set of terminating resistors. Designers may wish to generate V_{TT} on each end of the line. In this case, each supply only needs to provide one half of the current necessary to the GTL+ drivers. One larger GTL+ supply may suffice if both ends of the bus are fairly near each other.

When powering the bus from a single regulator, design techniques closely resemble those of Section 4. Motherboard designers should run a full analysis.

The low current required by each GTL+ regulator, in a dual regulator GTL+ supply, means that linear regulators can be used. Linear regulators are faster devices than switching regulators and therefore require less output decoupling. The dual regulator GTL+ supply also results in lower ESR and ESL of supply components. Intel recommends conducting a proper analysis as described in Section 4. Analysis techniques remain the same, reaction time of the supply and the current levels differ.

The two regulators need not track each other as long as each keeps V_{TT} in tolerance. The bus naturally performs an averaging function on these two supplies. V_{REF} (as discussed in Section 6.3) must also track these supplies.

6.2. Distributing V_{TT}

Only V_{REF} and the GTL+ termination resistors need V_{TT}. If the distance to the termination resistors is small, distributing V_{TT} with a wide trace should be sufficient. A wide trace to the V_{REF} generation point keeps inductance to a minimum.

Intel recommends a V_{TT} plane, especially if a single GTL+ supply is used and the ends of the GTL+ bus are not near each other. This helps offset V_{TT} distribution resistive and inductive losses. Again, separate smaller linear regulators at each end of the bus may alleviate the possible need for a power plane.

NOTE

When using resistor networks with single corner pin V_{CC} connections for GTL+ 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.

6.3. Generating and Distributing V_{REF}

V_{REF} is a low current input (about 15 µA per device) to the differential receivers within each of the components on the GTL+ bus. *Each Slot 1 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. V_{REF} should track the V_{TT} averaging of a dual regulated GTL+ supply.

Equation 11 uses $R_1 = 2 \times R_2$ to generate a V_{REF} set at a nominal value of 2/3 V_{TT}. Figure 23 illustrates using 1% resistors to generate the V_{REF} specification of 2/3 V_{TT} ± 2%.

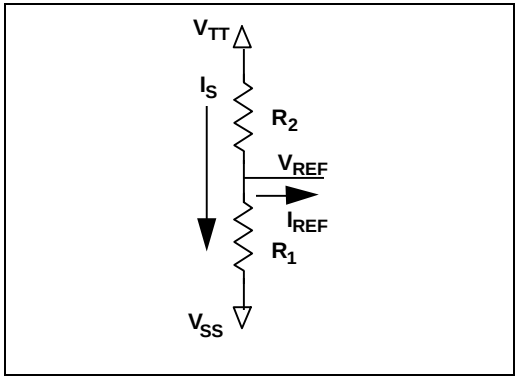

Figure 23. VREF

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

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

R_1 and R_2 should be small enough values that the current drawn by the V_{REF} inputs (I_{REF}) is negligible versus the current caused by R_2 and R_1.

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 R_1 and R_2 at the extremes of their tolerance specifications.

6.3.1. DISTRIBUTING V_{REF} OR V_{TT}

The system board designer may choose to distribute V_{TT} and place a resistor divider at each component or use a voltage regulator to generate V_{REF} and then distribute it to each of the devices. Use wide isolated traces on V_{TT} and V_{REF} to reduce noise and loss.

When using two regulators to generate V_{TT}, V_{REF} must track V_{TT} by averaging from both V_{TT} sources. Generate a separate V_{REF}, at each regulator, for every four loads. Connect the V_{REF}s together with a wide trace. The closer this V_{REF} signal tracks the path of the bus signals, the better it matches the average of the voltage on the GTL+ bus. Route V_{REF} on a separate layer to reduce cross-talk.

7.0. RECOMMENDATIONS

Intel recommends using simulation to design and verify Slot 1 processor based systems. With the above estimates, a model of the power source, and the model of the Slot 1 processor provided in Section 11, system developers can begin analog modeling. Intel

recommends the following as a starting point or benchmark:

7.1. VCC_L2

For VCC_L2, use a typical PC power supply with a 3.3V tap. Insure sufficient current on the 3.3V tap of the supply to power all of the system chipset, the GTL+ regulator (if run off of 3.3V), other 3.3V logic in the system and any possible L2 caches that may someday exist in the system. See the chipset specification for chipset power requirements. See the *voltage and current specifications* in the *Pentium® II Processor at 233 MHz, 266 MHz and 300 MHz* datasheet for the requirements of the L2 cache. Bulk decoupling requirements depend on the reaction time of the power supply. The Slot 1 processor contains sufficient high frequency decoupling, provided that the system has a well designed power supply.

7.2. V_{CC}_CORE

For VCC_CORE, Intel recommends starting with a socketed local DC-to-DC converter as shown in Figure 24. This removes cable inductance from the distribution, reduces board inductance, and allows for a low cost upgrade strategy as well. Regulator sockets can be provided for upgradable processor sockets rather than shipping with the full current capability already available. Another benefit 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. The output of this regulator should be adjustable to allow for changes in the voltage specification as new products become available.

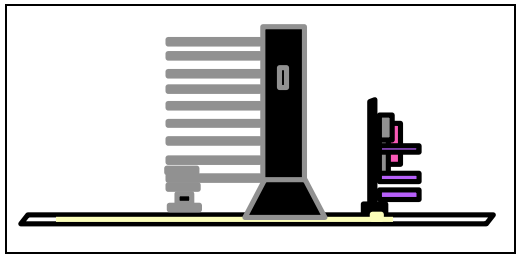

Figure 24. Local Regulation

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 Slot 1 processor contains all of the high frequency decoupling required for a properly designed system.

7.2.1. 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 5V or 12V outputs of the main supply. Voltage Regulator Module(VRM) DC-to-DC converter specifications describe the guidelines given to the DC-to-DC converter industry. Intel has located these specifications on its web site at:

http://developer.intel.com/design/PentiumII/applnots

Table 6 below defines the logic levels of the Slot 1 processor's VID pins. Intel recommends connecting the VID pins to the VRM as described in the appropriate VRM DC-to-DC converter specification.

7.3. V_{TT}

Intel recommends supplying V_{TT} to each end of the GTL+ bus using a separate linear regulator for each end. Since the losses in a linear regulator are directly proportional to $V_{IN}–V_{OUT}$, the 3.3V power supply makes a good choice for the input voltage to the regulator. The processor voltage (VCC_CORE) may seem like a better choice since it is lower than 3.3V, but it varies from one Slot 1 processor variant to the next. This would lead to a design change for each generation of Slot 1 processor. Also, Linear regulators require a minimum voltage drop in order to operate. This becomes an issue as Slot 1 processor voltages decrease.

Separate linear regulators contain the voltage distribution to a very local region. In a bus layout where both ends of the bus are physically near each other, one regulator can supply both sets of termination resistors. In this situation, a 50 mil trace (the wider the better) should distribute the power to the termination resistors.

Linear regulators are fairly common and produced by many vendors. See your local field applications engineer for assistance locating a vendor.

Table 6. Voltage Identification Code

Slot 1 Processor Pins					VCC_CORE (VDC)
VID4	**VID3**	**VID2**	**VID1**	**VID0**	
0	1	1	1	1	Reserved
0	1	1	1	0	Reserved
0	1	1	0	1	Reserved
0	1	1	0	0	Reserved
0	1	0	1	1	Reserved
0	1	0	1	0	Reserved
0	1	0	0	1	Reserved
0	1	0	0	0	Reserved
0	0	1	1	1	Reserved
0	0	1	1	0	Reserved
0	0	1	0	1	1.80
0	0	1	0	0	1.85
0	0	0	1	1	1.90
0	0	0	1	0	1.95
0	0	0	0	1	2.00
0	0	0	0	0	2.05
1	1	1	1	1	No Core
1	1	1	1	0	2.1
1	1	1	0	1	2.2
1	1	1	0	0	2.3
1	1	0	1	1	2.4
1	1	0	1	0	2.5
1	1	0	0	1	2.6
1	1	0	0	0	2.7
1	0	1	1	1	2.8
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

NOTES:

0 = Processor pin connected to V_{SS}.

1 = Processor pin open (system design may include pull-up resistor to voltage consistent with TTL input levels).

2 = To ensure the system is ready for all Slot 1 processors values in **BOLD** must be supported.

Reaction time specifications of the regulator determine bulk capacitance. The capacitance must hold-over the regulator during a switch from 0 to 5.4 amps, as estimated in Table 5, until the regulator reacts. In addition, Intel recommends one 1.0 µF capacitors for each termination resistor package for high frequency decoupling. Place these capacitors as near to the termination resistors as possible.

7.3.1. TERMINATION RESISTORS

Discrete resistors may be employed, however the assembly time associated with placing about 140 resistors in a UP system should be taken into account. A lower part count implementation uses resistor networks.

NOTE

When using resistor networks with single corner pin V_{CC} connections for GTL+ 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.

7.4. V_{REF}

Intel recommends one voltage divider at each component. The Slot 1 processor generates its own V_{REF} internally. Therefore, most motherboards require only two V_{REF} voltage dividers.

Tie all of the V_{REF} inputs of each component together. Assume a maximum of 15µamps 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 both V_{REF} Loads. Using 1% resistors for the voltage divider in Figure 23, make R_1 a 150Ω resistor, and use 75Ω for R_2. This creates a static usage of 7 mA (1.5V/225Ω) per voltage divider. After looking at all combinations of R_1 and R_2 (above and below tolerance) and I_{REF} (±30 µA), the worst case solution for Equation 14 can be found with I_{REF} at 30µamps, R_1 at the low end of its tolerance specification (148.5Ω), and R_2 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 1.00V, 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 R_1 and R_2 have additional benefits: The parallel combination terminates the V_{REF} line to 50 ohms. This generally available resistance value reduces resistor cost.

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 loads, use a 30-50 mil trace (The wider the better) and 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.

7.5. 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 Slot 1 processor model can be found in Section 11.

Table 7. Various Component Models used at Intel (Not Vendor Specifications)

Component of Simulation	ESR (Ω)	ESL (nH)	ESL+ Trace + Via (nH)
0.1µF Ceramic 0603 package	0.100	1.60	3.0
1.0µF Ceramic 1206 package	0.120	0.47	1.9
100µF MLC (2.05"x0.71")	0.005	0.30	1.7
47µF, 16V Tantalum D Case	0.100	0.602	2.0
330µF, 16V Aluminum Electrolytic	0.143	2.37	3.8
1000 µF, 10V Aluminum Electrolytic (20mm)	0.053	N/A	N/A
1000 µF, 25V Aluminum Electrolytic (25mm)	0.031	N/A	N/A
L_{BOARD}. One used for V_{SS} 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 Slot 1 processor filters them out.

9.0. EXISTING TECHNOLOGY FOR A SLOT 1 PROCESSOR SYSTEM DESIGN

9.1. Solutions for VCC_CORE

Intel has assisted in the development of many industry DC-to-DC converter modules. Designers should understand the use of one of these components in a specific design. Intel can not guarantee the use of a DC-to-DC converter in specific designs. In general, the vendor of any component assists designers in the usage of their component. See your local field office for a list of possible vendor solutions.

Another solution that is a simple extension to the discussion in this paper is to integrate the components of the DC-to-DC converter, including the bulk capacitance, onto the system PCB. Intel has helped power silicon vendors as well in designing Slot 1 processor specific solutions. Again, see your local field office for a list of possible vendor solutions.

9.2. Linear Regulators for V_{TT}

Linear regulators are widely available. Switching regulators can also be used to generate V_{TT}.

9.3. Termination Resistors

Intel recommends the use of resistor networks to reduce the part count of the processor board assembly. 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.

When using resistor networks with single corner pin V_{CC} connections for GTL+ termination, beware of inductive packages. These packages can cause significant voltage drops due to the inductance in the 24 pin SOIC packages being used for this purpose.

10.0 SLOT 1 PROCESSOR POWER DISTRIBUTION NETWORK MODELING

Intel provides the AC electrical models shown in Figure 25, Figure 26 and Figure 27 for use in the simulation of the AC transient response of the Slot 1 processor power delivery systems. Due to tool capability limitations, these models have been greatly simplified and are provided as a rough illustration of the Slot 1 power delivery systems.

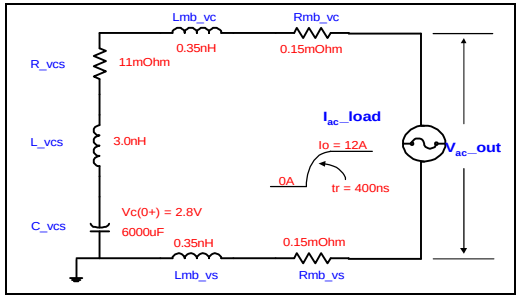

Figure 25. VCC_CORE Power Delivery Model for AC Transient Response

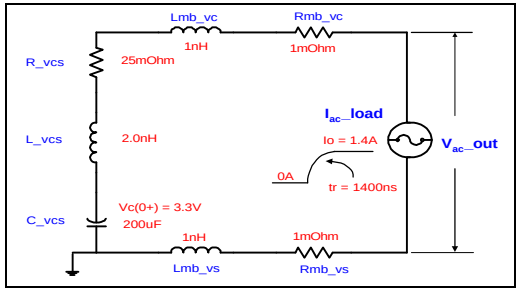

Figure 26. VCC_L2 Power Delivery Model for AC Transient Response

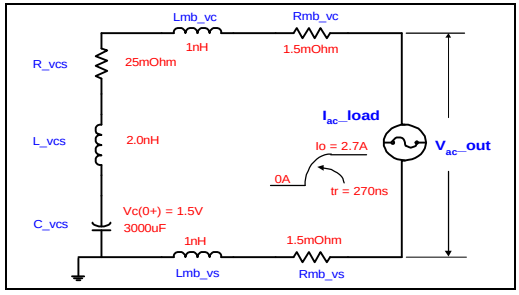

Figure 27. VTT Power Delivery Model for AC Transient Response

The exponential waveform of the form:

$$I = Io\ (1 - exp(-t/tr))$$

represents the Slot 1 processor as a current source. Motherboard inductance and impedance characteristics are modeled as Lmb_vc, Lmb_vs, Rmb_vc and Rmb_vs. Finally source impedance, inductance and capacitance characteristics are modeled as R_vcs, L_vcs, and C_vcs. The values provided below for the motherboard and source characteristics are for example purposes and all source and motherboard characteristics should reflect the actual values for the system under analysis.

11.0 RIGHT ANGLE CONNECTOR POWER DELIVERY CONSIDERATIONS

This chapter compares the effects of a right angle connector to the typical Slot 1 connector.

System developers considering using a right angle Slot 1 connector (right angle connector (RAC)) need to keep in mind the following constraints:

- The added delay of the RAC increases risk to both higher frequency implementations and implementations that use single ended termination.
- System designers must take into account the added resistance and inductance of the RAC.

11.1. Right Angle Connector Assumptions

The following assumptions provide the basis for the right angle analysis in the following sections:

- Comparison is between motherboard regulator solutions for both typical Slot and right angle connectors.
- The motherboard implementations will not use complex load offset circuitry prevalent on several VRM solutions.
- Static requirements: 2,800 ±100 mV.
- Specifications for worst case negative going transient (current load changes from 1A to 13A) must not go below 2,604 mV level (measured at Pentium II pads in bandwidth of 20 MHz).
- Specifications for worst case positive going transient (current load changes from 13A to 1A) must not exceed 2,900 mV for more than 2 us.

- Validation data taken on a production Slot 1 platform can be linearly scaled to worst case current levels.
- Set point accuracy is ±35 mV (1.2%) of level setting.
- Regulator decoupling equivalent to that used on VRM's.
- RAC resistance and inductance values three times greater than the typical Slot 1 connector.

11.2. Right Angle Connector Static and Dynamic Analysis

Table 8 and Table 9 illustrate the predicted impact of the RAC on both the high end of the Slot 1 processor voltage specification and the low end of the Slot 1 processor voltage specification.

Table 8. RAC Static Requirements, High End Comparison

	Typical Slot 1 Connector (mV)	Right Angle Slot 1 Connector (mV)
Set Point	2,835	2,835
Set Point Accuracy	35	35
Ripple	10	10
Total	2,900	2,900

Table 9. RAC Static Requirements, Low End Comparison

	Typical Slot 1 Connector (mV)	Right Angle Slot 1 Connector (mV)
Set Point	2,835	2,835
Set Point Accuracy	35	35
Ripple	10	10
Thermal Drift	20	20
IR Drop(13A) (1)	13 × (1 + 1) = 26	13 × (1 + 3) = 52
Total	2,764	2,738

NOTES:

1. R includes the motherboard trace resistance plus the Slot 1 connector resistance.

The RAC has little impact under static conditions. By re-targeting the voltage regulator's set point, system designers can easily overcome the minor voltage drop.

Table 10 shows that systems with RACs may have a tougher time meeting dynamic requirements. Systems with RACs need to have more decoupling to overcome the increase in IR drop due to longer pins on the connector.

Table 10. RAC Dynamic Requirements

	Typical Slot 1 Connector (mV)	Right Angle Slot 1 Connector (mV)
Set Point (Low)	2,820	2,820
Ripple	10	10
L\Req. cap.	150	150
Connector	30	90
Power Planes	26	26
Total	2,604	2,544 (60 mV below spec)

11.3. Right Angle Connector Power Distribution Comparison

Figure 28 provides a model with which to assess the RAC's impact on power distribution to the Slot 1 processor.

The power distribution design must ensure that the voltage at the processor connector gold fingers (V_{CD}) stays within specification limits during a processor transition from I_{MIN} to I_{MAX}. To accomplish this, sum up the variation caused by this load transient at each of the node pairs identified in Figure 28.

Changes to the voltage at the regulator output, V_{AF}, are a function of the amount of energy storage in the bulk capacitor, the response time of the voltage regulator, and capacitor parasitics. At this point the dominant parasitic is the resistance, as the rate of change of load current is slow enough to minimize the Ldi/dt drop. The voltage at the Slot 1 connector pads, V_{BE}, is a function of the path length of the power/ground plane segments connecting the regulator to the processor connector (for VRM's it would also include the L and R of the regulator connector).

The voltage drop across the connector (typical Slot 1 or RAC) is $V_{BE} - V_{CD}$. Based on estimated contact lengths for the RAC, its impedance would be three times that of the typical Slot 1 connector. The key impedance parameters are power pin to adjacent ground pin inductance and contact resistance.

11.4. Right Angle Connector Typical Implementation

Intel analyzed a typical voltage regulator scenario to determine the impact on decoupling requirements of the RAC.

The regulator had a set point accuracy of ±1.2% (over a temperature range of 0°C to 60°C, with input voltage variation of ±5%) and a response time on the order of 10 us. Intel sized the bulk decoupling to ensure meeting the V_{CD} specification under worst case conditions. Intel estimates the cost of this decoupling at $2.20 in high volumes.

Intel adjusted the nominal set point level to ensure that the regulator never exceeded the upper bound of 2.9V. At this upper bound, the effects of IR drops ($I = I_{MIN}$) could be neglected so that, with an overall set point accuracy of 1.2%:

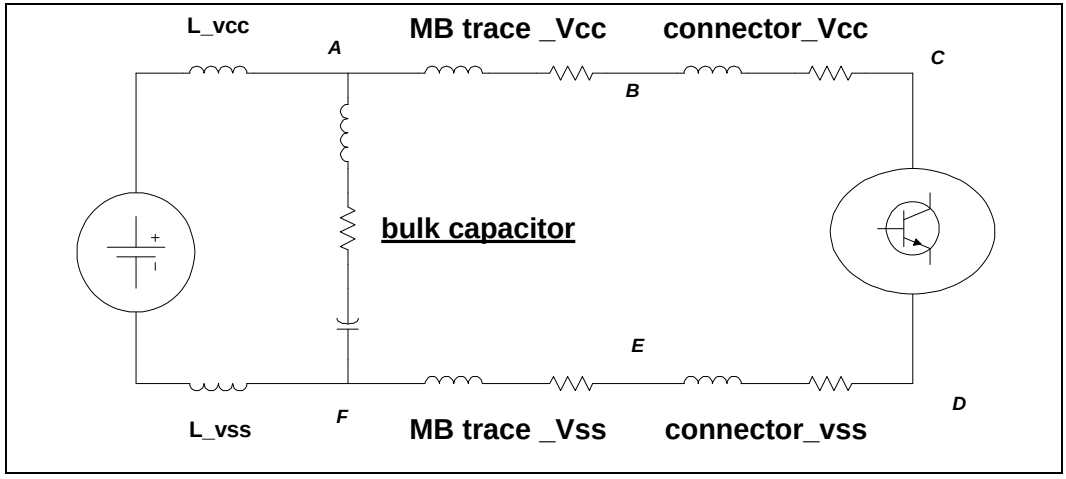

Figure 28. Power Distribution Path

Equation 16. Nominal Set Point for a Right Angle Connector

$$Vaf = Vcd = 2.900V - 0.012(Vcd)$$

$$Vaf = Vcd = 2.900/1.012 = 2.866V$$

For a load change from 1A to 13A (I_{MIN} to I_{MAX}) the lowest initial voltage setting will be:

Equation 17. Initial Voltage Setting for a Right Angle Connector

$$Vaf = 2.866/1.012 = 2.832V$$

Empirical testing of the typical Slot 1 connector shows that a 53 mV drop ($V_{CB} - V_{DE}$) will be developed across the connector for a 12A transient (from 1 to 13A). The worst case drop for a two inch planar interconnect from the regulator to the connector is 10 mV. Since the system specification for the minimum voltage level at the gold fingers, V_{DE}, is 2.605V (2.800 to 0.195), the transient drop, ΔV_{AF}, across the bulk capacitance may now be calculated:

$\Delta V_{AF} = \Delta V_{CD} - \Delta (V_{BE} - V_{CD}) - \Delta (V_{AF} - V_{BE})$ or $\Delta V_AF = (2.832 - 2.605) - 0.053 - 0.010 = 0.164V$

With a 3X increase in connector impedance for a RAC, the value of $\Delta (V_{BE} - V_{CD})$ will increase from 0.053 to 0.159V (3X increase). To keep the V_{CD} transient level unchanged, V_{AF} needs to decrease. So ΔV_{AF} must change from 0.164 to (0.164 - [0.159 – 0.053])V or $\Delta V_{AF} = 0.058V$

The ratio of the change in V_{AF} for a typical Slot 1 implementation to the change in V_{AF} for a RAC implementation is: 0.164/0.059 = 2.8.

A RAC implementation requires the amount of bulk decoupling to increase by 2.8X over a typical Slot 1 implementation. This increases decoupling costs by approximately 2.8X, or a net dollar increase of approximately $4 (2.8x$2.20 - $2.20).