Power Integrity and Energy aware Floor Planning

Raj Nair & Donald Bennett, Anasim® Corp. (www.anasim.com)

Roof Planning, Rooms and Walls in SoC's

We've heard so much about floor planning for integrated circuits – routing, timing awareness, and even leakage and temperature awareness; how often do we come across the term Roof Planning in SoC's? Yet, just as the foundation of any integrated circuit is its substrate, the roof that ensures that the various functional blocks perform as required is its power grid and power delivery system. While much attention has been focused upon the substrate through investigations of substrate noise conduction, little has been done on power grid and power delivery analysis other than the simple exercise of determining $i \bullet r$ drop that is substituted for comprehensive power integrity investigation. As integration continues in accordance with Moore's Law, and nanoscale SoC's include ever more functional blocks operating with switching edge rates of fast fabrication processes capable of multi-GHz operating frequencies, on-chip noise now includes significant $L \bullet di/dt$ content, or dynamic noise, propagated across the chip by the roof above that has two orders of magnitude lower resistivity or impedance as compared with the foundation below. How does one decide what is the girder width of such roofs, optimizing metal usage? How does one determine how best to position the rooms of this integrated mansion so as to minimize total noise? How does one determine how best to place minimal walls of isolating decoupling capacitors so as to protect sensitive rooms from the din created by the entire system? How does one determine how low must the roof be, so as to minimize power consumption that is quadratically (active power) and exponentially (leakage) dependent upon the roof height or the operating supply voltage? This paper details how these challenges may be addressed in SoC floor planning.

Low Power and the dominant design constraint

Extreme performance and frequency are no longer dominant design goals for SoC's. Recent years have witnessed right hand turns in microprocessors, with multi-GHz unicore CPU architectures abandoned in favor of low-frequency multi-core variants, and frequency becoming a forgotten memory. SoC's are also treading into the Ghz domain, but cautiously; great care is taken to ensure that a single dominant design constraint is not violated under any operational condition. So what is this most sacred aspect? It is the one exponential that has hit a wall for integrated circuits (IC's) of all varieties – Power. While scaling of transistor dimensions and the corresponding number of transistors per unit area continues along its 'Moore's Law' exponential for the next decade, power, which used to double roughly every 36 months for microprocessors of a specific architecture, can no longer do so. Power is the single dominant design constraint for SoC's today.

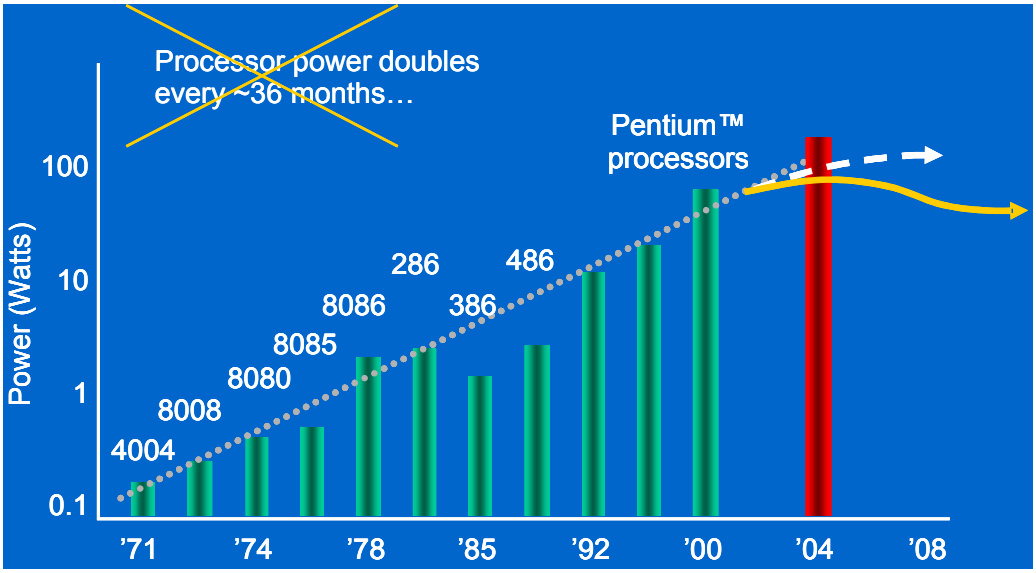

Figure 1: The exponential that hit a wall: Power

Not Frequency, but Voltage

Low power design is a key quality measure for any ASIC or SoC today. Minimizing active power may be accomplished simply by reducing the frequency of operation in CMOS IC's, since active power is proportional to frequency. But that does not benefit most systems, since it is energy that is the principal concern. Energy consumed directly relates to battery life in portables and handheld devices where a majority of SoC's are found today. Energy is Power • Time, and therefore lowering frequency, while reducing power, does not reduce energy used since the same task now takes more time for execution. Additionally, leakage, or static, wasted power consumption in CMOS is now about the same as active power, and it is minimally influenced, if at all, by the frequency at which a block is operated. The longer the block remains operational, more is the leakage static power consumed, and therefore more is the energy wasted. Hence the solution to minimizing energy consumption IS NOT reducing FREQUENCY, but IS reducing the VOLTAGE at which the SoC is operated, since both active and leakage power are strongly dependent upon operating voltage. Nanoscale CMOS processes offer an interesting opportunity in being able to reduce operating voltage while still maintaining the necessary frequency of operation [1] because of the near linear relationship of transistor saturation drive current to the applied control voltage.

Bringing the roof down: analysis complexity hits a wall

Reducing and optimizing operating voltage levels needs to be done with full cognizance of required chip performance and necessary noise margin. Simply put, operating voltage must be high enough that even with anticipated reduction due to on-chip noise, the chip and its functional blocks perform as specified under all process and voltage variations. At the same time, the operating supply differential must be the lowest that it can be made, while ensuring that critical pathways in high-performance blocks see minimal reduction in supply differential at all times. In other words, it is essential to have **spatial** and **temporal** awareness of **true dynamic noise** in the chip power grid. This is where currently available tools for power integrity face difficulties in providing rapid answers.

We believe this inability of current tools in providing sophisticated power integrity information rapidly stems in large part from how integrated circuits are physically represented within tools, in the form of

polygons. Polygons do not allow us to take advantage of any of the abstract aspects of physical layout that greatly simplify electrical analysis. Whereas finite element analysis takes specific advantage of this method, through representation of arbitrary shapes as being constituted by triangular or polygonal shaped structures, such methods require enormous computing resources in order to evaluate even static variations, much less their dynamic progression. As chips move into the nanoscale regime, the complexity of the analysis, with objects to be analyzed represented as polygons, explodes exponentially and also hits a wall. Hence tools that deal with physical objects on an integrated circuit as polygons resign themselves to doing justice to simple power integrity analyses such as $i \bullet r$ drop simulations.

Simplifying and speeding up analysis while improving accuracy

On the other hand, it is possible to take maximum advantage of the 'differential', 'symmetric' and 'periodic' nature of power grid design on an integrated circuit. In other words, one can comprehend both the physical aspect of power grid wires (and functional blocks and decoupling capacitance) as well as their principal electrical and architectural characteristics and employ these aspects into simplifying the analysis problem. For example, power grid wires comprise of metal of a given thickness and width. And power buses are always drawn out differentially for any high-frequency or broadband supply design. Most electromagnetic interaction of power bus wires is limited to being between the true and complement wires of a power bus pair. These aspects translate to power buses being represented as transmission lines, with frequency dependent inductance, resistance and capacitance captured during the input of the power grid into the analysis schematic. Such representation simplifies a long power bus or 3-D power grid that may consist of a very large number of polygons into a single, simple transmission plane with a single equation describing the relationship of charge and voltage along the structure.

A tool, **π-fp** [2] developed by the authors goes farther in its representation of the power grid, functional blocks, on- and off-chip decoupling capacitances and the entire power delivery stack. This is done through the fusing of the entire power grid into a single surface through a patent-pending 'effective current density' or ECD modeling method [3] that eliminates the specific construction of the grid from determining analysis complexity. In so representing the power grid, a distributed representation of the current consumption of any functional block, spread out within the area it occupies as it is in reality is facilitated, as also the capture of on-chip decoupling capacitance as a distributed capacitance array. These innovations bring the analysis to being very close to true conditions within a chip, providing a comprehensive, accurate view of noise and power integrity within the chip. Such capability combined with an order of magnitude faster analysis speed for the entire power delivery stack allows designers to bring the roof down, optimize supply voltage to the lowest possible level, thereby achieving the best low-power design feasible. With the ability to inspect noise distribution across the chip area and over time, designers can also make trade-offs between power integrity and low power, ensuring that their floor plans are such that critical paths in the chip are well protected from supply transients.

Advanced analysis: does capacitance really reduce noise?

In the nanoscale regime, where gate leakage is high and chip area is scarce, optimal use of on-die capacitance in reducing or deflecting noise is critical. Incorrectly configured capacitance arrangements can actually amplify noise in nanoscale SoC's. To investigate such effects, a tool that explores the electromagnetic, wave nature of on-chip noise is essential.

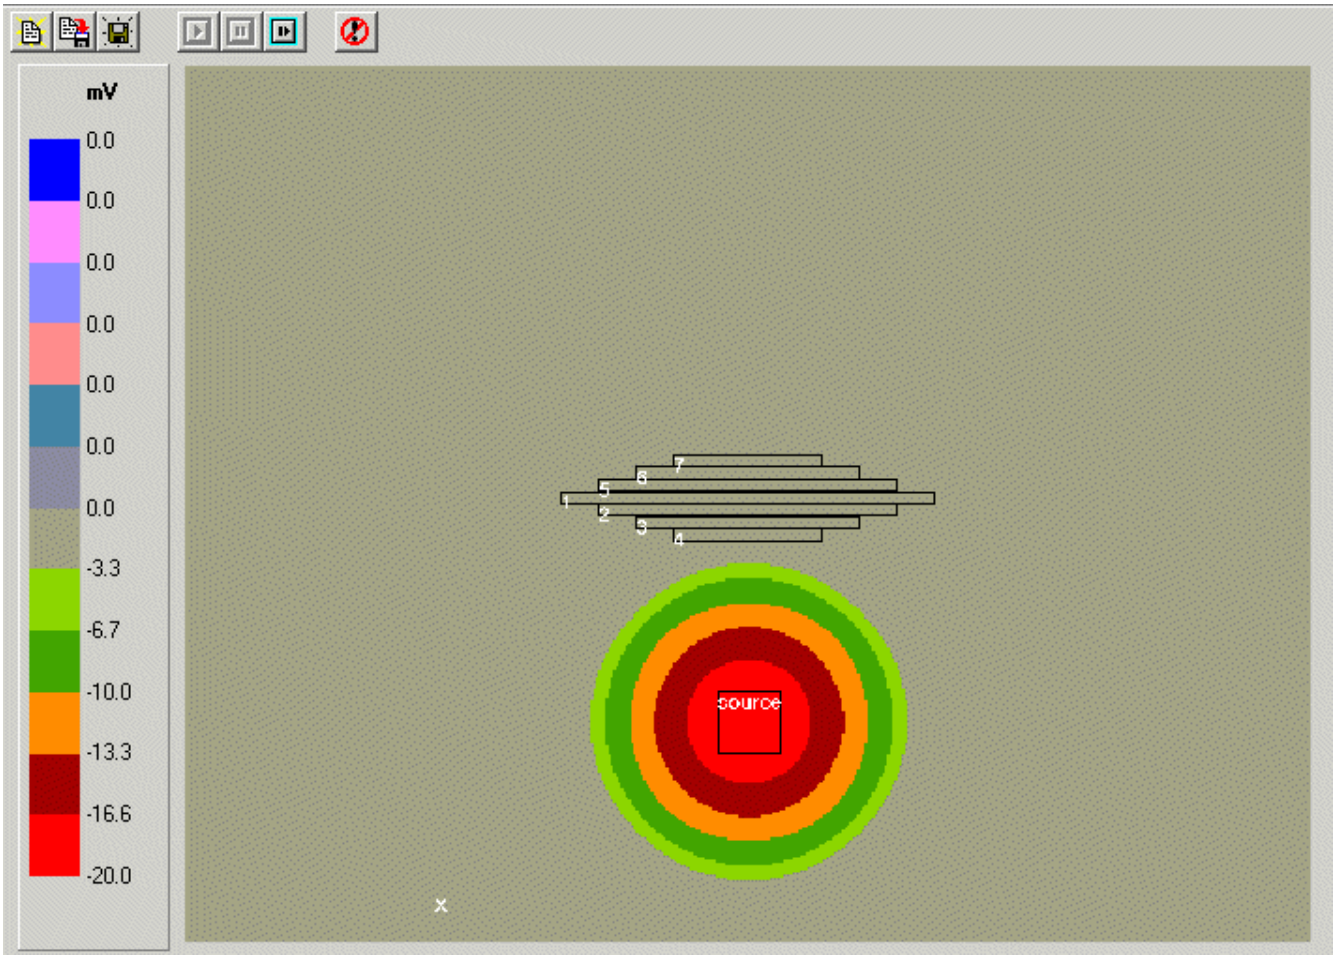

Figure 2: Noise propagation and interaction with a CAP array

Figure 2 shows a time dependent solution to a grid equation derived using the effective current density method (ECD, [2]). The plot shows the differential noise distribution across a 9mm x 7mm integrated circuit. Noise is injected into the power grid from the rectangular load current block at the bottom center of the image (100mA - pulse width 100ps). This noise propagates out from the noise source into a region containing distributed decoupling capacitance (5nF/cm2).

A lens shaped region is defined in the middle of the IC. This region contains additional capacitance. The electromagnetic wave velocity is slowed in this region (wave speed in a transmission line is given by v~=1/sqrt(LC) where L and C are inductance and capacitance per unit length).

The capacitance array (objects 1 through 7 in the picture) is seen to 'focus' the noise wave emanating from the load current block onto a region above it, creating a transient region of high noise in an area supposedly 'shielded' by the added capacitance. This effect seen through a 'true-EM' simulation is very similar to another more commonly observed effect, that of focusing of light through an optical lens.

This simulated solution illustrates how additional capacitance added to a power grid can result in both noise suppression and noise amplification in different parts of an IC. This is in stark contrast to the effect of decoupling capacitance in a low frequency system where any additional capacitance results in noise suppression only. Such systems are dominated by RC delay and rapid wave amplitude attenuation. High speed systems such as those in the nanoscale regime are dominated by transmission plane and LC resonance behavior that result in complex, difficult to predict, noise distributions.

Optimizing the girders of the SoC Roof

Routing metal resources are quite often scarce in SoC's. Power delivery grids take up a substantial amount of metal in at least two metal layers, and quite often in 4 layers in very complex SoC's such as microprocessors. Determining exactly how wide power bus wires need to be is therefore a critical optimization task.

Once again, $i \bullet r$ drop tools may not provide reliable indications in this analysis space. A significant part of supply noise in nanoscale SoC's comes from $L \bullet di/dt$ drops. The L here is not packaging inductance as is often incorporated by many tools, but is the local loop inductance in the power delivery system to the nearest reservoir of charge that supplies the local transient charge demand. In a regular SoC power grid, this is typically the inductive aspect of the wire pair in the local power bus. Fast edge rates and synchronous functionality in nanoscale SoC's produce very large, local di/dt values within a chip resulting in noise that is far greater than what an $i \bullet r$ drop analysis may estimate it to be.

Simulation studies using **π-fp** [2] have shown that increasing the power bus wire width provides noise reduction to a minimum noise floor determined primarily by $L \bullet di/dt$ drop. There is hence a maximum wire width beyond which increasing width provides no further benefit for a particular grid architecture. Noise also increases rapidly for wire widths substantially lower than this maximum. This ability to determine total noise through rapid what-if analyses enables designers in correct optimization of the use of metal resources for power delivery and integrity within the SoC.

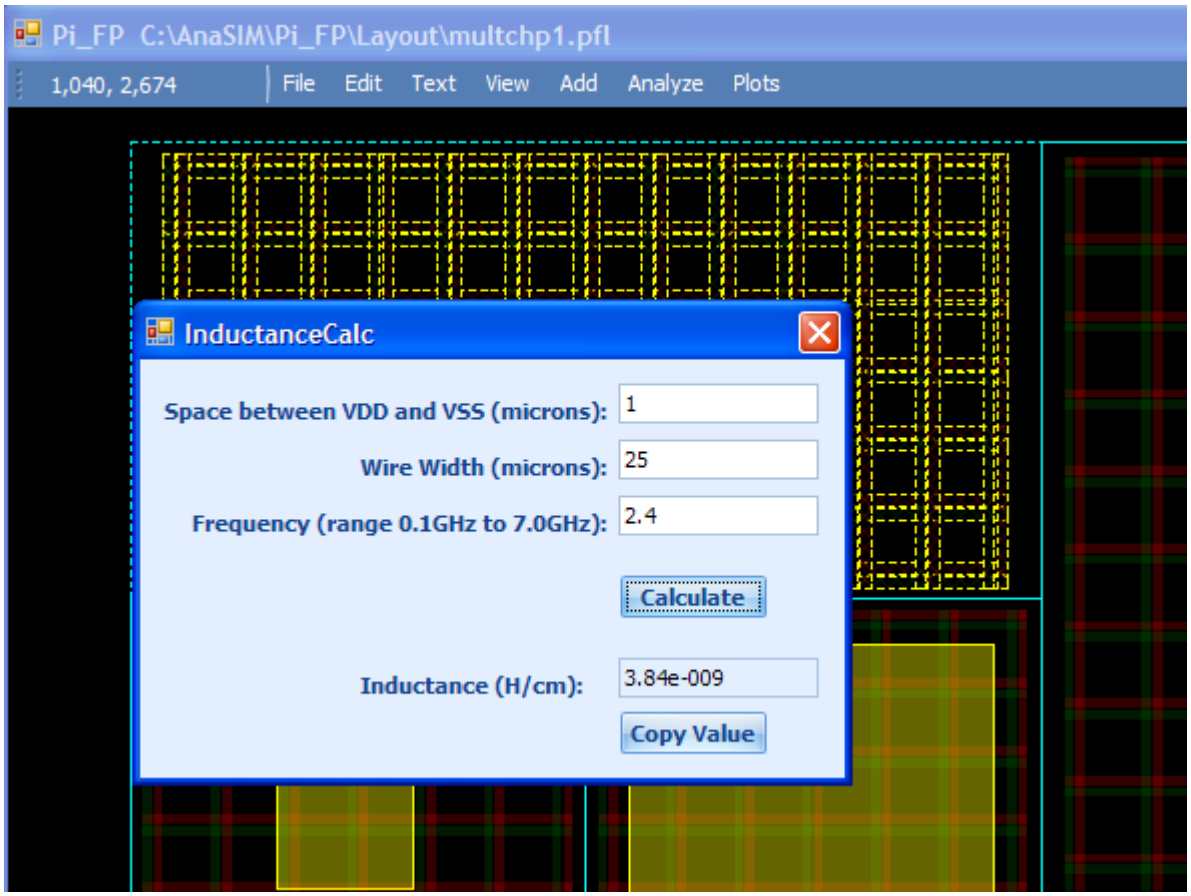

Figure 2: π-fp power grid inductance calculation

Advanced analysis: noise propagation and the effects of power gating

It is important to recognize that noise is not a point event, where a given instantaneous current combines with a local impedance to generate a local $i \bullet r$ drop. Just as noise is propagated through a semiconductor substrate, it is very effectively conveyed by the SoC power grid that is two orders of magnitude lower in impedance as compared with the substrate. While substrate noise progression occurs due to capacitive and resistive effects, power noise propagation follows L and C characteristics of the grid.

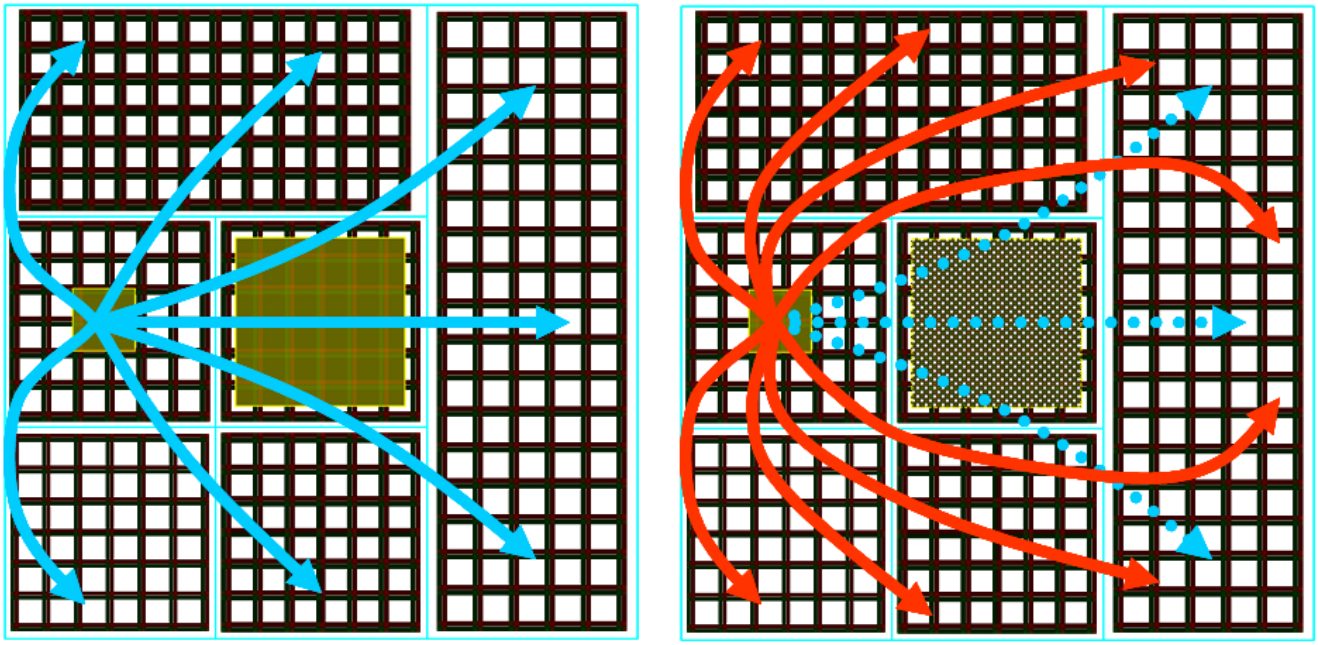

Figure 4: Nose propagation without and with power gating

Figure 4 illustrates the impact of power gating upon an SoC floor plan. Power gating is a common technique, akin to clock gating, employed in nanoscale SoC's to minimize leakage and wasted power in blocks that temporarily do not require to retain their operating state. But its impact upon power integrity is less understood. A power gated block presents much higher loop inductance and becomes ineffective in distributing noise or providing capacitive balancing for grid inductance. In simple words, noise is forced to crowd around a power gated block, thereby increasing dynamic noise in adjacent active circuit blocks.

Understanding such implications in an SoC is critical to comprehending the sum of static and dynamic noise components in order to arrive at an optimal floor plan. Such analyses helps designers eliminate uncertainties as to total noise content, minimize supply voltage margins effectively, and thereby arrive at the lowest power and energy design for the SoC.

Concluding thoughts

Power grids tend to be over designed, and floor plans under-analyzed, in designers' attempts to compensate for unknowns in the system. Complex noise distributions in systems dominated by capacitive and inductive effects are not so easily suppressed by simply adding metal or by adding capacitance. Rigorous investigation with advanced tools that provide pre- and post-synthesis power integrity what-if analyses will be key to arriving at low power, low energy SoC's.

A design methodology that increases the designers understanding, early in the design flow, of dynamic and total noise, makes achieving important design objectives such as reliability and lower power consumption easier. Designers need to get to 'power noise by design' instead of an over engineered system that may or may not work at tape out.

A designer would not consider a functional block to be complete if its output as a function of input was not completely understood. Why should power analysis be any different? Total noise in the power delivery system should be considered to be an important design objective, since voltage is most definitely now a design variable and an energy conservation controlled parameter.

As with many other aspects of architecture and design, a methodology and advanced tools will be required to comprehensively analyze and minimize noise in SoC's. Advanced power integrity analyses and pre-layout floor planning will provide much additional information to complement information obtained from static analyses or other types of simulation such as SPICE or fast SPICE derivatives. An increased understanding of system behavior as a whole, early in the design cycle, will make it easier to achieve important objectives such as reliability, low energy consumption and first silicon success.

References

1. Raj Nair, “Power Delivery, Integrity Analysis and Management in Nanoscale SoC/SiP”, International Symposium on System-on-Chip SoC 2007, November 2007 http://www.anasim.com/SoC07_Anasim_Raj_Nair.pdf
2. **π-fp**: Power Integrity aware Floor Planning; Anasim Corporation, www.anasim.com
3. Donald Bennett, “Effective current density and continuum models for conducting networks”, US Patent Application 11/714,427