

# Integrated Modeling Methodology for Core and I/O Power Delivery

Kaladhar Radhakrishnan\*, Yuan-Liang Li, and William P. Pinello  
Intel Corporation

5000 W. Chandler Blvd., CH5-157, Chandler, AZ 85226

[kaladhar.radhakrishnan@intel.com](mailto:kaladhar.radhakrishnan@intel.com), Tel: (480) 554-2146, Fax: (480) 554-7615

## Abstract

Traditionally core power delivery and I/O signal analysis were performed separately to analyze the performance of a microprocessor package. The coupling between I/O and core power delivery was assumed negligibly small. In this paper, we describe a methodology to for analyzing core and I/O power delivery using the same integrated model. Having an integrated model allows us to study the noise induced on the core power nets due to switching currents on the I/O nets and vice-versa. Having an integrated model also gives us the opportunity to study the relative merits of separating or combining the core and I/O power networks.

## I. Introduction

The power delivery networks in a system can be classified into core power and I/O power networks. In some processors, the core and I/O power nets are combined together while in others they are isolated. The power delivery loop starts with the VRM and includes the motherboard, the package and the die itself. Due to the inductance associated with the power delivery loop, any transient current through the loop will cause a drop in the voltage available to the die. The core power network acts as a power supply to the die and any drop in voltage will directly impact the maximum operating frequency of the processor. On the other hand, the I/O power network supplies power to all the I/O buffers on the die. Any drop in the voltage on the I/O power net will introduce delays in the signals and could affect the timing and in some extreme cases cause erroneous switching. One way to combat the inductive drop is by adding several stages of decoupling capacitors at various points in the power delivery loop. This helps lower the overall impedance of the power delivery loop over a broad frequency range.

Figure 1 represents a simplified model of core and I/O power networks. A separate I/O power net is assumed. Planes are used for local power distribution, which help minimize the lateral inductances. The load on the core power network can be modeled as a transient current source. For the I/O networks, the current is drawn from the source to drive the I/O buffers on the die. From Figure 1, we can see that both the core and I/O power networks share the same ground. Consequently, any transient current through either network will influence the other network by causing an inductive drop along the ground plane.

Traditionally, the power delivery problem has been modeled by using a two dimensional distributed model to represent the planes in the package [1]. This is usually a two step process. The first step is to extract the parasitics of

different sections of the packages. These parasitics are then imported into a Spice netlist. The transient simulations are then carried out in Spice. In this paper, we use the simulation tool Speed [2] to model both core and I/O power delivery. Speed features a field solver that is integrated with a circuit and transmission line solver. The field solver is based on the 2-D finite difference time domain (FDTD) scheme and is used to solve for the fields between the planes in the package. By including the entire package in the simulation model, it is possible to account for package resonances and via interactions. All the signal lines are solved using the transmission line simulator while the external circuit connections are solved using the circuit simulator. By constructing an integrated model in Speed, we can quantify the interactions between the core and I/O power networks.

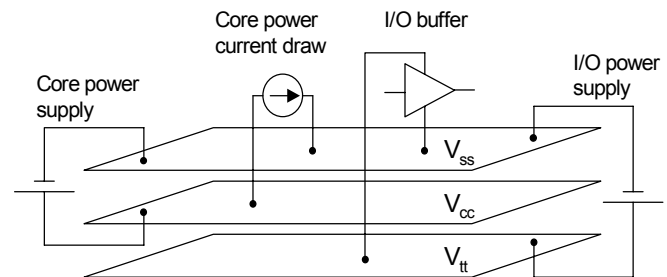


Figure 1. Core and I/O power distribution networks.

## II. Modeling Methodology

The first phase of the methodology is to model the package in Speed. Speed has a translator associated with it which allows the user to convert the files from Cadence format (.brd or .mcm) to one that is recognized by Speed. The user also has the option to specify the nets that need to be extracted into the Speed model. For an integrated model, the Vss, Vcc, Vtt and the signal nets would need to be extracted. For complex packages, the user needs to make some simplifications to keep the simulation time within reasonable limits. The grid size in the Speed model needs to be small enough to resolve every single via on the package. For a complex package, this might translate to a 500x500 grid. The simulation time for such a grid might be unacceptably long, especially if the package has a large number of layers. The simulation time for the FDTD portion can be expressed as

$$t_{FDTD} \propto N_x \cdot N_y \cdot N_t \cdot N_{grids}$$

Here  $N_x$  and  $N_y$  represent the number of grid points along the x and y directions.  $N_t$  is the number of time steps required and is proportional to  $N_x$  or  $N_y$  in order to satisfy the stability criterion.  $N_{grids}$  represents the number of FDTD grids required to solve the problem. There is one FDTD grid associated with each plane pair. The number of FDTD grids is usually one less than the number of planes in the package. From the above expression, we can see that a package discretized using a 500x500 grid will take 8 times longer to solve than the same package discretized using a 250x250 grid. Thus it is important to keep the grid size in check.

Usually it is the microvia pitch that determines the minimum grid spacing. However, the microvia inductance is not significant when compared to the overall inductance of the power delivery network from the die to the power supply. Consequently, one can move the locations of the microvias without any significant impact on the results. In some cases, even the number of microvias can be reduced. In such cases, the dielectric thickness was reduced by the same factor to keep the overall inductance constant. This approach was also used in [3]. Another approach used in [3] was to reduce the number of power and ground planes. For example, consider a package with four pairs of power and ground planes, each of them separated by 100 microns. They can then be approximated as two pairs of power and ground planes separated by 50 microns. However, this is useful only if the package has a large number of dedicated power and ground planes.

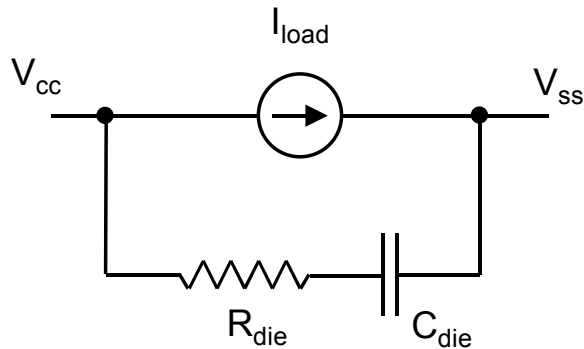


Figure 2. Current source representing the die load.

Once the package model is complete with a reasonable grid size, the next step is to hook up the external circuits. For the core power network, the die is usually modeled as a distributed grid of current sources of the kind shown in Figure 2. For typical core power delivery simulations, the current source is assumed to be piece-wise linear with a large current ramp initially and is constant thereafter. However, one of the objectives of creating the integrated model was to study the impact of transient currents through the core power network on the I/O power rails. Using a simple ramp function for the current source will not produce any transient current after the first few nanoseconds. As a result, the noise coupled to the I/O power nets will be minimal after the first few nanoseconds. Consequently, the current source is defined to be a periodic pulse with an amplitude smaller than the one used in the initial ramp. Each current source is hooked up with shunt on-die capacitance and resistance. The socket pins, the

mother board and the VRM are modeled as lumped elements which are hooked up to the bottom of the package. The decoupling capacitors associated with the package are connected as external circuit elements with an appropriate ESL and ESR.

For I/O power delivery, the goal is to model the noise induced on the I/O power rails as a number of drivers switch simultaneously. A sudden spike in the current tends to induce a voltage drop along inductive path as the current flows through the signal lines and the power and ground planes. It is crucial to keep this voltage drop within acceptable limits to ensure signal integrity. In our model, the  $V_{tt}$ ,  $V_{ss}$  and the signal nets are imported into Speed. The buffer model used to represent the drivers is shown in Figure 3. The output node is connected to the signal line while the  $V_{tt}$  and  $V_{ss}$  nodes are connected to the nearest I/O power and ground vias. The variable resistor  $R_2$  is controlled by the voltage applied at the input node.  $R_{open}$  is a large resistance used to complete the circuit. Each buffer model has some decoupling capacitance associated with it. The signal lines are terminated using 50 ohm pull up resistors. At the other end, just as in the core power delivery case, the socket pins, mother board and VRM are modeled as lumped elements.




Figure 3. I/O Buffer models.

### III. Numerical Results

To test the methodologies described in the previous section, a six-layer package was modeled on Speed. The package stackup is listed in Table 1.

Layer	Plane Function
L1	Signal Routing
L2	Ground (Vss)
L3	I/O Power (Vtt)
L4	Core Power (Vcc)
L5	Ground (Vss)
L6	PGA Pins

Table 1. Package stackup

Figure 4 shows the portion of the top layer underneath the die. The signal lines are located on the periphery of the die. The Vtt vias are also located on the periphery close to the signal lines. The Vss and Vcc nets are arranged in a checkerboard pattern on the top layer. Each column of vias is connected using vertical power bussing strips.




Figure 4. Top layer underneath the die.

A 5x5 distributed network of current sources is used to simulate the die load for the core power network. Each of these current sources is modeled as a pulse function with a period of 30ns. The I/O buffer model shown in Figure 3 is used to model the drivers on the die. These I/O buffers were driven with a switching waveform at a frequency of 66 MHz. The core and I/O power at all nodes were monitored for the following three cases:

1. Core switching with I/O idle,
2. I/O switching with core idle, and
3. Core and I/O switching.

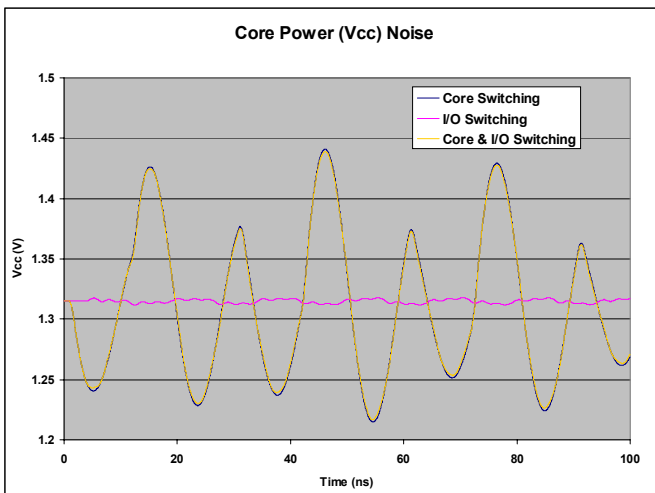


Figure 5. Core power noise for the three different cases.

Figure 5 plots the worst-case voltage differential between adjacent Vcc and Vss nodes for all three cases. The waveforms for cases 1 and 3 are almost identical. This indicates that the core voltage level is not very sensitive to I/O switching. The waveform for case 2 further validates this notion. There is hardly a ripple in the core voltage when the I/O buffers are switching with the core idle.

Figure 6 plots the worst-case Vtt voltage available to the drivers for all three simulation cases. The waveform for case 1 indicates that the I/O voltage is sensitive to transient currents through the core power network. A peak to peak swing of 40 mV is observed when the core is switching with I/O idle. When just the I/O drivers are switching (case 2), the noise on the I/O power rails is in the form of a high frequency ripple. The waveform for case 3 appears to be a superposition of the first two cases.

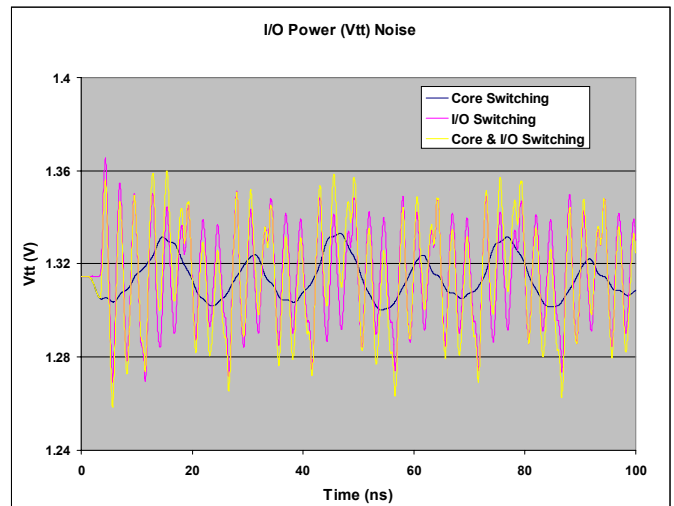
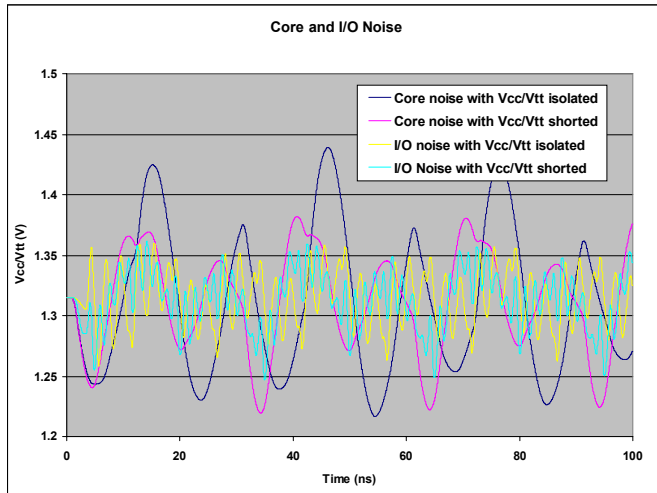


Figure 6. I/O power noise for the three different cases.

So far, all simulations were carried out with separate core and I/O power nets. By shorting the Vcc and Vtt nets, we can study the noise levels without a dedicated I/O power net. A fourth simulation run is carried out with the Vcc and Vtt net shorted together. Figure 7 plots the core and I/O voltage levels with Vcc and Vss shorted. The voltage levels with separate Vcc and Vss nets are included for reference. We can see that core voltage swings are significantly lower with the two nets shorted together. When Vcc and Vtt are isolated, the peak to peak swing of Vcc is about 220mV. On the other hand, when Vcc and Vtt are shorted together, the peak to peak swing of Vcc is just 160mV. From the figure, we can also see that the frequency of oscillations on Vcc has changed slightly. This indicates that the system is resonating at a different frequency. This can be attributed to the fact that shorting Vcc and Vtt would reduce the effective package inductance for the core power delivery loop. Since the resonant frequency is inversely proportional to the square-root of the package inductance, it follows that the new resonant frequency will be slightly higher. This is what we observe from the waveforms shown in the figure. However, shorting Vcc and Vtt increases the I/O voltage swings when compared to the case when Vcc and Vtt are isolated.



**Figure 7.** Core and I/O noise levels with Vcc and Vtt isolated, and with Vcc and Vtt shorted.

## V. Summary

An efficient methodology for simultaneous analysis of the core and I/O power delivery networks is presented in this paper. By using an integrated model, the coupling between the two networks through the common ground is quantified. Simulation results on a typical package indicates that the core power network is not very sensitive to switching currents through the I/O power network. However, transient currents through the core power network can couple significant noise onto the I/O power network. This result is not surprising since the total current through the core power network is usually much larger than the current through the I/O power network.

Simulations were also carried out to study the relative noise levels when the Vcc and Vtt nets are shorted together. The voltage swing on the core power rails went down from 220mV to 160mV when Vcc and Vtt were shorted together. However, the noise on the I/O power rails increased slightly when the two nets were shorted. Shorting Vcc and Vtt allows the two networks to share the package planes and the decoupling capacitors. However, shorting the two nets tends to equalize the noise levels for both core and I/O power. In some cases, the tolerance level for one power delivery network might be much smaller than that of the other. For such systems, it might be prudent to go with dedicated core power and I/O power nets.

## VI. References

- [1] K. Lee, and A. Barber, "Modeling and Analysis of Multichip Module Power Supply Planes," *IEEE Trans. Comp., Pkg., Manufact. Technol.-Part B*, vol. 18, no. 4, Nov. 1995, pp 628-639.
- [2] Simulation Package for Electrical Evaluation and Design (SPEED 2000), Sigrity Incorporated, Santa Clara.
- [3] B. Garben, and M. F. McAllister, "Novel Methodology for Mid-Frequency Delta-I Noise Analysis of Complex Computer System Boards and Verification by Measurements," 9<sup>th</sup> Topical Meeting on EPEP, pp 69-71, 2000.