Power Grid Analysis Benchmarks

Sani R. Nassif
IBM Research - Austin
11501 Burnet Road, MS 904-6G021, Austin, TX 78758, USA
nassif@us.ibm.com

I. Abstract

Benchmarks are an immensely useful tool in performing research since they allow for rapid and clear comparison between different approaches to solving CAD problems. Recent experience from the placement [1] and routing [2] areas suggests that the ready availability of realistic industrial-size benchmarks can energize research in a given area, and can even lead to significant breakthroughs. To this end, we are making a number of *power grid analysis* benchmarks available for the public. These are all drawn from real designs, and vary over a reasonable range of size and difficulty thereby making studies of algorithm complexity possible. This paper documents the format for the various benchmarks, and give details for their access.

II. Power Grid Analysis

Because of the strong dependence of circuit performance - both delay and power dissipation- on power supply voltage, there has been much work dedicated to the modeling, simulation, and analysis of integrated circuit power delivery systems (see [10, 5, 7, 6] to name a few). These efforts have resulted in computer-aided design (CAD) tools that estimate the power supply voltage delivered to every component of a complete integrated circuit. In such an analysis, it is possible to take into account the on-chip power distribution network, the on-chip decoupling capacitance, the package parasitics, as well as the parasitics associated with the board on which the integrated circuit is mounted.

The reason for the existence of these customized power delivery analysis CAD tools is the complexity of the power grid. To get a sense for this complexity, consider the following:

- In the design of a high performance processor, it would not be unusual to dedicate 10% or so of the overall wiring resources on all wiring levels to power delivery. This heavy investment in wiring is required because it not uncommon for a modern high-performance multi-core processor to dissipate 100 Watts at a 1 Volt supply. Thus a series resistance of $1m\Omega$ results in a voltage drop 100mV which is 10% of the voltage supply, which will typically cause a change of about 6% in frequency and about 18% in power dissipation.
- Assume the processor in question is created in an 8-level metal process, occupies a 1cm $\times$ 1cm area, and uses wires 1μ wide at a pitch of 9μ[1]. This means that every metal level has about 10^3 wires. If we further assume that we will model each *via* (i.e. intersection between metal wires) as a resistor, and the metal segment between two vias as a single resistor, then each metal level will define 10^6 nodes. This results in a power grid circuit with about 8 million nodes and 24 million resistors.

As different components within an integrated circuit process data, thus drawing power from the power delivery (power grid) network, the voltage supply at all points of the network fluctuates. This time domain behavior can be quite complicated, especially with the recent introduction of various power reduction techniques such as clock gating [3] and power gating [4]. Furthermore, for a complex system that includes software-programmable parts, there will be a large dependence of instantaneous power on the actual software program and data being processed. The complexity of these interactions makes it quite difficult to determine the appropriate set of conditions under which to perform a power delivery variability analysis, and has resulted in most of the work in this area being split between two problems: (a) a detailed steady state (DC) problem [5, 7], and (b) a simplified time domain problem [8].

1. In steady state power delivery analysis, we restrict our attention to the DC voltage drop at each component in the circuit, and our goal is to analyze the full network in order to assess the statistical and spatial variations of the power supply. Simple engineering checks can be made to insure -say- that all grid voltages are within specified bounds, and some have even tried to formally show how this can be ensured under various DC loading conditions [9].

2. For the time domain problem, power supply variation tends to be dominated by the resonant interaction between the predominantly capacitive chip and the predominantly inductive package. It is fair to say that these time domain interactions tend to be more global in nature, i.e. with a lower spatial frequency across the chip, and can thus be well estimated from simplified models of the system [10].

[1] Since each metalization level has slightly different width and pitch dimensions, these numbers should be considered to the the average across the complete metal stack.

978-1-4244-1922-7/08/$25.00 ©2008 IEEE

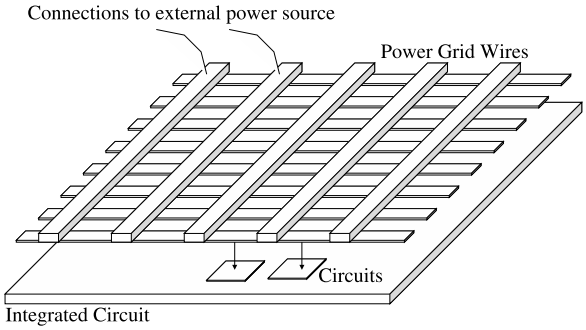

Figure 1: Illustration of a small portion of a typical power grid.

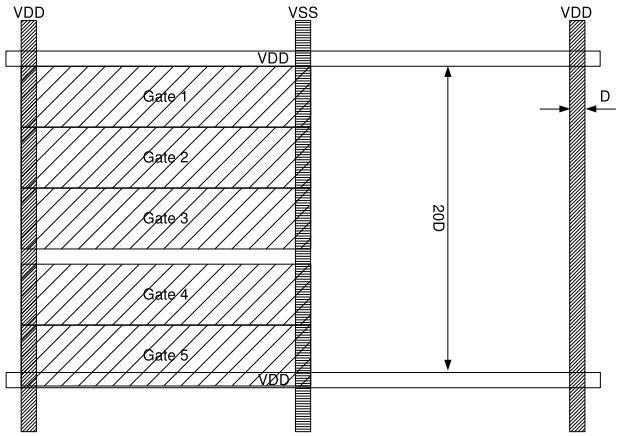

Figure 2: Justification for typical power grid wire aspect ratio.

For this set of benchmarks, we will focus on the steady state (DC) problem. The reason is that the complexity of DC problem is largely independent of circuit activity, and therefore enables us to provide realistic power grid designs without the intellectual property complications that would arise from having to share the other details of integrated circuit. In the near future, and depending on the response from the research community, we hope to extend the benchmarks to include time and frequency domain operation as well.

III. Power Grid Modeling

In this section we will describe the modeling assumptions that were made in producing the benchmarks. The assumptions were used in developing the equivalent electrical model for the power grid, for the integrated circuit, and for the package. The next three subsections describe each of these components.

A. Power Grid Extraction

A power grid is typically composed of an orthogonal mesh of wires, a somewhat idealized representation is illustrated in Figure 1. In a realistic design, the mesh is not complete, i.e. some wires may be missing or truncated. Also, the periodicity and density of the wires may vary, so areas of the chip which require less power may have fewer and narrower power grid wires than areas which are more power hungry.

The translation of the geometric shapes representing the power grid wires to an equivalent circuit useful for simulation is done via *circuit extraction*, which has always been a standard part of the integrated circuit design process. In the context of the steady state analysis of power grids, however, the circuit extraction procedure can be drastically simplified. First, we are only interested in the *resistance* of the power grid, so there is not need to look into the capacitance and inductance of the power grid wires. Second, power grid wires are typically quite long. To understand why that it, let us go back to the assumption that roughly 10% of the wiring resources are used for the power grid. Let us also assume that on any given wiring level, the width and spacing of wires are equal and let us denote that dimension by D. Finally, assume we have a power grid with two nets, V_{DD} and V_{SS} (also commonly referred to as ground). Since one of every 10 wires is a power wire, then one of every 20 wires is a V_{DD} wire. Assuming that the metal level above has the same dimension D, this means that a typical power grid segment will be of width D and length $20D$. This situation is illustrated in Figure 2.

Given the above observations, an efficient scheme for power grid extraction would rely on modeling each intersection as a via, and each via-to-via segment as a single resistor. Such an extraction scheme is significantly simpler than the fully general algorithms used for ”real” circuit extraction, and can be implemented for full-chip design quite efficiently.

Two further simplifications serve to complete the electrical equivalent model of the power grid:

- First, we assume that connection between circuits and the power grid occur at the lowest metal level (obviously), which we will refer to as the **M1** level. The reason for this choice is obvious, since no contact can occur to the devices that make up the integrated circuit except at the lowest level of metal.
- Second, we assume that the connection between circuits and the power grid only occurs at the intersection between the lowest and next-to-lowest level (which we will refer to as the **M2** level). The rationale for this choice has more to do with managing the complexity of the overall problem than anything else. Consider the case where -say- 5 individual gates are connected between the wires (as illustrated in Figure 2). If we were to model the individual connections between the gates and the power grid we would have 5 times the number of nodes.

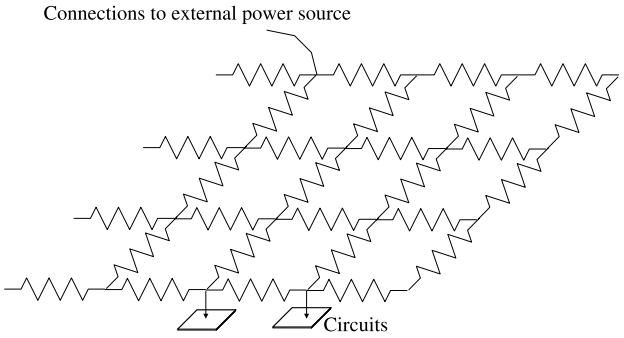

Figure 3: Equivalent circuit of a small portion of a typical power grid.

- Third, and perhaps most controversially, we assume that via resistance can be ignored. This is, in fact, optional and the benchmarks will be made available with and without this assumption. In cases where vias are included, the vias within one intersection are considered to be in parallel and reduced to one equivalent via. The advantage of assuming zero via resistance is obvious, an instant reduction in total node count by a factor of 2!

With these simplification, a typical equivalent circuit for the power grid is illustrated in Figure 3.

B. Load Modeling

Since we are focusing on DC benchmarks, we only need to model the DC loading associated with the integrated circuit. The subject of power estimation, i.e. the science and art of approximating the power consumed by individual portions of an integrated circuit, is one in which much research has been done. An excellent overview is in [11]. Our goal in this benchmarking effort is only to provide sufficiently interesting load models to ensure that algorithms developed to perform power grid analysis are accurate and robust with respect to such models. Also, since the actual load model will be highly dependent on the detailed description of the design, we wish to substitute a simpler but effective equivalent in order to avoid having to share needless design detail.

To this end, we create a scalable load model as follows:

1. Assign a total power for the design, P_{tot}.
2. Grid the integrated circuit into N_x by N_y regions.
3. Generate a $N_x \times N_y$ random numbers associated with each region and normalize them such that $\sum_{i=1}^{N_x} \sum_{j=1}^{N_y} Z_{ij} = 1$.
4. Assign each region a power of $Z_{ij} P_{tot}$, and convert it to a DC current equal to $I_{ij} = V_{DD}^{nom}/(Z_{ij} P_{tot})$.
5. Determine the number of power grid connections (i,e, **M1 M2** intersections) within each region, denote that number by N_{ij}.
6. Apply I_{ij}/N_{ij} at each power grid connection, that is, distribute the power uniformly withing the i, j^th region.

The benchmarks are defined with DC loads as above. For a more realistic model, however, two important factors need to be considered and researchers are encouraged to explore the implications of these extensions to the overall power delivery analysis modeling and analysis methodology.

The first refinement has to do with the fact that I_{DD} depends on V_{DD}. A well known approximation to the power dissipated by a logic circuit is $P \approx \alpha V_{DD}^2 f C_{sw}$, where P is the power, α is the so-called *activity factor* and ranges from zero to one, f is the frequency of operation, and C_{sw} is the total switched (i.e. discharged and charged) capacitance within the circuit. From this simple formula, we see that the current -at constant frequency and switching factor- can be expressed as $I_{DD} \approx \alpha V_{DD} f C_{sw}$. This implies that the current is, approximately, a linear function of the voltage. This model, of course, is only valid for small changes in V_{DD}, but is immediately suggests modeling the current associated with a circuit in the form:

$$I_{DD} = I_0 + G_{DD} V_{DD} \tag{1}$$

which would provide the first order correction to the dependence of power supply current on power supply voltage.

The second refinement addresses the increasing influence of *leakage* current as a component of I_{DD}. For technologies at and beyond the 65nm node, it is expected that the leakage component of total power supply current is 30% or more. Leakage current has a strong super-linear dependence on the power supply voltage [12], but for moderate local changes in V_{DD} we would expect that Eqn. 1 would suffice to include leakage as well as the intrinsic dependence of I_{DD} in V_{DD}.

These two refinements imply that the load model can be modified from a constant DC current source to a current source in parallel with a conductance of value G_{DD}. Due to the dependence of I_0 and G_{DD} on technology, this refinement is *not included* in the current benchmark definitions. It is hoped, however, that researchers will explore these modeling issues further and develop efficient algorithms for including them in the analysis.

C. Package Modeling

Since we are focusing on DC analysis only, the package model (which would normally include a complex system of self and mutual inductances) is reduced to a simple per-connection parasitic resistance. In real packages, these resistances are most definitely *not constant* [13]. For this benchmarking effort, however, they are defined as constant.

Table 1: Benchmark summary descriptions.

Name	Number of Nodes	Number of Elements	Metal Levels
IBMPG1	30638	55109	2
IBMPG2	127238	246581	5
IBMPG3	851584	1603581	5
IBMPG4	953583	1838583	6
IBMPG5	1079310	2156735	3
IBMPG6	1670494	3246725	3

IV. Typical Power Grid Performance Metrics

For any benchmarking effort, it is important to define the expected result of the benchmarks. In the context of the DC analysis of power grids, the following results are required:

- The *voltage* are each **M1 M2** intersection. This is needed in order to insure that all circuits have a sufficiently high voltage to insure proper operation. Under certain conditions, it may suffice to solve an alternate problem which ensures that each of these voltages is above some specified threshold V_{min}. In general, V_{min} is circuit specific.
- The *current density* through each metal segment, defined simply as the current per unit width. This is needed in order to insure that no part of the power grid exceeds electro-migration reliability limits. Similar to the point above, it may be possible to define an alternate problem where the result is simply the assurance that every current in the power grid is below the defined limit, which -of course- is function of the conductor width, but also of the metal level.
- The *current* through each via. This is similar to the item above. Note that in these benchmarks, some circuits are modeled with resistive vias, and others are modeled with shorted vias.

Any proposed power grid analysis algorithm must be able to provide at least some of the results above.

V. The Benchmarks

The current benchmark set includes six designs (more may be added by the time of the final publication of this work). The designs are simply named IBMPG1 through IBMPG6. Table V summarizes the important features of these designs. The table include the benchmark name, the number of nodes in the circuit representing the benchmark, the number of circuits elements, i.e. resistors, voltage and current sources, and finally the number of levels of metals included in the analysis. It can be seen that overall benchmark size varies by about two order of magnitude.

The benchmarks will be provided as *annotated Spice format files*. This was a difficult decision to make, since it should be obvious that the Spice [14] input description language is not the most efficient method to encode circuits with millions of nodes and elements. Nevertheless, it is a well known format requiring the minimum of extra documentation.

Within each Spice file, the user will find additional comments and a specific naming and numbering scheme which will aid in understanding the linkage between the circuit and the original geometry of the power grid. These specific features are documented below.

1. Node Names are of the form:

   `n<net-index>_<x-location>_<y-location>`

2. The data associated with each layer in the file is preceded with a line of the form:

   `* layer: <name>,<net>_net: <net-index>`

   since each layer/net combination is treated separately (i.e. **M1** V_{DD} and **M1** V_{SS} are modeled distinctly of each other in terms of nodes, resistors, contacts etc...) then each layer/net combination is associated with a unique index, which is the `net-index` field on the line above. It also corresponds to the same field on node names.

3. The collection of vias between one layer/net combination and another are preceded with a line of the form:

   `* vias from: <net-index> to <net-index>`

   where the `net-index` field will correspond to the same field on node names. Vias will be implemented as resistors, when modeled as such; or as zero-valued voltage sources when modeled as shorts.

4. Each circuit file will have a global V_{DD} voltage source, and each package connection can be recognized as a resistor connected to that global source. No special comments are placed in the file for package connections.

5. Finally, the current sources representing the loads have the form:

   `iB<block-number> <node> 0 <value>`

   and

   `iB<block-number> 0 <node> <value>`

   Note that each current source is split into two components, from V_{DD} to *ideal* ground, and from V_{SS} to ideal ground. This allows modeling certain situations where asymmetric switching can occur.

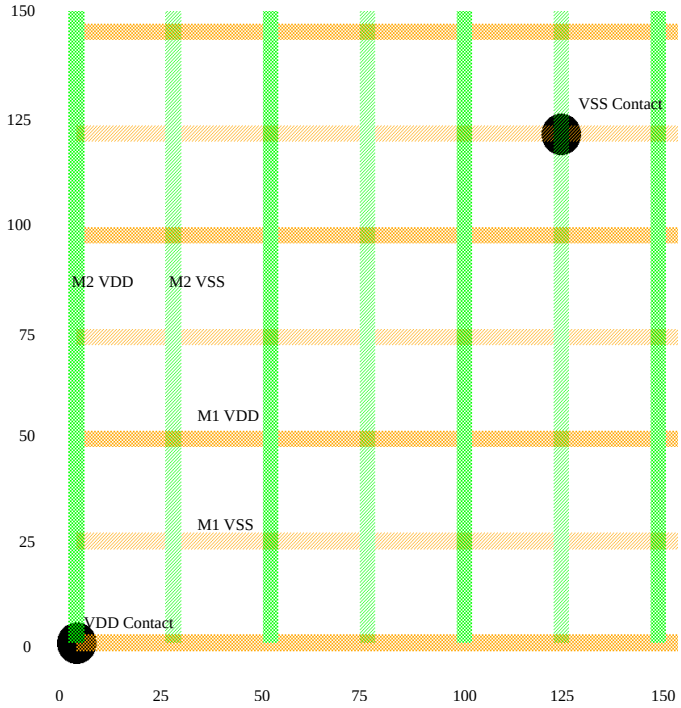

Figure 4: A small power grid to serve as an example.

VI. An Example

Perhaps the best way to make the benchmark description concrete is with a simple example. Figure 4 shows an example of a trivially small power grid composed of two levels of metal, **M1** is horizontal and **M2** is vertical. Each of the metal levels has 4 V_{DD} wires and 3 V_{SS} wires. There are two package connections (denoted by the circles), one for power and one for ground.

The resulting Spice format file for this simple design is below, formatted slightly to fit within the format of this paper.

```
rr0 n3_0_0 _X_n3_0_0 0.5
v1 _X_n3_0_0 0 1
rr2 n2_125_125 _X_n2_125_125 0.5
v3 _X_n2_125_125 0 0
* layer: M1,VDD net: 1
R4 n1_0_0 n1_50_0 1.25
R5 n1_50_0 n1_100_0 1.25
R6 n1_100_0 n1_150_0 1.25
R7 n1_0_50 n1_50_50 1.25
R8 n1_50_50 n1_100_50 1.25
R9 n1_100_50 n1_150_50 1.25
R10 n1_0_100 n1_50_100 1.25
R11 n1_50_100 n1_100_100 1.25
R12 n1_100_100 n1_150_100 1.25
R13 n1_0_150 n1_50_150 1.25
R14 n1_50_150 n1_100_150 1.25
R15 n1_100_150 n1_150_150 1.25
* vias from: 1 to 3
V16 n1_0_0 n3_0_0 0.0
V17 n1_0_50 n3_0_50 0.0
V18 n1_0_100 n3_0_100 0.0
V19 n1_0_150 n3_0_150 0.0
V20 n1_50_0 n3_50_0 0.0
V21 n1_50_50 n3_50_50 0.0
V22 n1_50_100 n3_50_100 0.0
V23 n1_50_150 n3_50_150 0.0
V24 n1_100_0 n3_100_0 0.0
V25 n1_100_50 n3_100_50 0.0
V26 n1_100_100 n3_100_100 0.0
V27 n1_100_150 n3_100_150 0.0
V28 n1_150_0 n3_150_0 0.0
V29 n1_150_50 n3_150_50 0.0
V30 n1_150_100 n3_150_100 0.0
V31 n1_150_150 n3_150_150 0.0
* layer: M2,VDD net: 3
R32 n3_0_0 n3_0_50 1.25
R33 n3_0_50 n3_0_100 1.25
R34 n3_0_100 n3_0_150 1.25
R35 n3_50_0 n3_50_50 1.25
R36 n3_50_50 n3_50_100 1.25
R37 n3_50_100 n3_50_150 1.25
R38 n3_100_0 n3_100_50 1.25
R39 n3_100_50 n3_100_100 1.25
R40 n3_100_100 n3_100_150 1.25
R41 n3_150_0 n3_150_50 1.25
R42 n3_150_50 n3_150_100 1.25
R43 n3_150_100 n3_150_150 1.25
* layer: M1,GND net: 0
R44 n0_25_25 n0_75_25 1.25
R45 n0_75_25 n0_125_25 1.25
R46 n0_25_75 n0_75_75 1.25
R47 n0_75_75 n0_125_75 1.25
R48 n0_25_125 n0_75_125 1.25
R49 n0_75_125 n0_125_125 1.25
* layer: M2,GND net: 2
R50 n2_25_25 n2_25_75 1.25
R51 n2_25_75 n2_25_125 1.25
R52 n2_75_25 n2_75_75 1.25
R53 n2_75_75 n2_75_125 1.25
R54 n2_125_25 n2_125_75 1.25
R55 n2_125_75 n2_125_125 1.25
* vias from: 0 to 2
V56 n0_25_25 n2_25_25 0.0
V57 n0_25_75 n2_25_75 0.0
V58 n0_25_125 n2_25_125 0.0
V59 n0_75_25 n2_75_25 0.0
V60 n0_75_75 n2_75_75 0.0
V61 n0_75_125 n2_75_125 0.0
V62 n0_125_25 n2_125_25 0.0
V63 n0_125_75 n2_125_75 0.0
V64 n0_125_125 n2_125_125 0.0
*
iB0_0_v n1_0_0 0  0.3125m
iB0_0_g 0 n0_25_25  0.3125m
iB0_1_v n1_0_50 0  0.3125m
```

```
iB0_1_g 0 n0_25_25  0.3125m
iB0_2_v n1_0_100 0  0.3125m
iB0_2_g 0 n0_25_75  0.3125m
iB0_3_v n1_0_150 0  0.3125m
iB0_3_g 0 n0_25_125  0.3125m
iB0_4_v n1_50_0 0  0.3125m
iB0_4_g 0 n0_25_25  0.3125m
iB0_5_v n1_100_0 0  0.3125m
iB0_5_g 0 n0_75_25  0.3125m
iB0_6_v n1_50_50 0  0.3125m
iB0_6_g 0 n0_25_25  0.3125m
iB0_7_v n1_50_100 0  0.3125m
iB0_7_g 0 n0_25_75  0.3125m
iB0_8_v n1_100_50 0  0.3125m
iB0_8_g 0 n0_75_25  0.3125m
iB0_9_v n1_100_100 0  0.3125m
iB0_9_g 0 n0_75_75  0.3125m
iB0_10_v n1_50_150 0  0.3125m
iB0_10_g 0 n0_25_125  0.3125m
iB0_11_v n1_100_150 0  0.3125m
iB0_11_g 0 n0_75_125  0.3125m
iB0_12_v n1_150_0 0  0.3125m
iB0_12_g 0 n0_125_25  0.3125m
iB0_13_v n1_150_50 0  0.3125m
iB0_13_g 0 n0_125_25  0.3125m
iB0_14_v n1_150_100 0  0.3125m
iB0_14_g 0 n0_125_75  0.3125m
iB0_15_v n1_150_150 0  0.3125m
iB0_15_g 0 n0_125_125  0.3125m
.op
.end
```

This simple example shows all the significant components of the larger examples. The collection of benchmarks will include this benchmark for reference.

VII. Summary

We presented a number of power grid benchmark problems which we hope will motivate new research in the area, and result in breakthroughs in this challenging problem. The benchmarks will be made available at the following web site:

`http://dropzone.tamu.edu/ pli/PGBench/`

The benchmark files, of course, are quite large and are therefore stored in compressed form. Any future corrections and documentation will be placed on the same web site. Comments, suggestions, and requests concerning these benchmarks are welcome, and should be sent to the author.

Acknowledgments

The author is indebted to many colleagues within IBM who have enabled, collaborated, and driven this effort. Prominent among them are Raju Balasubramnian, Erich Schanzenbach, Ivan Wemple, Joseph Kozhaya, Howard Smith and Chenggang Xu. The author also thanks Dr. Peng Li of Texas A&M University, and Mr. Zhiyu Zeng for setting up the web site for the benchmarks.

References

[1] S. N. Adya and I. L. Markov, "Consistent Placement of Macro-Blocks using Floorplanning and Standard-Cell Placement," *Proceedings of ISPD*, 2002.

[2] G. Nam, M. Yildiz, D. Z. Pan, and P. H. Madden, "ISPD Placement Contest Updates and ISPD 2007 Global Routing Contest," *Proceedings of ISPD*, 2007.

[3] R. Patel, S. Rajgopal, D. Singh, F. Baez, G. Mehta and V. Tiwari, "Reducing Power in High-Performance Microprocessors," *Proceedings of DAC* 1998

[4] J. P. Halter, and F. N. Najm, "A gate-level leakage power reduction method for ultra-low-power CMOS circuits," *Proceedings of CICC* 1997

[5] A. Dharchoudhury et. al., "Design and Analysis of Power Distribution Networks in PowerPC Microprocessors", *Proceedings of DAC* 1998

[6] E. Chiprout, "Fast Flip-Chip Power Grid Analysis Via Locality and Grid Shells", *Proceedings of ICCAD* 2004

[7] S. R. Nassif, and J. N. Kozhaya, "Fast Power Grid Simulation," *Proceedings DAC* 2000

[8] R. Panda, D. Blaauw, R. Chaudhry, V. Zolotov, B. Young and R. Ramaraju, "Model and analysis for combined package and on-chip power grid simulation," *Proceedings of ISLPED* 2000

[9] D. Kouroussis and F. N. Najm, "A static pattern-independent technique for power grid voltage integrity verification," *Proceedings of DAC* 2003

[10] H. H. Chen and D. D. Ling, "Power Supply Noise Analysis Methodology for Deep-Submicron VLSI Chip Design," *Proceedings of DAC* 1997

[11] F. N. Najm, "A survey of power estimation techniques in VLSI circuits," IEEE Trans. VLSI, Dec. 1994

[12] H. Su, Y. Liu, A. Devgan, E. Acar, S. Nassif, "Full Chip Leakage Estimation Considering Power Supply and Temperature Variations," *Proceedings of ISLPED* 2003

[13] S. R. Nassif, "The Impact of Variability on Power," *Proceedings ISLPED* 2004

[14] L.W. Nagel, *SPICE2: A Computer Program to Simulate Semiconductor Circuits*. PhD thesis, University of California, Berkeley, 1975.