



System-level Power Distribution Noise Closure: Looking Beyond the SoC Power Integrity Challenge

Aveek Sarkar

Apache Design Solutions, Inc., San Jose, CA, USA

Notice of Copyright

This material is protected under the copyright laws of the U.S. and other countries and any uses not in conformity with the copyright laws are prohibited. Copyright for this document is held by the creator — authors and sponsoring organizations — of the material, all rights reserved.

DESIGN AUTOMATION CONFERENCE

WHITE PAPER: Power Closure

System-level Power Distribution Noise Closure: Looking Beyond the SoC Power Integrity Challenge

Aveek Sarkar

Apache Design Solutions, Inc., San Jose, CA, USA

Abstract — This article takes a comprehensive look at the subject of power, and reviews its various facets, such as power estimation and optimization, power grid noise prediction, power delivery network design and power-induced failure mechanisms, both on the die and in the system. Through the use of various examples of prior work done and to be presented in DAC 2010, this article attempts to tie together all these seemingly disparate topics into an integrated topic of system-level noise modeling, whose challenges confront IC, package and PCB engineers as they consolidate ever-increasing functionality requirements in their designs through advanced circuit, layout and processing technologies.

Index Terms— Power grid noise, power optimization, PDN design and fixing, system-level design.

I. INTRODUCTION

The emphasis on power as a requirement for design closure is increasingly apparent. However, power by itself is a rather general term, with different meaning and requirements depending on where in the design phase it is being considered. Most commonly power is defined as the amount of current the chip draws during its operational mode. For chips used in handheld non-tethered devices, standby current is a key metric. Reducing the amount of current consumed by devices during their operational or standby modes is now a key engineering requirement, primarily to extend the battery life of the device. In addition to the current drawn, the power dissipated by the devices needs to be reduced.

The reduction in current consumption is achieved through several micro-architectural and design (register-transfer or circuit) level changes. The reduction in power dissipation is achieved through the reduction of the current drawn and through the reduction of the operating voltage of the chip. However, reducing the nominal supply voltage of a chip affects its performance, which is often compensated through the use of circuit, layout and process technology changes. These techniques, such as the use of fast libraries available in advanced process nodes, or the use of heavily-gated and highly-tuned clock tree structures, can adversely impact the quality and stability of the voltage seen by the devices. These circuits manifest sudden spikes in the current they require, as a result of the temporal and spatial localization of on-chip switching activity. Voltage regulators must provide this current through the board, package and on-die power delivery network (PDN). These highly inductive networks prohibit the sudden supply of current, even if the circuits demand it; hence, this current must be supplied primarily by on-die decoupling capacitance. The possible unavailability of on-die charge from these decoupling capacitances, coupled with the resistive nature of the PDN mesh, forces the voltage seen by the devices to be significantly reduced from what is intended. Figure 1 illustrates the scenario of voltage droop occurring due to a high transient current demand event.




Figure 1: Impact of high switching current on power supply integrity.

Today's circuit designers are thus faced with the following decisions and tradeoffs:

- Reduce the current drawn by the circuits, both in operational and in standby modes.
- Reduce the power dissipated, through a combination of reduced current and lower supply voltage.
- Achieve similar or higher performance at this (reduced) voltage by using circuit, layout and fabrication techniques (which in turn may adversely impact the quality of the supply voltage).

The following table outlines the steps designers must take to achieve these goals:

Engineering step	Impact
Power optimization	Reduces the overall consumption of current drawn and power dissipated in the chip
On-die PDN optimization	(a) Low resistive and inductive drop (b) Optimizes on-die decap to reduce high frequency switching noise (c) Reliability requirement compliance (EM/ESD)

Package and board design	<ul style="list-style-type: none"> (a) Robust power delivery, with minimum voltage fluctuation, using the fewest number of layers and discrete components. (b) Manage thermal and mechanical stress requirements, ensuring that the die or die(s) are not adversely affected (c) Meet EMI/EMC regulatory requirements
--------------------------	--

These design choices and analysis steps are not only targeted at designs that go into low-power handheld devices, but also impact high-performance designs that go into servers and desktop platforms. The problems created in these designs – from the power drawn by the chip, and from power delivery through board, package and on-die PDN – are no longer just the chip designer’s issues. Rather, choices must be made that impact all the different aspects of making a system function properly. The design tradeoffs undertaken affect not only the power delivered to the device, but also affect its long-term reliability, its thermal signature, its ability to transmit signals, and its ability to meet EMI/EMC regulatory requirements. Thus “power” is no longer an issue for chip designers to solve but for all the parties in the design of the system to comprehend and address from the beginning of the design process. Power as a problem is significant enough for it to be addressed as “a system-level noise problem” affecting the design, performance, manufacturability, and ultimately the marketing of the system.

This article outlines key challenges associated with each of these topics, as well as ongoing work that target these challenges. Section II discusses the need to reduce power consumption by devices, as well as available solutions. Section III discusses power delivery network design techniques and reviews several available modeling and analysis methods. The consumption of current by the chip is a source of noise in the system, which is propagated through the chip, package and board PDN. Section IV discusses the impact of this propagated noise on system-level timing and EMI radiation. The article concludes by outlining some of the upcoming challenges in this area of work.

II. OPTIMIZING CHIP POWER CONSUMPTION




Figure 2: Trend in leakage versus dynamic power through various technology nodes [1].

Current consumed on a chip can be broken down into two categories: operational and standby power. Both of these are on an increasing trajectory (Figure 2) and need to be reduced, especially if the device is operating off a battery. Traditional low-power approaches have mostly targeted circuit and layout changes through the use of clock and power gating, voltage islands, and similar techniques. However, these changes, other than appearing too late in the design process, may not address inherent limitations of the design, or its architecture, in reducing its power envelope. Hence a power-aware design methodology should consider power as a design target even as early as the architectural definition phase. In addition, a power-aware methodology should be employed throughout RTL design and continuing into the layout implementation phase. The strategy must employ rigorous analysis and

verification techniques to ensure that such design and layout changes do not adversely affect the circuit operation.

By starting early in the RTL stage of the design process the designer can ensure that (a) power is a key consideration in the design process, and (b) optimizations otherwise not possible without significant changes in a synthesized netlist are achieved early on. Figure 3 summarizes the trend seen from a study done on several designs [2], showing that the biggest impact on a low-power design is seen by targeting changes early in the RTL phase, with potential power savings being as high as 80%.

A power-aware methodology should cover the following design steps:

- Micro-architectural level power prediction and exploration.
- Incremental refinement of RTL through block-level and full-chip analysis and optimization.
- Support for design partitioning for power and voltage islands.
- Full-chip RTL power analysis.
- Verification of power consumption in the final gate-level design.




Figure 3: Steps in design phase where power reduction is possible [2].

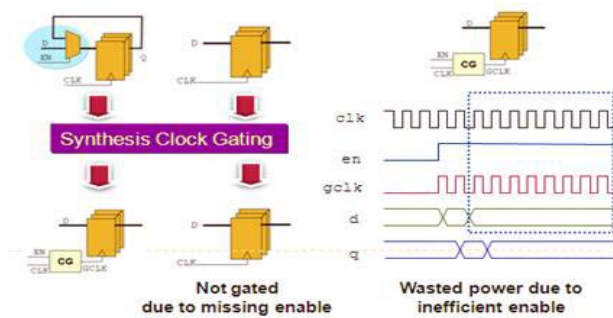


Figure 4: Issues in clock gating that cause wasted power.

EDA solutions that target such design requirements must provide multiple levels of functionality in addition to the desired accuracy. Tools such as the PowerArtist-XP™ RTL power analysis and optimization platform from Apache are used by design teams to obtain insight into the power consumption of their circuit at the RTL stage, and to optimize power through multiple techniques that target not only the clock tree network, but also the memory and datapath portions of a design. For example, Figure 4 shows how improper clock gating may result from an incorrect RTL specification, or how power consumed

in a register may be wasted at the clock pin when the input does not transition. Figure 5 shows how using an XOR-OR of the input along with the flop output to clock-gate a register locally can help reduce power consumption, provided that the cost (area/timing/power consumption impact) associated with this change can be justified.




Figure 5: Low-activity non-enabled register to reduce wasted power.

A power optimization solution must be based on an accurate estimation and analysis engine that provides insight as to where power savings can be achieved. For example, PowerArtist-XP leverages the PowerTheater™ engine to provide accurate power estimation to drive an analysis-driven optimization at all stages of the design. At the RTL netlist stage, advanced circuit construction (datapath and clock tree) and wire load prediction algorithms are used to mimic a gate-level design and provide accurate power estimation. This subsequently helps to achieve

circuit optimizations and “power bug” isolation. It also allows for the early exploration of different power and voltage island configurations to reduce the power signature of the design. Once the suggested design changes are deemed useful, an automatic RTL rewrite option provides the designer with an updated netlist description that is optimized for current draw and power dissipation, and that will prevent “power surprises” later in the design phase. After a gate-level netlist becomes available, PowerArtist-XP allows the designer to verify the predicted savings. Depending on the type of circuit/library used, the experience of the engineering teams, and the design techniques used, the power savings using RT-level optimizations can range from 10% to 60% of the original power.

III. MODELING AND ESTIMATING POWER SUPPLY NOISE

An accurate analysis of on-die PDN dynamic switching voltage drop noise must consider several effects from the die: switching and leakage current from the devices, the presence of various forms of decoupling capacitance (intentional decaps, diffusion and gate capacitance, signal wire load capacitance, and well capacitance), and power grid parasitics (resistance, inductance and capacitance). The package and board parasitics need to be included in this analysis as well. In this article, the RedHawk™ (SoC Power Noise and Reliability analysis) and Totem™ (custom circuit Power Noise and Reliability analysis) platforms from Apache will be used to highlight the technologies available when addressing the challenges associated with solving full-chip dynamic power noise issues [3,4].

The design, analysis and optimization of the power delivery network ensure that supply voltage noise is controlled. These steps also ensure that the low-power design techniques being utilized do not adversely affect the operation of the circuit. Low-power consumption is often achieved by operating parts or all of the design at a reduced supply voltage. On top of this, low-power design techniques such as clock and/or power gating can introduce significant levels of dynamic voltage drop (DvD) noise if not properly designed. Since the combination of lower operating voltage and heightened noise levels can impact the performance and functionality of the chip, an accurate estimation and prediction of voltage drop when using low-power design techniques is of considerable importance.

Clock gating is intended to reduce the amount of power consumed by non-essential registers and circuits. However, the power savings achieved through clock gating may not be fully achievable due to the introduction of clock buffering required for timing and clock skew constraints. In addition, the increased number of buffers firing simultaneously induces high dynamic voltage drop, especially if these buffers are not placed properly (that is, they all share the same power and/or ground rails).

Designers also need to pay careful attention to the design and implementation of power gates, since these may not only affect design performance and functionality, but also may fail to provide the standby mode leakage current savings they were expected to deliver. As noted in several prior works, improper sizing of the power gates can cause either high leakage in standby mode or cause high voltage drop or reliability issues during the operational mode. During transition of the circuit from standby to operational state the process must be carefully controlled to ensure that the rush-current and/or the noise coupling to an already on-circuit is minimized. As noted by Yong et al. [5], the noise fluctuation during a “un-gate” process (seen from simulation and measurement studies) can be quite significant. Additionally designers need to ensure that the variation in voltage across a power gated circuit is tightly controlled – otherwise parts of a circuit may reach full-rail compared to other parts, causing contention and high crowbar current. Figure 6 illustrates such a design condition automatically identified using the Explorer capability in RedHawk wherein during the power-up process the receiver gate powers on before the driver does, causing a short-circuit current condition.




Figure 6: Short-circuit current condition in a power-gated design.

The magnitude and shape of the on-die switching current determines the severity of voltage drop and power noise and is directly related to the temporal and spatial distribution of the switching activity and device placement in the chip. For example, if the clock buffers are tightly clustered to minimize skew, their simultaneous switching will cause localized hot-spots due to high current demand in those portions of the die. On the other hand, during scan testing, simultaneous firing of a large number of registers can induce high current demand during a very short period of time, leading to the collapse of the power grid voltage. These current demand spikes with mostly high frequency components can only be addressed by on-die decoupling capacitance.

However, decaps placed specifically to reduce PDN noise can provide only a limited amount of charge. Additionally, as seen in Figure 7 which plots the ESC (effective series capacitance) and ESR (effective series resistance) of a decap cell in 90/65/45nm, normalized to the value of the ESC/ESR of the same cell in 130nm, the efficiency of these intentional decap cells is reduced due to the combined effect of reduced cap (from the use of thicker oxides to reduce leakage), and higher intrinsic resistance. This increases the time it takes a decap in 45nm to respond to a demand for charge. Most of the on-die decap comes from the other parasitic capacitances from other devices that may not be switching, or may have the charge available. The charge from these decaps flows to the switching devices through the on-die PDN. R and C from the on-die PDN have been traditionally modeled. However, fast device transition times in 40nm and below technologies make inclusion of on-die PDN inductance (L) necessary, especially for higher metal layers. Figure 8 highlights the difference in the worst dynamic voltage drop with and without the inclusion of on-die PDN inductance for a 40nm SoC. As noted by Ergin et al. [6], the two dominant effects that require the inclusion of PDN inductance are: (a) large currents and the di/dt of such currents seen in the higher level metals, especially in the RDL (redistribution layer), and (b) the inductance of such structures.




Figure 7: ESC/ESR of an intentional decap cell through several technology nodes normalized to ESC and ESR of the same cell in 130nm.

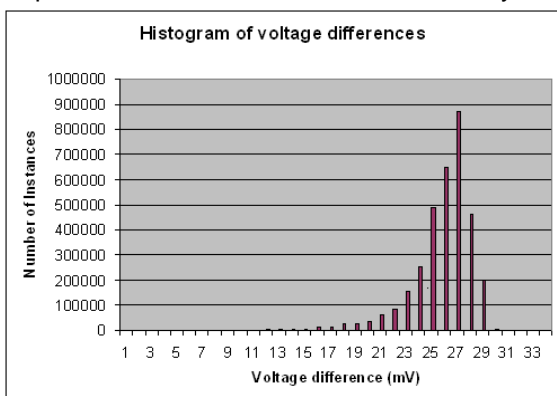


Figure 8: Histogram comparing worst drop difference with and without considering on-die inductance for a full-chip SoC design.

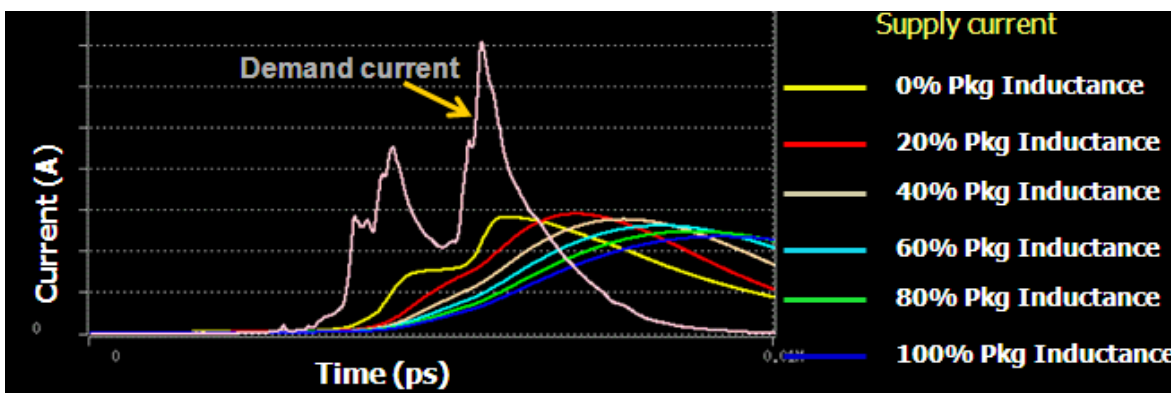


Figure 9: Impact of package inductance on supply current. The white curve shows the total current demanded by the switching circuit. The other colored (yellow, red, white, cyan, green, and blue) waveforms represent the current supplied by the regulator when impeded by different levels of package inductance. As the inductance value increases, the current flows even more slowly, requiring more charge to be supplied by the chip. Even without any package inductance (yellow curve), the regulator does not supply all the demand current, with the on-die decap providing most of the charge required. If there is not enough charge on the chip, the supply voltage collapses until the regulator/package-board decaps provide the charge at a later time.

On-die voltage drop analysis is incomplete if the package and PCB geometries are not considered in the simulation. These structures not only cause increased drop but also provide additional “redistribution” layers for current flow. The board parasitics are seen to impact the dynamic noise in some cases by 20% of the total drop. The drop across the package can be as much as 70-80% of the total drop, depending on the package layout (L) and the on-die switching (di/dt) signature. Figure 9

shows how the current flowing from the supply is modulated based on the value of the package inductance. As the current flowing through the package slows down due to increasing L, the voltage drop at the switching cells increases unless the on-die capacitance has sufficient charge available.

A critical aspect in the modeling of on-die power noise is the switching scenario or the excitation that is used to stimulate the circuit. The use of gate-level test-benches is prohibitively expensive, both in terms of generation and usage, especially for larger designs. Most design teams use tools such as the RedHawk VectorLess™ engine, which provides them with an alternative method to model on-die switching. The RedHawk VectorLess engine uses statistical techniques, along with design parameters such as exclusions, logic properties, and toggle activities, to create switching scenarios to reflect on-die operations. Unlike static (or average) analysis, where all cells in the design draw some current (scaled by some toggle activity), a dynamic analysis reflects a true switching operation in which some cells' transition drawing the full current they need for the transition (for example from 0 to 1 in their outputs), while other cells maintain their states. In many cases, design teams use the activity from RTL test-benches to control the RedHawk VectorLess engine to switch the circuit in the specific mode they want. The RTL-based PowerArtist-XP engine can prune a large RT-level test-bench through a multi-level filtering procedure and identify the specific clock periods or groups of periods that highlights voltage drop-induced issues in the circuit. Using the activity information at the registers and primary inputs from the selected clock periods, RedHawk can derive a switching scenario for the circuit by propagating the state through the logic.

Traditional vectorless approaches have focused on having power (as consumed by the switching circuit) as a controlling parameter when generating the switching scenario on the chip. Since the goal of this analysis is to predict voltage drop-induced failure scenarios and to address them through chip, package and board redesigns, advanced modes of RedHawk VectorLess take a system-level approach when determining the appropriate on-die switching scenario [7-9]. Figure 10 shows results from one such study [8]. Also as noted by Dobre et al. [9], the consideration of the system PDN resonance as a criterion of the switching scenario creation helps to provide a mechanism to perturb the system under resonance and also helps to identify areas of high frequency localized drops in the chip.

CPM #	FFT Sum	Avg Cur	Drop (p2p)
CPM3.3	1.782135	0.61	345
CPM3.4	1.473784	1.302	358
CPM1.3	1.2549734	2.91	141
CPM1.4	1.247208	2.93	143
CPM2.1	0.8116995	2.28	107

Figure 10: In an ASIC targeting a networking application, increased voltage drop is seen when the energy from the on-die switching activity is concentrated around the resonance frequency of the system. In this study, several chip power models (CPMs) of the chip were created, each representing a unique switching scenario of the chip [8].

Each CPM was simulated along with the package and board layout, and dynamic noise was measured at the C4 bumps. For each operating mode (or CPM), the energy around the resonance frequency (FFT sum), the average switching power (Avg Cur) and the worst dynamic drop (p2p) are shown [8].

As noted earlier, the design, optimization and verification of the chip PDN must be performed in conjunction with the design and optimization of the package in which the chip will reside, and the PCB on which the packaged chip will be mounted. Since the extraction and modeling technologies used for on-die geometries (static extraction) and package/PCB structures (quasi-static/full-wave extraction) are different [7], a model-based approach is used to analyze these designs separately. This methodology involves providing accurate models of the package and board to the chip team, and of the chip to the package and board teams. Several technologies exist to create accurate models of the package and board structures. Most package structures have fragmented routing to accommodate multiple signal and power/ground nets in few package layers. To accurately model these shapes, 3D full-wave extraction technologies are increasingly required to capture fringing effects, the impact of imperfect ground planes, and boundary reflections. The next generation of 3D full-wave extraction tools leverage advanced meshing and multi-CPU solver technologies to rigorously solve Maxwell's equations and generate accurate broadband models. Chen et al. detail the use of one such tool in a chip-package-system analysis that included a six-layer package, a 14-layer PCB and a 45nm chip targeting mobile handset applications [10].

The accuracy of the chip models used by package and PCB designers to optimize their designs has been significantly improved through the use of Chip Power Model (or CPM™) technology [11]. CPM creates a SPICE representation of the chip, with ports created at the C4 bump/pad locations (or inside the chip as necessary) to capture both activity (current) and parasitics (RLC) of the chip. The generation of an accurate model of the die requires the following tool support: (a) capacity to handle an entire chip layout along with device information, (b) extraction of the on-die PDN parasitics and coupling, (c) inclusion of all device and other capacitances, and (d) reduction of this large netlist (with 100+ million nodes for most designs) into a behavioral model that can be interpreted and simulated by SPICE engines. Despite such a reduction, the model must provide the accuracy of the full-chip layout that it represents. A common technique to validate the accuracy of the model is to simulate the layout of a chip with its package model and compare the bump-level current and voltage waveforms to a Spice simulation of its equivalent die model (or CPM) with the same package model. Figure 11 shows the result from this “self-consistency” check, indicating that very close results to the full-layout database can be achieved using a highly reduced SPICE-based netlist description provided by a CPM of the chip. In most cases, the results are within 5-10% of silicon-measured numbers. Compared to heuristic-based approximations, CPM provides silicon-correlated die models [12] that allow package and board designers to validate and optimize their designs.

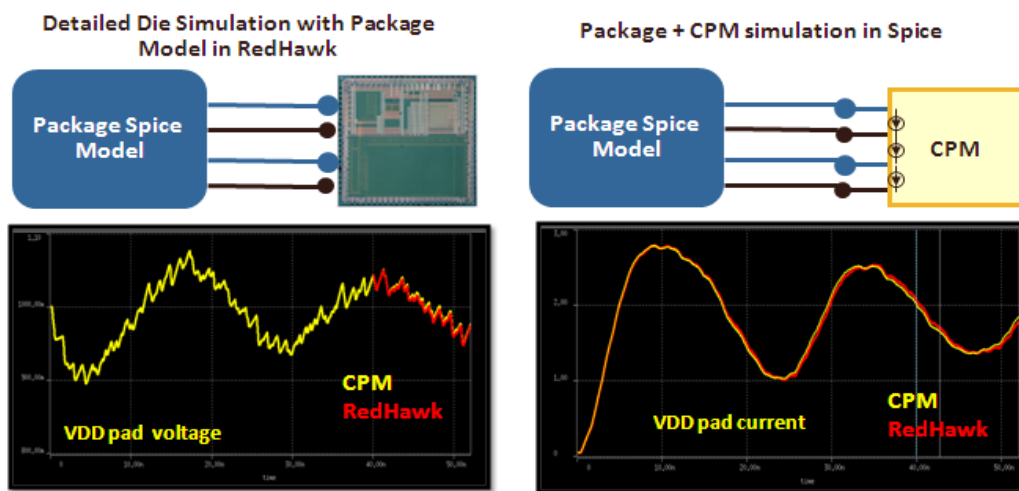


Figure 11: “Self-consistency” check indicating close match in pad voltage and current between full-layout-based (in RedHawk) and CPM-based (in SPICE) simulations with a package model.

CPM can be used in multiple ways for PDN verification, given that it provides both parasitic and activity information for the die. The former enables impedance analysis of the system PDN, while the latter enables DC and time-domain simulations. The activity on the chip can be generated using VCD and/or vectorless techniques that consider the resonance of the system [7-9], providing system PDN engineers with die models that can stress their designs. As noted earlier, figure 10 shows the impact on dynamic voltage noise from customizing on-die switching scenarios with a system-level focus, giving designers more flexibility in testing their chip-package-PCB designs [8].

For circuits using smaller technology nodes such as 40nm and 28nm, the concern about the long-term reliability of the wires and vias is well-founded. However, the traditional focus on power delivery network electromigration using “average” or DC current flow is limiting. Advanced techniques and EM reliability rules need to be considered, such as those pertaining to I_{RMS} , Blech and I_{PEAK} effects, when designing circuits in these technology nodes. Huard et al. detail the use of Totem™ for establishing a methodology for performing extensive back-end degradation mode analysis [13]. Static and dynamic power noise simulations need to be performed in addition to analyzing the current flow in signal nets (especially in the clock tree and in custom designed circuits [4]). It is often necessary to model the impact of temperature gradient across the chip, given the severe impact of temperature change on

electromigration current density limits. Design teams may prefer to incorporate a distributed temperature profile over their chip, instead of a single pessimistic number, to be able to waive the EM violations they are seeing in areas of lower temperature.

Additionally, devices fabricated in these technology nodes targeting consumer applications are increasingly susceptible to ESD event-induced failures. The circuits must be made immune to ESD charging (HBM/MM) and discharging (CDM) events. Circuit and layout analysis must be performed to ensure the proper design, placement and connectivity of clamp and protection circuits. As noted by Johari et al. [14], comprehensive analysis is needed at the SoC level to ensure that the circuit meets ESD design guidelines. However, most design teams do not use verification technologies to sign off their chips as far as ESD protection is concerned, instead relying on manual plot checks. By using full-chip layout based ESD verification technologies, chip engineering teams can make sure that their ESD protection circuits are optimally designed, inserted and connected in their designs.

IV. IMPACT OF SUPPLY NOISE ON CHIP AND SYSTEM PERFORMANCE

Power supply noise not only affects the functionality and performance of the chip, but also impacts the way it interacts with its environment and other chips. Inside the chip, power supply noise can affect the launch path differently than the capture path, thereby reducing timing slack from the value that was calculated in the absence of such noise. Voltage degradation is the most common cause for parts not achieving volume shipments at higher speeds. Additionally, when dynamic voltage drop noise affects hold-time margins, the failure of the circuit is catastrophic, requiring circuit and/or layout changes through a mask re-spin. In previously published studies, 1% voltage drop degradation caused about 1% impact on timing through increased delay. In smaller technology nodes such as 45nm, similar drop levels have larger impacts, in the range of 1.5 – 2% delay degradation.

Static timing analysis (STA) tools leverage the added functionalities in libraries (as in the CCS models) to model the impact of dynamic voltage drop noise. RedHawk for example provides for every cell in the design different voltage drop parameters, such as min-avg-max drop from its dynamic voltage drop analysis. Depending on the location of the cell in a specific timing path, the appropriate voltage value for that cell is taken to model the degradation effect. Additionally, SPICE-based critical path or clock tree simulations can be performed by incorporating the dynamic noise waveform from RedHawk analysis. Dynamic voltage drop noise is not constant from one cycle to another, being affected by the difference in switching activity between the cycles and by the LC resonance impact of the package/PCB on the chip. Due to the fluctuating supply voltage, signals, especially those belonging to the clock tree network, have different arrival times from one clock period to the next. As noted by Grenet [15] clock tree simulations using PsiWinder™ is necessary to identify clock tree jitter induced failures. PsiWinder incorporates the impact of varying power supply noise (introduced by the combined effect PDN RLC and on-die switching) on a clock tree behavior through a Spice based analysis of the clock tree network.

Off-chip, the impact of dynamic power noise can be even more significant. As noted by Murugan et al. [16], comprehensive analysis is needed during the I/O placement and package design stage to address the concerns associated with high-speed memory interface signal transmission in the presence of power grid noise. In their study, the authors used the Apache Sentinel-SSO™ tool-suite to simulate a 45nm high-speed memory interface by incorporating on-die, package and board parasitics (signal and power), the current demand of the switching circuit, and the associated capacitances. The tool takes in the I/O ring layout for the memory interface, SPICE data (netlist/models), and switching test-bench, along with the package and PCB layouts, to create a unified simulation setup, which captures the impact of power grid noise and signal cross-talk on the propagation of signal from a DDR2 interface to a memory chip.

The accuracy of the analysis comes from consideration of the impact of the simultaneous switching of the entire I/O bank (such as a 32-bit LPDDR2 interface) on the I/O ring power supply (modeled in full RLC detail), and the inclusion of broadband S-parameter based models of the package and board parasitics that capture not only the power/ground and signal net parasitics by themselves, but also model their cross-coupling to each. The benefit of this modeling technology is illustrated in Figure 12, in which the results from Sentinel-SSO pre-silicon analysis tracks silicon measured signal waveforms [16].




Figure 12: Correlation of simulated results to measured waveforms for a propagated signal in a DDR interface [16].

Another area of concern, especially for automotive and consumer applications, is the increasing impact of power/ground noise on the EMI signature of the system. The chips in a system act as the generator of noise, which is then propagated outside the chips through their package and board power and signal traces, and radiated by the cabling. Knowledge of the chip noise signature enables accurate near and far field simulations, providing options to change the system to meet EMI requirements before the silicon is fabricated and the system put together. As noted by Steinecke et al. [17], use of layout-based power noise analysis tools (in this case RedHawk) to generate full-chip models like CPM to provide simulation results provide this early visibility. With this methodology, system engineers do not need to wait for fabricated parts to predict the EMI signature for their system as they have done traditionally. Additionally, if a system fails EMI requirements, presently changes are mostly done at the PCB and system-level. However, design teams can also use RedHawk and/or Totem to understand the frequency components of the on-die noise and the distribution of this noise over the chip. This analysis can be performed both for digital and custom/analog designs through time-domain simulations and by profiling the current signature noise in the various wires that provide the pathways for conduction and radiation-based emission. Redesign of the PDN, along with the optimization of the on-die switching, can address many of the noise profiles that contribute to EMI failure in the first place, helping the system to meet its regulatory specifications.

V. UPCOMING CHALLENGES AND CONCLUSIONS

As design teams explore more advanced structures, such as 3D IC (using through-silicon-vias or TSVs), the issues caused by power are only getting exacerbated. Analysis of such 3D stacked die structures can be performed either in a concurrent mode (*with all layouts simulated at the same time*) or in the model-based mode (*with one or more dies coming in as models such as a CPM*) [18]. As die model creation becomes standardized through the generation and use of CPM, the sharing of data between different design teams (inter- and intra-company) for both chip-package-system and 3D IC analyses will become more streamlined. Additional model requirements exist for signal integrity and EMI simulations. Better models for I/O circuits can help account for power noise coming from the combined impact of simultaneous switching and chip/package/PCB inductances as in a DDR interface. Also, more advanced modeling of the chip for comprehensive EMI simulations is needed. These models need to include power/ground and signal line noise, along with substrate noise coupling, to provide a complete noise signature for the EMI simulations [17].

As seen as this article, power is not an issue that can be addressed by any one engineer. Power and its associated challenges have to be addressed holistically in every step of the design process, from the product definition step to the final design signoff stage by the respective owners. Nor can power as an issue be partitioned. It needs a comprehensive chip-package-board modeling, analysis and optimization methodology to ensure that the design objectives of power consumption, circuit performance, reliability and regulatory compliance are met.

REFERENCES

- [1] N.S. Kim, T. Austin, D. Blaauw, T. Mudge, K. Flautner, J.S. Hu, M. J. Irwin, M. Kandemir and V. Narayanan, "Leakage Current: Moore's Law Meets Static Power", *IEEE Transactions on Computers*, Vol. 36, No. 12, December 2003, pp. 68-77.
- [2] Cadence Design Systems, "A Practical Guide to Low-Power Design: User Experience with CPF", *Cadence Design Systems whitepaper*, May 2008.
- [3] A. Sarkar, "Power Noise Analysis for Next Generation ICs", *Apache Design Solutions whitepaper*, June 2009.
- [4] A. Shanmugavel, "Power and Noise Integrity for Analog / Mixed-Signal Designs", *Apache Design Solutions whitepaper*, May 2009.
- [5] L. K. Yong, F. Tan and C. S. Lee, "Power Noise Mitigation Strategy from RTL Perspective on MTCMOS Design", *DAC User Track presentation*, DAC 2010, to appear.
- [6] E. Ergin, T. Todesco, M. Frankovich, F. Guo, G. Wong, R. Sequeria, and T. Yogarasa, "Delivery Network Design in Presence of Inductance and Dynamic IR Drop Analysis", *DAC User Track presentation*, DAC 2010, to appear.
- [7] A. Shayan, K. Bowles, S. Dobre, M. Popovich, X. Chen and C. Pan, "Resonance-Aware Methodology for System-Level Power Distribution Network Co-Design", *Proc. IEEE Electrical Performance for Electrical Packaging*, October 2009, pp 29-32.
- [8] W. Cheng, A. Sarkar, S. Lin and J. Zheng, "Worst Case Switching Pattern for Core Noise Analysis", *DesignCon 2009*.
- [9] S. Dobre, A. Shayan, M. Popovich, K. Bowles, X. Chen and C. Pan, "Package/PCB Aware On-Die Power Grid Noise Analysis", *DAC User Track presentation*, DAC 2010, to appear.
- [10] X. Chen, B. Patra, S. Dobre, P. Vennam, and M. Elmore, "Block-Level Analysis of Chip and System-Level Resonances", *DesignCon 2010*.
- [11] E. Kulali, E. Wasserman and J. Zheng, "Chip Power Model - A New Methodology for System Power Integrity Analysis and Design", *Proc. IEEE Electrical Performance of Electronic Packaging*, October 2007, pp. 259-262.
- [12] X. Liu and Y. Liu, "The Extraction and Measurement of On-Die Impedance for Power Delivery Analysis", *Proc. IEEE Electrical Performance of Electronic Packaging*, October 2009, pp 195-198.
- [13] V. Huard, R. Chevallier, C. Parthasarathy, A. Mishra, N Ruiz-Amador, F. Persin, V. Robert, A. Chimento, E. Pion, N. Planes, D. Ney, F. Cacho, N. Kapoor, V. Kulshreshta, S. Chopra and N. Vialle, "Managing SRAM Reliability from Bitcell to Library Level", to appear in *Proc. IEEE International Reliability Physics Symposium*, May 2010.
- [14] P. Johari, Y. Liao, N. Chang, and A. Sarkar, "ESD Robustness Verification for System-On-Chip Designs", to appear in *International ESD Workshop*, May 2010.
- [15] V. Grenet, "Clock Tree Analysis in the Light of Dynamic Electrical Effects using Apache's PsiWinder Tool", *DATE 2007*.
- [16] R. Murugan, S. Mukherjee, V. Subramanian, and J. Zheng, "Timing and Noise Analysis of a Simultaneously Switching IO Bank with Silicon Correlation", *DAC User Track presentation*, DAC 2010, to appear.
- [17] T. Steinecke, M. Gokcen, J. Kruppa, P. Ng, and N. Vialle, "Layout-Based Chip Emission Models using RedHawk", *International Workshop on Electromagnetic Compatibility of Integrated Circuits*, November 2009.
- [18] D. Riquet, L. Doyen, J. Antonijevic, N. Vialle, and A. Sarkar. "Analysis of Power Delivery network of Multiple Stacked ASICs using TSV and Micro-bumps", *DAC User Track poster*, DAC 2010, to appear.