

Beyond IR Drop: Dynamic Voltage Droops & Total Power Integrity

Raj Nair & Donald Bennett, [Anasim](#) Corp., March 31, 2008

Most power supply and system power delivery designers today are intimately familiar with concepts of power network impedance, bypass or decoupling capacitor frequency responses, and important roles played by device electrical aspects such as effective series resistance and loop inductance. As entire systems are integrated onto silicon chips, these system-level, frequency-domain, 'analog' design and verification aspects are becoming increasingly important, particularly in nanoscale SoC design. Power Integrity analysis in SoC's is hence moving from traditional IR Drop to *total power integrity*, and *true-electromagnetic simulations*, comprehending all aspects of interactions between integrated circuit blocks and the common power delivery network. This article sheds light on key differences such as voltage droops and noise wave propagation resulting from on-chip load interaction with power network impedance and discusses how total power integrity may be rigorously inspected through rapid analyses and physics-based simulations with corresponding benefits to SoC cost and time-to-market/money.

So what is Total Power Integrity?

Prior articles on the Power Integrity Wall [1] and Power Integrity and Energy aware Floor Planning [2] discuss the looming problem with power integrity (PI) in nanoscale ULSI and how a comprehensive PI methodology may be employed in optimal SoC floor planning without providing a theoretical basis that clarifies total power integrity (TPI).

In brief, TPI is comprehensive, true-electromagnetic analysis of interactions between load circuits in an SoC/SiP with power distribution network partitions and delivery system that provides a spatial and temporal view of the degradation of the ideal nature of power supply voltage provided.

The EDA industry enabling ULSI design has been moving towards TPI, albeit slowly. Traditional IR Drop investigations are augmented by 'Dynamic IR Drop' in recent years, employing intra-cycle load current profiles (current spikes at clock edges, in other words) in order to isolate 'peak-noise' values rather than just the multiplied product of an average current and power path resistance. The availability of intra-cycle current profiles enables a move to total power integrity investigation that includes all electromagnetic aspects of the interaction of load currents with power network attributes. In TPI analysis, the impedance of a power network is fully represented in addition to power path resistances, providing a complete picture of the response of the network to excitation stimuli such as load currents.

Power distribution network Impedance and Voltage Droops

A lossy transmission line is often represented by its characteristic impedance as $\sqrt{\frac{R+j\omega L}{G+j\omega C}}$. In a nanoscale SoC, the power distribution network and integrated circuits can be abstracted into a similar characteristic impedance, with resistance and inductance of power bus wires forming R, L respectively, and leakage conductance of distributed semiconductor devices as well as their intrinsic capacitance forming G and C. Given that leakage is highly non-linear with voltage, a simpler model without G is more useful in understanding the behavior of a power network to stimuli changing network voltage. This model considers the resistance and inductance of power bus wires and a distributed capacitance

corresponding to the density of integrated semiconductor devices. When excited by a step current source, the transient response of such a network is given by:

$$V_f = V_i - \left(\Delta I \sqrt{\frac{L_p}{C_d}} \right) \sin(\omega_1 t) e^{-\alpha t} - I r_s$$

where V_i is the initial voltage, ΔI is the current step, L_p is the effective loop inductance of the power path, C_d is the distributed capacitance, $\omega_1 = \frac{1}{\sqrt{L_p C_d}}$ is the natural oscillation frequency, $\alpha = \frac{r_s}{2L_p}$ is the damping factor, with r_s being the effective series resistance in the power path.

There are two distinct parts to the time-domain response of the power network to a change in stimulus. One is the '**transient**' first part that contains a damped sinusoidal response, and the other is a '**static**' or DC response. In a more complex power network, the transient part may contain multiple sinusoidal responses at different natural frequencies corresponding to distinct filter stages in the network.

The static portion of the response is what traditional IR Drop tools evaluate. Under the assumption that power bus resistance is by far the dominant impedance encountered by supply currents, traditional IR Drop analysis suffices to determine voltage variation in different portions of the power network. With the inclusion of peak instantaneous currents into the analysis in 'Dynamic IR Drop', maximum noise values are also inspected with some accuracy.

But the assumption that resistance and resistive voltage drops dominate breaks down for wide, low-resistance, global power distribution in an SoC, where resistance tends to be a small component, and loop inductance tends to dominate given the sparseness of the distribution network. In this scenario, voltage variation in the power supply is a combination of both transient and static aspects of response to a stimulus, with transient behavior seen primarily in the wide, low-resistance buses, and static drop seen in the high-resistance, low-inductance, dense low-level interconnect. *This behavior leads to a partitioning of an SoC power distribution network, where both static and true-dynamic (transient) analyses are necessary, and the possibility of optimization of metal usage between the two partitions ensuring lowest overall voltage variation.* The maximum amplitude of transient response is given by:

$$\Delta V = \Delta I \sqrt{\frac{L_p}{C_d}}$$

which represents **dynamic voltage droop**. Given that transient response is a sinusoidal superimposed variation, a corresponding overshoot also manifests, which may affect reliability and lifetime of the integrated semiconductor devices. *Inclusion of dynamic voltage droops and overshoots provides total power integrity*; these aspects of power network behavior are distinctly absent in IR Drop analysis.

The quantity $\sqrt{\frac{L_p}{C_d}}$ represents an approximate power distribution network impedance that determines dynamic voltage droop. This impedance may be diminished in any region of an SoC by appropriate addition to capacitance, or reduction in inductance through power bus and distribution architecture, thus reducing voltage droops and enhancing total power integrity.

Dynamic Voltage Droop: getting better, or worse?

A simple derivation to understand the progression of dynamic voltage droops with scaling is helpful to recognizing a critical need for total power integrity verification in nanoscale ULSI design.

Consider a **Roots of Two Scaling** [1] scenario:

- Capacitance-per-unit-area, C_a , scales by $\sqrt{2}$, operating voltage scales by $\frac{1}{\sqrt{\sqrt{2}}}$, frequency scales by $\sqrt{2}$, and chip area scales by $\frac{1}{\sqrt{2}}$.

Following this constant-power scaling direction, and inspecting the change in Voltage Droop in a unit area (/ua) of integrated silicon, we get:

- Since C/ua scales by $\sqrt{2}$ and voltage scales by $\frac{1}{\sqrt{\sqrt{2}}}$, ΔI scales by $\sqrt{\sqrt{2}} * \sqrt{2}$.
- Assuming the effective L/ua doesn't change, $\sqrt{\frac{L}{C}}$ reduces by a factor of $\frac{1}{\sqrt{\sqrt{2}}}$.

Multiplying ΔI and $\sqrt{\frac{L}{C}}$ in the scaled process generation, the voltage droop amplitude:

$$\rightarrow \Delta I \sqrt{\frac{L}{C}} \text{ scales as } \sqrt{\sqrt{2}} * \sqrt{2} * \frac{1}{\sqrt{\sqrt{2}}} \text{ or by a factor of } \sqrt{2}.$$

Note that this droop trend calculation is based upon filter step responses that are independent of the operating frequency, and is different from discrete $L \cdot di/dt$ noise calculations [1], because it considers the droop absorbing impact of distributed capacitance as well. Nevertheless, it is clear that **transient voltage droops are increasing** at least by the inverse of the process scaling factor in constant-power scaled designs. Additionally, per unit area, $\Delta I * r$ drop increases by $\sqrt{\sqrt{2}} * \sqrt{2}$, assuming effective resistance does not change in the scaled process, which is not true since higher frequency currents tend to crowd within low-inductance pathways and thus do not take advantage of full metal widths of the power bus wires.

Some engineering directions stand out from these total power integrity methodology studies: 1) employ **total power integrity** investigations for nanoscale ULSI chips, 2) minimize power bus loop inductance values, 3) optimize and balance metal usage between global and local power distributions in SoC's, and 4) address second order effects stemming from high-frequency currents induced by faster switching edge-rates in nanoscale processes: don't let skin effect become an unanticipated issue.

Optimal on-chip power network design

Simple examples simulated in [\$\pi\$ -fp](#) illustrate a beneficial aspect of total power integrity investigation, the trade-off between wire width and power bus spacing in the global power distribution network.

The simulation netlist and stimulus current profile employed in an example whose results are shown in Table-1 follow below. This example has been run on Anasim's [\$\pi\$ -fp](#) [6], a tool developed for total power integrity studies and power integrity aware floorplanning. Note that such qualitative studies on TPI are carried out very early in the design cycle, and are meant to assist IP Core and chip designers in arriving at an initial floorplan that is correct by design and optimized for metal resource usage, block placement, and power network noise.

Netlist

```
.TRAN 100e-12
.PLOT 10
.ACC 0.0040
Gchip1 0.5 0.5 0.0010 0.0100 0.020 10e-9 10e-9
Ichip1 0.2 0.2 0.05 0.05 ramp.txt 1
```

ramp.txt

```
0 0
100e-12 0.5
```

In the above simulation input, the single line 'Gchip1 0.5 0.5 0.0010 0.0100 0.020 10e-9 10e-9' describes a power grid that is 5mm by 5mm in area, power grid wires of 10 μ width and 100 μ bus pitch, 20m Ω wire sheet resistance, and distributed electromagnetic properties of 10nH/cm wire inductance as well as 10nF/sq. cm capacitance. *This simplicity in power grid description, unavailable in traditional SPICE, is key to 10X faster true-electromagnetic simulations in [\$\pi\$ -fp](#)'s 'what-if' analyses.*

A current source region 50 μ by 50 μ in area, located at (2mm, 2mm) in the physical layout and ramping up to 500mA in 100ps as defined in file 'ramp.txt' is used as stimulus in the experiment.

Table-1: Maximum voltage droop with power grid dimensions

Width (μ)	Pitch (μ)	dVmax (mV)
10	100	223
20	100	215
40	100	211
10	50	132
10	25	76

Results from this qualitative experiment indicate that increasing wire width has minimal benefit in total power integrity that considers on-chip inductive voltage droops along with the traditional IR Drop. On the other hand, decreasing the power bus pitch with thinner wires provides substantial benefits in grid voltage noise. In accordance with this simple result, published research on power network optimization and [the impact of on-chip inductance on power distribution network design](#) [4] claims as much as 30% improvement or reduction in metal area usage for power grids in the 90nm process, and approximately 60% improvement in the 45nm process with accurate modeling of on-chip power grid inductance.

Published research also indicates that this type of on-chip inductive effect becomes significant for load current edge rates at or below 100ps. In this example noise was reduced by increasing the number of transmission lines in the periodic grid. We find that an increase in the number of parallel inductors (with approximately the same inductance) is an effective way to reduce overall grid impedance in systems with high edge rates, and, in general, SPICE power grid analysis is much more difficult when

inductance can no longer be ignored. Noise propagates out from the source at a velocity given by:

$$v = \frac{1}{\sqrt{LC}}$$

where L is inductance per unit length of transmission lines in the grid and C is capacitance per unit length. This capacitance can be very large on-chip as it includes the capacitance of every component connected to the main current carrying bus (including the capacitance of many other wires on the lower layers). In a uniform, periodic grid, the capacitance per unit length is given by $C = sC_A$ where s is the periodic space (bus pitch) and C_A is capacitance per unit area. Thus, wave velocity is given by

$$v = \frac{1}{\sqrt{sLC_A}}$$

This means that, when we reduce noise levels by increasing the number of grid wires (reducing the pitch s), we also change the overall noise distribution by increasing the wave velocity. We can see this effect by modifying the load current in the previous example to the profile shown in figure 1. [π-fp](#) simulation results using this profile are shown in figures 2 and 3. Again, we see significant overall noise reduction when the same metal area is used for a larger number of parallel transmission lines. However, the reduced grid impedance also has the effect of increasing the rate at which the noise wave front spreads out across the chip. This wave combines (is superposed) with waves propagating from other sources to produce the overall chip noise. Waves from each source interfere constructively and destructively in different parts of the grid. This type of noise distribution can not be predicted using a static analysis or dynamic IR drop tool as it is the result of complex interactions between edge timings, block placement, grid impedance and variable wave velocities (caused by variable capacitance per unit area in different parts of the chip). Again, this type of analysis is best done using tools that perform realistic simulations in minutes rather than hours or days. This makes it possible to complete the large simulation based experimental designs, covering the large parameter space (capacitance, width, space, placement etc.) required for this type of complex power grid design, within a reasonable timescale.




Figure1: Triangular load current profile for grid noise investigations

Viewed from another perspective, critical metal resources for routing and compact SoC physical design can be freed through total power integrity analyses of chip floorplans and power grids. Moving beyond IR Drop based design, we can not only improve verification of on-chip noise, enabling low energy design, but also optimize the SoC floor plan, improving routing and corresponding signal integrity, saving area, cost, and potentially, total design effort by minimizing synthesis and floor plan iterations.




Figure 2. On-chip power noise distribution for grid wire width 20um and periodic space 100um.




Figure 3. On-chip power noise distribution for grid wire width 10um and periodic space 50um.




Figure 4. Superposition of power noise from two sources (wire width 10um, pitch 50um)

Figure 4 shows superposition of power network response to two loads. Here, cumulative voltage droop in between the two load regions is seen to be substantially higher as compared with figure 3. Additional capacitance is added in the region marked C1, and voltage droop and noise propagation can be seen to have reduced here. Such effects, and electro-migration by current crowding through low-inductance, higher-resistance pathways (with low resistance, longer loop, global pathways in parallel) are entirely missed out in IR Drop based or other analyses that ignore inductances or time/space distributed loads.

Moving beyond IR Drop to Total Power Integrity

Power integrity related failures tend to fall into the occult basket as far as most digital designers go. In part, this is because power supply non-idealities have not been as clearly quantified as signal-to-noise ratio in signal integrity, and partly because most circuits and chips are designed to exhibit very good PSRR, or power supply rejection ratio, minimizing attention paid to power supply integrity. It used to be common to specify a power supply with $\pm 10\%$ variation, with power simulations and verification

conducted at a system or package level assuming a certain power consumption profile in the chip supported by such system or package. Such luxuries are no longer available in the nanoscale regime. Chip-Package co-design has become a gating aspect of product design and development. Power supply voltage values are specified with a much tighter variance band, in part because achieving performance with increasing device variations is that much more difficult with substantial supply voltage variations. In other words, PVT methodology is seeing so much 'P' variation with no reduction in 'T' variation that the budget for 'V' variation is that much lower. Energy consumption is also a critical and product-differentiating design constraint; the one dominant variable that impacts energy is voltage, which must be reduced to the absolute minimum possible. That can only be done if the noise band anticipated within the chip embedded in the system is very well understood and controlled.

There is clear and present motivation to move to advanced, true-electromagnetic, total power integrity investigations. There also is clear and present danger that if such a transition is delayed, there could be significant yield impact in nanoscale SoC's that could be hard to debug, and even harder to correct absent comprehensive understanding or verification methodology.

Nevertheless, the EDA industry is known to move cautiously. Dynamic IR Drop, for example, is said to have taken more than 5 years before being accepted as a necessary aspect of chip power integrity and decoupling capacitance sufficiency verification. In the case of on-chip inductance inclusion, the barrier is even higher, given that all of today's device extraction techniques focus on polygonal R, C and L extraction, and L extraction requires an understanding of current return pathways which makes it all the more challenging. Advanced modeling techniques such as [PEEC](#) have shown some promise, but simulations with any large number of inductors tend to take many hours or days, with convergence and accuracy problems, and constructing multiple PEEC-based models for different grid architectures is very time-intensive. Therefore, despite multiple research results having shown that including on-chip inductance and $L \cdot di/dt$ noise results in significantly better optimization of metal resources ([3] and [4]), the EDA industry has been hesitant to move forward with this essential advancement to front-end total power integrity verification and optimal floorplanning.

Fortunately, solutions to ease of inductance inclusion and rapidity in early 'what-if' analyses have been developed. One needs only to understand that trade offs exist between degrees of freedom and design/verification capability just as it is with optical proximity correction and restricted physical design rules, or with nanoscale device sizing and predictability of device parameters. Employing restricted design and symmetry rules to power network synthesis as well, one can not only ensure low loop inductances, but also develop simplified representations of partitioned on-chip power networks that facilitate greatly reduced computational complexity for true-electromagnetic (Maxwell's field equations based) dynamic simulations. This is the underlying technology in Anasim's power integrity aware floor planner, [\$\pi\$ -fp](#), that represents power distribution grids as surfaces with simulation complexity independent of the number of wires or circuit/capacitance blocks within the areas represented.

From a methodology transformation perspective as well, [\$\pi\$ -fp](#) provides a non-disruptive pathway to benefiting from comprehensive, true-electromagnetic simulations on block and chip power distribution networks and floorplans. Working with conceptual floorplans, power grid dimension specifications, limited process information (sheet resistance and capacitance), and intra-cycle block-level load current profiles available from pre-processors to Dynamic IR Drop analysis, this tool provides spatio-temporal, dynamic, total power integrity information that permits optimization of power distribution grid dimensions, and circuit block / decoupling capacitance physical layout and placement. Optimization results may then be manually input into the physical design flow, with static IR Drop analysis detecting and eliminating any current or electro-migration hot-spots as physical layout matures.

The future could see true-electromagnetic simulations replace Dynamic IR Drop – or should it be now?

References

- [1] Raj Nair, "[A Power Integrity Wall follows the Power Wall!](#)", SOCcentral.com article, March 2008
- [2] Raj Nair & Donald Bennett, "[Power Integrity and Energy-Aware Floorplanning](#)", SOCcentral.com featured article, January 2008
- [3] A. Muramatsu, M. Hashimoto & H. Onodera, "[Effects of On-chip Inductance on Power Distribution Grids](#)", ISPD 2005
- [4] N. Srivastava, X. Qi and Kaustav Banerjee, "[Impact of On-chip Inductance on Power Distribution Network Design for Nanometer Scale Integrated Circuits](#)", ISQED 2005
- [5] Kendall Castor-Perry, Genuum Corp., and Tamara Schmitz, Intersil Corp., "[Know the sometimes-surprising interactions in modelling a capacitor-bypass network](#)", Planet Analog 10/15/2007 article
- [6] Anasim [\$\pi\$ -fp](#), Power Integrity and Energy aware Floorplanner, [Anasim](#) Corp., 2008 release
- [7] Richard Goering, "[Inductance nags, tools lag in 0.13-micron ICs](#)", EETimes 09/28/2001 article

Author Biographies



Raj Nair is an expert consultant in IC power delivery and power integrity, and co-founder of [Anasim](#) Corp. developing and marketing [\$\pi\$ -fp](#). Prior to Anasim, he founded ComLSI, Inc. that develops and licenses patents and silicon IP in advanced power delivery as well as high-speed signaling. Formerly, he was with Intel, where he conceived and implemented distributed on-chip voltage regulation in microcontrollers and championed a fully integrated CMOS voltage regulator μ -processor power delivery solution. He holds 36+ issued patents and has translated many of his inventions to commercial products.



Donald Bennett is a device physicist, IC design engineer, EDA algorithms and software developer, and co-founder of Anasim Corp. Prior to Anasim, he directed Modeling and EDA at ComLSI, Inc. and developed and marketed a power grid simulation tool Picosim (also known as RLCSim) at Quantum Design Automation in the UK. His corporate experience includes many years at ST Microelectronics. Donald is the inventor of the 'effective current density' (ECD) modeling method that is the core technology enabling Anasim's [\$\pi\$ -fp](#).