Analysis and Optimization of Power Grids *

Sachin S. Sapatnekar
ECE Dept, Univ. of Minnesota
200 Union St. SE
Minneapolis, MN 55455
sachin@ece.umn.edu

Haihua Su
IBM Corp.
11501 Burnet Rd.
Austin, TX 78758
haihua@us.ibm.com

Abstract

This paper presents an overview of techniques used for the analysis and optimization of power grids. After an initial motivation and an examination of technology trends that predict the growing importance of the problem of building reliable supply networks, several techniques for power grid analysis are described. Next, optimization methods for building better grids, using techniques such as topology selection, wire width optimization and the insertion of decoupling capacitors, are surveyed. Finally, a discussion on issues related to resource contention between power grid lines and signal lines, and close correlation between power grid and gate delays, are presented, motivating the need for the codesign of power, clock and signal lines.

1 Introduction

The ability of integrated circuits to implement correct logic at desired speeds is contingent on the design of reliable supply grids. The dependency on supply levels stems from the fact that a digital circuit is merely a practical implementation of an abstraction that models two binary logic levels with two distinct voltages, typically the supply voltage, V_{dd}, and the ground potential of zero, and the gates in a circuit are essentially sets of switches that steer either the V_{dd} or ground signal to the output, depending on the logic level. The ability to supply reliable V_{dd} and ground levels not only dictates the correctness of the logic value at the gate outputs, but also determines the speed of signal transitions.

Imperfections in the supply grid arise due to the non-ideal nature of the conductors that carry the supply and ground signals from the supply pins to the gates. These on-chip problems manifest themselves in several ways:

IR drops are caused by the flow of currents through wires that are imperfect conductors with nonzero resistances. From one technology generation to the next, wire resistances of proportionately scaled wires have increased significantly, causing IR drops to become major factors

*This research was supported in part by the SRC under contract 99-TJ-714 and by the NSF under award CCR-0205227.

in the performance of supply networks. A typical voltage contour for a power grid is illustrated in Figure 1.

L dI/dt effects result from the increasing inductive behavior seen in global lines, particularly in the uppermost levels of metal, in newer process generations. These inductance effects, which include both self and mutual inductances in the supply lines and in the package, could result in unwanted ringing and noise spikes.

Electromigration is an aging effect that appears on supply net wires that often carry currents in the same direction over many years, resulting in the migration of metal over an extended period of time. It has been empirically observed that the mean time to failure for a wire is dependent on the current density through its cross section and this can be controlled by limiting this current density.

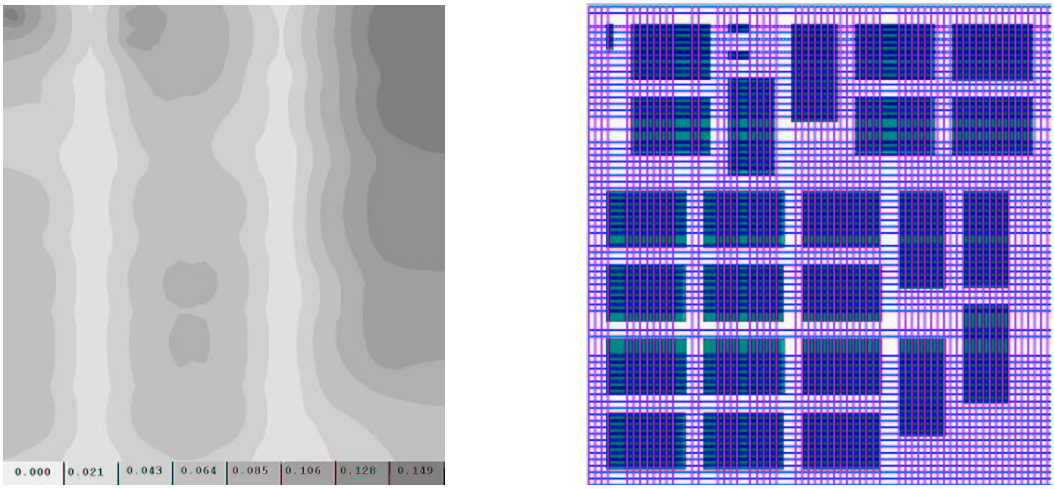

Figure 1: A typical IR drop contour on a ground network for benchmark ac3 and the corresponding grid on layers M3 and M4.

The design of the package also plays an important role in the integrity of the supply grid [1], but in this paper we restrict our attention to on-chip problems. However, some packaging issues must be incorporated into on-chip analysis by considering factors such as the spatial distribution of the pins and the parasitics. For example, the supply grid pins may either be distributed on the periphery of the chip, or may use C4 connections distributed over the entire surface of the chip. In the former case, the voltage contours typically show a degradation as the signal travels the periphery to the interior of the chip, while the latter allows for a more even distribution, as in the figure at left in Figure 1.

Advanced technologies have seen increasing hindrances to providing correct and noise-free supply and ground signals, and therefore, the correct design of supply grids is an important problem for

deep submicron and nanometer technologies. Two trends can be deduced from projections in the International Technology Roadmap for Semiconductors (ITRS) [2]: firstly, the supply voltage levels go down as technology progresses, and secondly, the average current entering the chip increases significantly over time.

Both of these trends have serious implications. The reduced levels for the "perfect" supply signal that is supplied to the chip imply reduced noise margins, necessitating tighter controls on the voltage levels as supply signals are distributed over the chip. On the other hand, the increased currents that flow through the chip, together with the increasing speeds at which gates switch (as reflected by the predicted clock rates) imply potentially large IR and L dI/dt effects, in addition to increased susceptibility to electromigration. This implies that careful design of the supply grid will be more important than ever in the future.

Feature size (nm)	250	180	130	100	65	45
Nominal $V_{dd}(V)$	2.1	1.8	1.2	1.0	0.7	0.6
ΔDelay(0.1V)	6.9%	18.2%	24.6%	11.2%	63.3%	79.8%
ΔDelay(10%)	19.3%	32.0%	28.1%	11.2%	29.5%	29.0%

Table 1: Effect of a variation in the supply voltage on the delay of an inverter.

Apart from its influence on circuit correctness, the fidelity of the power grid also affects the predictability of circuit performance. The value of the supply voltage levels strongly impacts the delays of gates, and variations in the supply level can cause unwanted variations in the delays associated with combinational logic. Table 1 shows the effects of voltage level variations in the supply and ground lines for an inverter driving a fanout of four in various technologies, based on SPICE simulations, using model files for existing technologies and for the Berkeley Predictive Technology Model [3]. Two scenarios are illustrated: one in which the supply and ground grids are each permitted to permit errors of up to 0.1V for each technology, and another in which the errors are limited to 10% of V_{dd}. In the former case, it is unsurprising that the effect becomes more visible with reduced voltage levels; for example, at the 45nm node, the uncertainty in the supply voltage becomes a very significant fraction of the V_{dd} level. However, if the voltage uncertainty is maintained to be within 10% of the supply voltage level, the uncertainty is more controlled.

2 Analysis of supply networks

2.1 Formulation of equations

Current-day supply networks may contain tens or hundreds of millions of nodes, and analyzing them accurately is therefore acutely computational as it requires the solution of systems of linear equations in about as many variables. The formulation of these equations first requires the elements of the supply grid to be modeled using basic electrical elements. The wires in the V_{dd} grid can be modeled as a linear RLC network that is connected through pads to multiple V_{dd} sources, and to switching elements in the circuit that draw current through the power grid. These switching

elements are most often modeled either as current sources or as RC switch elements [4, 5]. In either case, the simulation of the network requires the solution of the following system of differential equations, which may be formulated using modified nodal analysis (MNA) [6]:

$$\mathbf{G} \cdot \mathbf{x}(t) + \mathbf{C} \cdot \mathbf{x}'(t) = \mathbf{b}(t), \tag{1}$$

where $\mathbf{G}$ is a conductance matrix, $\mathbf{C}$ is the admittance matrix resulting from capacitive and inductive elements, $\mathbf{x}$(t) is the time-varying vector of voltages at the nodes, and currents through inductors and voltage sources, and $\mathbf{b}$(t) is the vector of independent time-varying current sources. This differential system is efficiently solved in the time domain by reducing it to a linear algebraic system

$$(\mathbf{G} + \mathbf{C}/h) \cdot \mathbf{x}(t) = \mathbf{b}(t) + \mathbf{C}/h \cdot \mathbf{x}(t-h), \tag{2}$$

using the Backward Euler (BE) technique with a small fixed time step, h. The coefficient matrix can be constructed to be symmetric and positive definite when the power grid is represented by RC elements and current sources; if inductances are used, a K matrix [7] formulation can be used to maintain this property. Since the coefficient remains constant under a fixed time step, transient analysis entails only one Cholesky factorization followed by forward/backward substitution at each time point.

2.2 Efficient solution using hierarchical methods

Even with the use of the most efficient methods that exploit sparsity and positive definiteness, the solution of a system of a hundred million equations is highly computation- and memory-intensive. To alleviate this problem, it is imperative that hierarchy be used in the analysis of these large systems. Several techniques have been suggested in this connection. We will describe two approaches from [8] and [9], respectively, in the description that follows. In addition to these approaches, other methods [10, 11] using model order reduction techniques have been proposed for the hierarchical analysis of supply nets.

2.2.1 Exploiting existing hierarchy

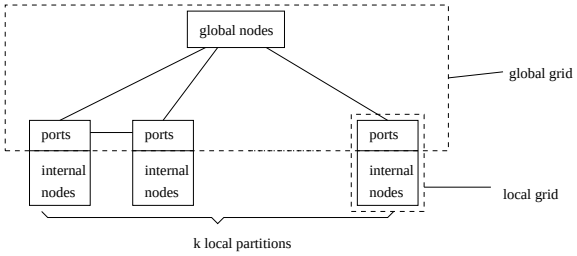

Figure 2: Hierarchical power network analysis

The macromodel-based approach in [8] is illustrated in Figure 2. The philosophy of this approach lies in the fact that the power grid is inherently hierarchical since it is created as a part of a hierarchical design process, where individual blocks with locally constructed power are first designed individually and then assembled at the chip level.

Based on this structure, the power grid can be divided into k local partitions, corresponding to blocks, and a global partition that connects the power grids within these blocks. A node in a local partition having links only to other nodes in the same partition is called an *internal node*, a node in the global partition is called a *global node*, and a node in a local partition that is connected to some node outside the local partition (i.e., in the global partition or in another local partition) is called a *port*. The *global grid* is then defined to include the set of nodes that lie in the global partition and the port nodes, while the grid in a local partition constitutes a *local grid*. The runtime improvements of this approach arise from the fact that existing design hierarchies often permit the power grid to be divided into partitions with a small number of port nodes, relative to the number of internal nodes. The technique consists of the following steps:

Step 1: Macromodeling Typically, the number of port nodes is much smaller than the total number of internal nodes, and therefore each of the k local grids may be modeled as a multi-port linear element represented by a macromodel of the type $\mathbf{I} = A \cdot \mathbf{V} + \mathbf{S}$, where $\mathbf{I}$ and $\mathbf{V}$ are the vectors of port voltages, and A and $\mathbf{S}$ are, respectively, a constant matrix and a constant vector.

Step 2: Solution of the global grid Once the macromodels for all the local grids are generated, the entire network is abstracted simply as the global grid, with the macromodel elements connected to it at the port nodes. The equations of each local grid may be stamped into nodal equations of the global grid, so that the system to be solved at the global level is considerably smaller than the original system as all internal nodes have been removed from consideration. However, the A matrices for each partition are typically dense, which could severely limit the gains of the method. Hence it is important to sparsify these matrices while preserving the positive definiteness of the system. The approach in [8] presents a sparsification scheme that formulates the problem as an integer knapsack problem and bounds the error obtained from sparsification.

Step 3: Solution of the local grids For each partition, $\mathbf{I}$ is obtained from the above solution and from using the port voltages. Equation (2) is solved for each partition using $\mathbf{I}$ on the right hand side, to obtain voltages at the internal nodes of partitions.

Experimental results in [8] show that this technique is capable of simulating networks with up to 63 million nodes, providing savings in the run time as compared to the non-hierarchical approach. Even more importantly, the non-hierarchical is limited by memory requirements and cannot simulate networks of 20 million nodes or more. In contrast, the peak memory utilization of this approach is much smaller, providing it a strong advantage and making the algorithm much more scalable in practice.

2.2.2 "Introducing" hierarchy into the solution

The approach in [9] implicitly introduces hierarchy into a supply grid by solving it in two steps: first, by creating a coarsened form of the network with a reduced number of nodes, which can be solved efficiently, and then by propagating the result of this solution to the full network. This technique is inspired by the multigrid approach for solving partial differential equations, where the coarse solution controls the low frequency error components in the solution, and the finer solution corresponds to a relaxation step that reduces the high frequency error. The technique consists of four steps:

Grid reduction, in which the large power grid is coarsened by selecting a subset of nodes that are to be maintained, while the other nodes are removed. The number of variables is therefore significantly reduced from from n to m.

Interpolation, in which an $n \times m$ interpolation operator matrix P is defined to map the original grid to the coarsened grid. This interpolation operator relates the voltages on the removed nodes to those on the coarsened grid, thereby allowing the solution of the coarsened grid to accurately reflect that of the original grid.

Solution of the coarsened grid, in which a solution is found for the voltages in the coarsened grid by solving the above linear equations.

Mapping the solution from the coarsened grid to the original grid, by applying the interpolation operator concludes the process.

The proposed multigrid-like method is a fusion of the algebraic multigrid (AMG) method that forms a reduction matrix based on directly on the A matrix, and the standard multigrid (SMG) method that forms the reduction matrix based on geometric information of the power grid. While both AMG and SMG are iterative methods starting from the initial solution predicted from the coarsened grid, the multigrid-like method in [9] requires no relaxation.

Experimental results in [9] on two real ASIC designs with about half a million nodes each show improvements of about 16–20× over an analysis of the flat network for DC analysis, and even better for transient analysis, where the speedups are of the order of 600×.

3 Optimizing supply networks

An analysis of the dependence of the voltage drop using ITRS parameter projections is presented in [12]. It is shown that the projected trends can be overcome using one of the following techniques:

- Adding *decoupling capacitances* to the layout, which can be done at an area cost.
- Increasing the effective conductance from the blocks that draw current from the supply grid to the decoupling capacitance by placing the capacitance closer to the load and by increasing the conductivity of the power grid by using *wider wires* or *denser topologies*.

- Increasing the grid conductance, again by using wider wires or denser topologies, as above.

We consider several techniques that may be used to reduce noise on the power grid: specifically, those of selecting the power grid topology, of widening wires, and of inserting decaps.

For purposes of optimization, it is important to develop metrics that consider not only the peak power grid noise, but also the duration of the noise pulse. A useful metric to estimate power-grid-induced noise at a node is the integral of the voltage droop below a user specified noise ceiling [13]:

$$z_j(p) = \int_0^T \max\{NM_H - v_j(t,p), 0\}dt = \int_{t_s}^{t_e} \{NM_H - v_j(t,p)\}dt, \tag{3}$$

where p represents the tunable circuit parameters, which may include the values of the decoupling capacitors or wire widths. The voltage droop integral beyond the expressed by Equation (3) represents the shaded area in Figure 3. We define the measure of goodness for the whole circuit as the sum of the individual node metrics:

$$Z = \sum_{j=1}^{K} z_j(p), \tag{4}$$

where K is the number of nodes. This metric penalizes more harshly transients that exceed the imposed noise ceiling by a large amount for a long time, and has empirically been seen to be more effective in practice than one that penalizes merely the maximum noise violation. Intuitively, this can be explained by the fact the metric incorporates, in a sense, both the voltage and time axes together, as well as spatial considerations through the summation over all nodes in the circuit.

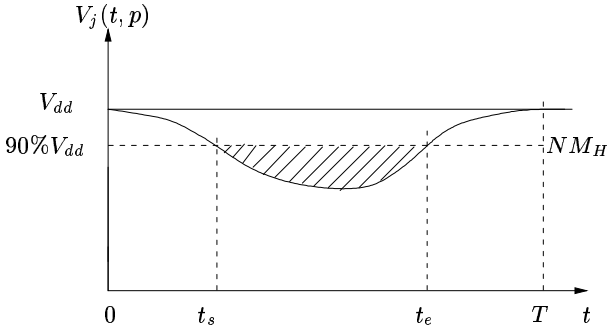

Figure 3: Illustration of the voltage droop at a given node in the V_{dd} power grid. The area of the shaded region corresponds to the integral z at that node.

3.1 Topology selection

The earliest techniques for supply grid optimization presented results on tree structures. While tree structured grids are easy to analyze, they have been found to be incapable of providing reliable

signal delivery. Subsequent work has focused on mesh optimization which however, requires longer analysis times. The work in [10, 14] presented a method that is intermediate to the two extremes, and uses a mixture of tree and mesh structure in attempting to find the best of both worlds. An overlying coarse mesh structure of a user-specified topology provides global distribution of the P/G signals across the chip. From various nodes of this mesh, tree structures of user-specified topologies originate and distribute the supply voltage to the utilization points, each of which is modeled as an equivalent RC branch. Other types of hierarchical power grids may also be constructed with, for example, meshes driving meshes. As long as the number of connections between meshes is relatively sparse, fast analysis in both the time domain (see Section 2.2.1) and frequency domain [10] is possible.

3.2 Wire widening

Various techniques for wire widening have been proposed in the literature. These include techniques for flat netlists [15, 16] as well as those for hierarchical grids [10]. As an example that illustrates an approach to wire widening, we describe here the method used in [15], where the power grid sizing problem is directly formulated as a nonlinearly constrained nonlinear programming problem as follows:

$$\begin{aligned} \text{minimize} \quad & Area(w_j) = \sum_{j=1}^{N_{wire}} l_j \times w_j \\ \text{subject to} \quad & w_{min} \leq w_j \leq w_{max},\ j = 1 \cdots N_{wire} \\ \text{and} \quad & Z(w_j) < \epsilon \end{aligned} \tag{5}$$

where w_j and l_j are, respectively, the width and length of the jth wire, N_{wire} is the total number of wires in the grid, w_{min} and w_{max} are the lower and upper bound on the allowable wire width, Z is the noise metric described in Equation (4), and ϵ is a very small number.

The objective function is to minimize the total area that the power wire occupies. The first constraint restricts every power wire width to lie within a realistic range that is technology dependent. The second constraint requires the definition of the parameter Z, defined earlier. The evaluation of the nonlinear constraint function Z can be performed by transient analysis of the power grid circuit, and its sensitivity with respect to all the variables w_j can be calculated using the adjoint method [6, 17]. A standard Sequential Quadratic Programming (SQP) solver is used to solve the optimization problem. This solver requires users to provide subroutines to evaluate the objective and constraint functions and their derivatives with respect to each decision variable.

It is possible to solve a modified version of the optimization problem using linear programming methods, as is done in [16], where a nonlinear program formulation adapted from [18] is cleverly solved using an iterative sequence of linear programs. The objective function there, however, is the worst-case voltage drop and extending the solution to handle the integral of the voltage violation is nontrivial.

3.3 Decoupling capacitance insertion

The placement of decaps can have a large influence on the integrity of the supply signals. In this section, we will describe a decap optimization and placement algorithm for row-based standard-cell design typical of Application Specific Integrated Circuits (ASIC) where each row has a fixed height. An approach for decap insertion during floorplanning for custom layouts is described in [19].

The ASIC layout here has N rows, with the ith row having M_i cells (blocks). A ratio $r_i (\leq 100\%)$ of the ith row is filled by cells, and the remainder of the row, corresponding to a fraction of $(1 - r_i)$, may be used for placing decaps. This corresponds to a post-placement optimization, after cells have been assigned to rows, but the rows may be perturbed slightly in order to insert decaps, with the objective of minimizing the power grid noise metric Z described in Equation (4).

The problem of decoupling capacitor optimization is therefore formulated as:

$$\begin{array}{rcl} \text{Minimize} & Z(w_j) & j = 1 \cdots N_{decap} \\ \text{Subject to} & \sum_{k \in row_i} w_k \leq (1 - r_i) W_{chip} & i = 1 \cdots N_{row} \\ \text{and} & 0 \leq w_j \leq w_{max} & j = 1 \cdots N_{decap} \end{array} \tag{6}$$

The scalar objective Z (defined by Eqnation 3) is a function of all the decap widths and N_{decap} is the total number of decaps in the chip. The first constraint states that the total decap width in a row cannot exceed the total amount of empty space in that row, and W_{chip} and N_{row} denote, respectively, the width of the chip and the number of rows in the chip. The second constraint restricts the decap widths within a realistic range. An upper bound w_{max} for a cell in row i is easily seen to be $(1 - r_i)W_{chip}$, which is the largest empty space in row i; while the lower bound of each decap width is zero.

Equation (7) represents a linearly constrained nonlinear optimization problem. The objective function Z can be obtained after the transient analysis of the power grid circuit, and its sensitivity with respect to all the variables w_j can be calculated using the adjoint method [6, 17]. A standard quadratic programming (QP) solver [20] for solving large nonlinear optimization problems is employed, starting the optimization with an initial guess that uniformly distributes the vacant space in each row to each decoupling capacitor in each row. At the end of the procedure, decaps that become smaller than a given threshold may be removed.

4 Codesign of supply, signal and clock networks

4.1 Codesign of supply and signal networks

The need for dense power grids with widened wires implies that a major consumer of these resources is the supply network. With the role of interconnect becoming increasingly critical in nanometer designs, the need to meet stringent performance constraints has resulted in strong contention for scant routing resources. Global signal wires also compete for the same routing resources, as they often require shortest-path routes to meet their own performance requirements. Traditionally, supply and signal nets have been designed independently, with the routing needs for a regular

power grid being determined first, after which the remaining resources are calculated to provide routing resource budgets for the signal nets.

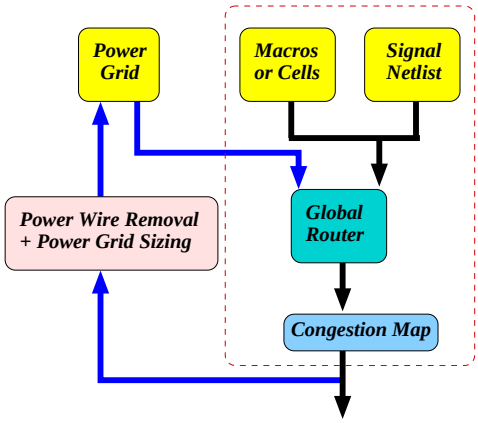

Figure 4: Congestion-driven power grid design and global routing.

As the number and criticality of global signal wires becomes more dominant, such a methodology becomes unsustainable as the initial budgets may often be entirely unreasonable. Therefore, in nanometer designs, there is a strong need for a unified approach to the design of signal wires and power grids, with an integrated approach to routing resource management.

While it is convenient to build a regular power grid with a constant pitch (defined as the distance between adjacent wires in the grid), some degrees of freedom exist and it is desirable that they be exploited. For instance, in regions where the demand for routing resources from signal nets is high, a sparser power grid may be used as long as the performance constraints on the supply and ground lines can be met; likewise, signal nets are well advised to avoid the hot spots of the chip if possible, since these may need a locally dense power grid.

The procedure in [15] presents a method for the codesign of supply and signal networks. The overall flow is illustrated in Figure. 4. As in global routing, the entire chip is tessellated into an array of grid cells. The width of a boundary between two neighboring grid cells represents the limited resources that must be shared on each layer by the supply lines and the signal lines that traverse the boundary. The procedure consists of the following steps:

Initialization An initial uniform width power grid is built by choosing sufficiently wide wires, and this corresponds to a feasible solution that satisfies the supply network constraints. Although this solution can be found quickly, it is likely to utilize more resources than a more carefully sized solution. Using the remaining routing resources, the global signal nets are then routed while attempting to satisfy all boundary capacities. In general, it is likely that this will result in overflows at several boundaries, which leads to the next iterative steps.

Supply net readjustment and sizing The adjustments made to the power grid consist of wire removal from the grid in congested regions, and sizing of the power grid to compensate for this removal. The wire removal heuristic identifies a set of critical wires that traverse "hot spots" of the chip with high voltage drops, and considers only noncritical wires as candidates for removal. The removal of each such wire presents new opportunities for rerouting the other wires, including those that cross congested boundaries that the removed wire may not traverse, and hence the procedure removes a limited number of power wires in each iteration.

The removal of any power wires is likely to degrade the quality of the power grid, and therefore, the removal step is followed by a power grid resizing step in which the remaining wires in the grid are individually sized using an SQP-based formulation to compensate for the loss of the removed wires.

Signal net rerouting The signal nets that cross the overflowing boundaries are then ripped up and rerouted to utilize the newly available capacity. If this satisfies all constraints, the procedure ends; otherwise, it iterates through another step of adjusting and resizing the supply net, followed by rerouting the signal nets.

The structure at the right in Figure 1 shows an example power grid topology that is created by this method for the benchmark circuit ac3. It can be seen that the power grid is sparser in some places than in others, and these correspond to regions where the demand from signal wires is large while the requirements on the power grid are less stringent.

4.2 Codesign of supply and clock nets

In newer technologies, the ever higher clock frequencies cause the clock network to draw larger currents from the power grid, and therefore consume more power. The close-coupling between the clock and power distribution networks demand codesign of both networks. While the cosynthesis of supply and clock network remains an open topic, recent publications have shown research focusing on the analysis of power grid noise effects in static timing [21, 22].

A practical methodology that is closely related to clock network design, for clock skew verification in the presence of power grid noise is presented in [23]. The power grid, a large linear portion of the circuit, and the clock network, a relatively small nonlinear portion of the circuit, are solved separately. Given this partitioning, assuming ideal power supply (zero IR-drop) in the initial clock circuit, the iterative analysis process alternately analyzes the clock network and the power grid until convergence. It is empirically shown that a weighted average of the first two iterations produces a relatively accurate result and no further iteration is necessary in practice.

5 Conclusion

The problem of analyzing and optimizing power grids will become extremely important for future designs. Techniques that use hierarchy are likely to be the most successful in finding good

solutions to the analysis and optimization problems. In future, supply grid analysis and optimization will also have to consider effects due to packaging. Substrate effects have long been known to be important [24], and they will play a more major role as noise issues become more critical.

A trend that will be increasingly visible in the future will be the need to codesign signal, clock and supply lines. In the future, more lines will require shielding from capacitive and inductive crosstalk, and this shielding is essentially carried out by running supply lines next to the line to be protected. With this trend, the design of supply and signal/clock lines will become even more tightly coupled in the future. Although there is little work in this area that specifically addresses supply network issues at this time, initial forays in the direction of shield design have been proposed, for example, in [25, 26].

6 Acknowledgements

The authors would like to acknowledge the help of prior collaborations in this area with David Blaauw, Kaushik Gala, Jiang Hu, Sani Nassif, Rajendran Panda and Min Zhao.

References

[1] L. Polka *et al.*, "Package-level interconnect design for optimum electrical performance," *Intel Technology Journal*, Q3, 2000.

[2] Semiconductor Industry Association, *The International Technology Roadmap for Semiconductors*, 2001. `http://public.itrs.net/Files/2001ITRS/Home.html`.

[3] Device Group, University of California, Berkeley, *Berkeley Predictive Technology Model*, 2002. `http://www-device.eecs.berkeley.edu/~ptm`.

[4] H. H. Chen and D. D. Ling, "Power supply noise analysis methodology for deep-submicron VLSI chip design," in *Proceedings of the ACM/IEEE Design Automation Conference*, (Anaheim, CA), pp. 638–643, June 1997.

[5] A. Dharchowdhury, R. Panda, D. Blaauw, R. Vaidyanathan, B. Tutuianu, and D. Bearden, "Design and analysis of power distribution networks in PowerPC^{TM} microprocessors," in *Proceedings of the ACM/IEEE Design Automation Conference*, pp. 738–743, June 1998.

[6] L. T. Pillage, R. A. Rohrer, and C. Visweswariah, *Electronic and System Simulation Methods*. New York, NY: McGraw-Hill, 1995.

[7] A. Devgan, H. Ji, and W. Dai, "How to efficiently capture on-chip inductance effects: Introducing a new circuit element K," in *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design*, (San Jose, CA), pp. 150–155, Nov. 2000.

[8] M. Zhao, R. V. Panda, S. S. Sapatnekar, and D. Blaauw, "Hierarchical analysis of power distribution networks," *IEEE Transactions on Computer-Aided Design*, vol. 21, pp. 159–168, Feb. 2002.

[9] J. Kozhaya, S. R. Nassif, and F. N. Najm, "A multigrid-like technique for power grid analysis," *IEEE Transactions on Computer-Aided Design*, vol. 21, pp. 1148–1160, Oct. 2002.

[10] H. Su, K. H. Gala, and S. S. Sapatnekar, "Fast analysis and optimization of power/ground networks," in *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design*, (San Jose, CA), pp. 477–480, Nov. 2000.

[11] Y. Cao, Y.-M. Lee, T.-H. Chen, and C. C.-P. Chen, "HiPRIME: Hierarchical and passivity reserved interconnect macromodeling engine for RLKC power delivery," in *Proceedings of the ACM/IEEE Design Automation Conference*, (New Orleans, LA), pp. 379–384, June 2002.

[12] H. Su, S. S. Sapatnekar, and S. R. Nassif, "An algorithm for optimal decoupling capacitor sizing and placement for standard cell layouts," in *Proceedings of the ACM International Symposium on Physical Design*, (San Diego, CA), pp. 477–480, Apr. 2002.

[13] A. R. Conn, R. A. Haring, and C. Visweswariah, "Noise considerations in circuit optimization," in *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design*, (San Jose, CA), pp. 220–227, Nov. 1998.

[14] J. C. Shah, A. A. Younis, S. S. Sapatnekar, and M. M. Hassoun, "An algorithm for simulating power/ground networks using Padé approximants and its symbolic implementation," *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, vol. 45, pp. 1372–1382, Oct. 1998.

[15] H. Su, J. Hu, S. R. Nassif, and S. S. Sapatnekar, "Congestion-driven codesign of power and signal networks," in *Proceedings of the ACM/IEEE Design Automation Conference*, (New Orleans, LA), pp. 477–480, June 2002.

[16] X.-D. Tan, C.-J. R. Shi, D. Lungeanu, J.-C. Lee, and L.-P. Yuan, "Reliability-constrained area optimization of vlsi power/ground networks via sequence of linear programmings," in *Proceedings of the ACM/IEEE Design Automation Conference*, (New Orleans, LA), pp. 78–83, June 1999.

[17] S. W. Director and R. A. Rohrer, "The generalized adjoint network and network sensitivities," *IEEE Transactions on Circuit Theory*, vol. 16, pp. 318–323, Aug. 1969.

[18] S. Chowdhury, "Optimum design of reliable IC power networks having general graph topologies," in *Proceedings of the ACM/IEEE Design Automation Conference*, (Las Vegas, NV), pp. 787–790, June 1989.

[19] S. Zhao, K. Roy, and C.-K. Koh, "Decoupling Capacitance Allocation for Power Supply Noise Suppression," in *Proceedings of the ACM International Symposium on Physical Design*, (Napa, CA), pp. 66–71, Apr. 2001.

[20] C. Zhu, R. H. Byrd, and J. Nocedal, *LBFGS-B: Fortran subroutines for large-scale bound constrained optimization.* EECS Department, Northwestern University, Evanston, IL, 1994.

[21] G. Bai, S. Bobba, and I. N. Hajj, "Static timing analysis including power supply noise effect on propagation delay in VLSI circuits," in *Proceedings of the ACM/IEEE Design Automation Conference*, (Las Vegas, NV), pp. 295–300, June 2001.

[22] L. H. Chen, M. Marek-Sadowska, and F. Brewer, "Coping with buffer delay change due to power and ground noise," in *Proceedings of the ACM/IEEE Design Automation Conference*, (New Orleans, LA), pp. 860–865, June 2002.

[23] R. Saleh, S. Z. Hussain, S. Rochel, and D. Overhauser, "Clock skew verification in the presence of IR-drop in the power distribution network," *IEEE Transactions on Computer-Aided Design*, vol. 19, pp. 635–644, June 2000.

[24] B. Stanisic, R. A. Rutenbar, and L. R. Carley, *Synthesis of Power Distribution to Manage Signal Integrity in Mixed-Signal ICs.* Boston, MA: Kluwer Academic Publishers, 1996.

[25] P. Saxena and S. Gupta, "Shield count minimization in congested regions," in *Proceedings of the ACM International Symposium on Physical Design*, pp. 78–84, Apr. 2002.

[26] J. D. Ma and L. He, "Toward global routing with RLC crosstalk constraints," in *Proceedings of the ACM/IEEE Design Automation Conference*, pp. 669–672, June 2002.