

Hierarchical Analysis of Power Distribution Networks

Min Zhao, Rajendran V. Panda, Sachin S. Sapatnekar, David Blaauw

Abstract—Careful design and verification of the power distribution network of a chip are of critical importance to ensure its reliable performance. With the increasing number of transistors on a chip, the size of the power network has grown so large as to make the verification task very challenging. The available computational power and memory resources impose limitations on the size of networks that can be analyzed using currently known techniques. Many of today's designs have power networks that are too large to be analyzed in the traditional way as flat networks. In this paper, we propose a hierarchical analysis technique to overcome the aforesaid capacity limitation. We present a new technique for analyzing a power grid using macromodels that are created for a set of partitions of the grid. Efficient numerical techniques for the computation and sparsification of the port admittance matrices of the macromodels are presented. A novel sparsification technique using a 0-1 integer linear programming formulation is proposed to achieve superior sparsification for a specified error. The run-time and memory efficiency of the proposed method are illustrated on industrial designs. It is shown that even for a 60 million node power grid, our approach allows for an efficient analysis, whereas previous approaches have been unable to handle power grids of such size.

I. INTRODUCTION

With the increase in the complexity of VLSI chips, designing and analyzing a power distribution network has become a challenging task. A robust power network design is essential to ensure that the circuits on a chip operate reliably at the guaranteed level of performance. A poorly designed power network can become the cause for a variety of problems such as loss of circuit performance, noise generation, and electromigration failures. With the increased power level and device densities of microprocessors in sub-micron technologies, these problems are more likely unless serious attention is given to power network design. Critical to obtaining a robust design is the ability to analyze the network efficiently several times in the design cycle. Several previously published research works [1], [2], [3], [4], [5], [6], [7], [8] discuss methodologies and techniques to accomplish this task efficiently.

The difficulty in power network analysis stems mainly from three sources: (i) the network is very large, typically 1 million to 100 million nodes, (ii) the network is nonlinear as it contains switching devices, and (iii) the voltage and current distribution in the network is dependent on the instruction executed on the processor. Our work, presented in this paper, addresses the first problem. The second

problem is circumvented traditionally [1], [9] by performing nonlinear simulation of individual circuit blocks without including the parasitics in the power interconnects, and then simulating the power interconnect as a whole using the time-variant current profiles, obtained in the nonlinear simulation as the excitation sources. The third problem is one of obtaining a good coverage of all possible worst-case power demand situations. Manually generated “hot loops,” an extensive set of input vectors, and statically generated worst-case current profiles [10], [11], [12], [13], [6] are some of the alternatives that address the worst-case coverage issue. Recently, frequency domain based construction of worst-case current profile including inductive effects has been proposed [14].

The first problem becomes a critical issue due to the rapid increase of the number of transistors on a chip. At the current technological level, it is seen that the available computing resources are insufficient to simulate very large power grids of today's microprocessors using a flat model. The size of the power grid of a typical high performance microprocessor in 0.18 micron design, and using 6-9 layers of metal, is in the range of 30 million to 120 million nodes. Moreover, the power grid simulation would require solving a linear system of similar size at multiple time points. Clearly, the speed and memory capacity of a typical computing environment is insufficient to solve such a large system even with the most efficient linear system solution techniques.

The size based complexity of the problem has been addressed in several works [1], [2], [8]. The approaches in [1], [2] proposed the usage of very efficient sparse linear system solution techniques. Cholesky factorization (direct method) [15] and conjugate gradient techniques with preconditioners (iterative method) [15] have been used to solve the linear system associated with the power grid. These specialized techniques operate very efficiently by exploiting the special structure and properties of the underlying linear system. However, the solutions proposed earlier have applied these techniques to a flat (nonhierarchical) model of the power network. As a result, there is a serious limitation on the size of the problem they can solve, the limitation being imposed by the amount of memory available for computation. The work in [8] proposed a PDE-like multi-grid method for the simulation of large power grids. This method, which is particularly attractive for mesh grids, reduces the size complexity by solving several coarser meshes and then extrapolating the results to the original fine mesh. Therefore, this method solves the network approximately and can result in an unpredictable error, especially for non-uniform (non-mesh) grid structures. Finally, in [5] a frequency based analysis formulation is proposed. While this approach can be efficient for repetitive signals, it will

This work was supported in part by the Semiconductor Research Corporation under contract number 99-TJ-714 and by the National Science Foundation under contract number CCR-9800992.

M. Zhao and R. Panda are with with Motorola SPS, Austin, TX 78729, USA

S. S. Sapatnekar is with the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55455, USA.

D. Blaauw is with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109, USA.

not perform well for irregular simulation vectors. Also, it suffers from the same size limitation as the time domain method described in [1], [9].

In this work, we propose a hierarchical analysis technique to overcome the limitations of the earlier approaches. Our approach comprises of the following steps: (1) Partitioning of the power grid into local and global grids using the hierarchical structure in the design or automatic partitioning techniques, (2) Generating macromodels for the local grids using efficient numerical methods, (3) Sparsifying the port admittance matrices of the macromodels, while maintaining the accuracy of the solution, (4) Simulating the global grid after augmenting it with the macromodels of the local grids, and finally, (5) Simulating the local grids where desired. Of these, Steps (2),(3) and (5) are parallelizable.

The basic strength of the proposed approach is derived from the well-known strategy of “divide and conquer,” which is realized through partitioning. The approach also enables the parallel computation of the task. However, the efficiency and usefulness of the hierarchical approach is sensitive to several factors, such as the partitioning technique, the memory and runtime costs involved in generating the macromodels, and the size and density of the macromodels. Our work in this research addresses these problems in order to realize a practical and efficient implementation of the hierarchical analysis strategy. We propose a partitioning technique that can handle extremely large power grids and effectively reduce the memory and computation time required. Moreover, a novel matrix sparsification technique based on 0-1 integer linear programming is proposed to further reduce the memory requirements. Additionally, an efficient numerical procedure for calculating the macromodels is given. The computation takes advantage of the fact that the underlying linear system is symmetric and positive definite. The proposed approach has been applied to the analysis of the power grid of a number of high performance microprocessors and DSP chips, obtaining significant memory and runtime advantages over the flat model analysis approach.

The approach can be applied to power grid models that include package inductance without modification using the method described in [7]. In addition, the approach can be applied to RLC model as long as the inductance interacts between the partitions are not considered. In RLC model of the flat method, the modified nodal formulation will no longer be positive definite, and can not longer be solved using efficient Cholesky factorization, hence leading to increased run time. The hierarchical analysis method proposed here suffer from the same inefficiency degradation as the flat method. However, in practice, inductance effects usually are important for the upper layer of metal, which can be easily put into the global partition, and the inductance of lower layer of metal can be ignored. In this case, the similar method as [7] can be applied to solve the RLC model efficiently. In this paper, we will restrict our discussion to RC power grid models for the purpose of explanation.

The remainder of the paper is organized as follows. In

Section 2, we present the concept of macromodeling, the partitioning strategy, and the computational techniques for generating the macromodels. In Section 3, the matrix sparsification technique is explained. Section 4 reports the performance results of the proposed approach for a set of benchmark designs, followed by conclusions in Section 5.

II. MACROMODELING APPROACH

A. Overview of Power Grid Simulation

Before presenting the macromodeling approach, we present an overview of power grid simulation in general. A chip’s power distribution system is modeled as a linear RLC network with independent time-varying current sources modeling the switching currents of the transistors. Simulating the network requires solving the following system of differential equations, which are formed in a typical approach such as the Modified Nodal Analysis (MNA)[16] approach:

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

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 very efficiently solved in the time domain by reducing it to a linear algebraic system

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

using Backward Euler (BE) technique with a small fixed time step, h . The BE reduction with fixed time stepping is advantageous for transient simulation since the left hand side (LHS) matrix $(\mathbf{G} + \mathbf{C}/h)$, referred to as the coefficient matrix, is rendered stationary, allowing either pre-processing or factoring of the matrix for a one-time cost and reusing it efficiently to solve the system at successive time points.

When \mathbf{x} consists only of node voltages, as in the case of a modified nodal formulation of a network with R’s, C’s, and only of current sources, the coefficient matrix can be shown to be symmetric and positive definite. The symmetric positive definiteness of the coefficient matrix, which is also very sparse, is especially attractive as the system can now be solved very efficiently using specialized linear system solution techniques, such as Cholesky factorization (direct method) and Conjugate Gradient (iterative method) techniques. The direct method through Cholesky factors is very cost-effective for simulations at multiple time points, as the expensive step of factoring is performed only once and its cost is amortized over multiple time point solutions. Successive solutions would involve only inexpensive forward and backward substitution procedures. Although the macromodeling techniques presented in this paper are suitable for use with either type of solution approach, direct or indirect, we will assume, for simplicity of presentation, that the underlying linear solver is direct.

B. Proposed Approach

The run time and memory requirement for solving a linear system is determined primarily by the size, sparsity, and structure of the coefficient matrix. If the network is very large ($10^7 - 10^8$ nodes), the available physical and virtual memory of the system is insufficient even for loading in the data associated with the network. Even if the base memory requirement is met, memory demand quickly grows during the matrix factorization process, due to new fills being created. Given a reordering scheme, the number of fills created is determined by the initial sparsity and structure of the matrix. The sparsity is given by the ratio of the number of elements in the network to the number of nodes. While tree-like network structures have few fills, mesh structures which are common in power grid networks generally tend to have a large number of fills during factorization. The amount of matrix computation being very sensitive to the sparsity and fill pattern, it is very desirable to have the initial matrix as sparse as possible. The objective of the proposed approach is, hence, twofold - (i) to reduce the size of the problem, and (ii) to maintain a high degree of sparsity in the reduced problem.

The first objective is met by partitioning the given network into subnetworks of manageable size, and solving the network by solving the sub-pieces individually. Since the entire network is tightly connected, we cannot ignore the interaction between the various partitions without incurring significant error. Hence, in order to account for the interactions between partitions, while at the same time not enlarge the size of the problem at hand, we use models for the partitions that capture their behavior as observed at their interface nodes (also called ports). We refer to these models as macromodels. A macromodel is a multi-port linear circuit element that has the same linear relation between the voltages and currents at its ports as the partition itself. Using the macromodels, the original (unpartitioned) network is efficiently solved by replacing the partitions by the respective macromodels, and then solving the combined reduced model.

Besides a memory advantage, the macromodeling approach provides a significant speedup as the creation of macromodels for the partitions can be performed in parallel. However, the gains made from partitioning can be quickly lost if the partitions generate very dense macromodels, and thus increase the storage and computational complexity of the problem. Noting that the port admittance matrix can be fully dense and that the number of matrix entries grows as the square of the number of ports. Our approach addresses this issue in two ways. First, the partitioning is performed strategically to reduce the number of ports, as explained in Section II-E. Then, an optional step of sparsification can be applied to the generated models. The key issue in sparsification is not to compromise accuracy of the final solution. The sparsification technique is covered in Section III.




Fig. 1. Hierarchical power network analysis

C. Hierarchical Modeling

The macromodel approach for power grid analysis is illustrated in Figure 1. Let us consider a division of the entire power network into one global partition and k local partitions. 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*.

Each of the k local grids is modeled as a multi-port linear element with a transfer characteristic given by an equation of the type:

$$\mathbf{I} = \mathbf{A} \cdot \mathbf{V} + \mathbf{S}, \quad \mathbf{I} \in \mathbf{R}^m, \mathbf{A} \in \mathbf{R}^{m \times m}, \mathbf{V} \in \mathbf{R}^m, \mathbf{S} \in \mathbf{R}^m \quad (3)$$

where m is the number of ports in the local grid, \mathbf{A} is the port admittance matrix, \mathbf{V} is the vector of voltages at the ports, \mathbf{I} is the current through the interface between the local and global grids, and \mathbf{S} is a vector of current sources connected between each port and the reference node vector. Vector \mathbf{S} essentially has the effect of moving all the current sources internal in a local grid to the ports of the multi-port model. We refer to the set (\mathbf{A}, \mathbf{S}) as the macromodel of the respective local grid, which is illustrated in Figure 2:




Fig. 2. Macromodel (\mathbf{A}, \mathbf{S})

The macromodel (\mathbf{A}, \mathbf{S}) in equation (3) is obtained through a reduction procedure starting from the modified nodal equations of the local grid. The procedure of deriving the transfer characteristic in Equation (3) from the modified nodal equation is referred to as macromodeling, and will be addressed in detail in Section II-D.

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. This is achieved by combining the coefficient matrix and the right hand side (RHS) current vector of the global grid with the macromodel, (A, \mathbf{S}) ; Equations (3) of each local grid may be stamped into the modified nodal equations of the global grid as follows.

$$\begin{bmatrix} G_{00} & G_{01} & G_{02} & \dots & G_{0k} \\ G_{01}^T & A_1 & G_{12} & \dots & G_{1k} \\ G_{02}^T & G_{12}^T & A_2 & \dots & G_{2k} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ G_{0k}^T & G_{1k}^T & G_{2k}^T & \dots & A_k \end{bmatrix} \begin{bmatrix} \mathbf{V}_0 \\ \mathbf{V}_1 \\ \mathbf{V}_2 \\ \vdots \\ \mathbf{V}_k \end{bmatrix} = \begin{bmatrix} \mathbf{I}_0 \\ -\mathbf{S}_1 \\ -\mathbf{S}_2 \\ \vdots \\ -\mathbf{S}_k \end{bmatrix} \quad (4)$$

In the above equation,

- global nodes are labeled as partition 0
- G_{ij} represents the conductance links between partition i and partition j .
- \mathbf{I}_0 is the vector of currents that flow out of the global nodes.
- \mathbf{S}_i is the constant vector of partition i .
- \mathbf{V}_i is voltage vector of partition i .
- A_i is the port admittance matrix of partition i , where $i \in [1, k]$.

The left hand side matrix in equation (4) should be sparse to permit fast solution. Of the blocks in the matrix, the G_{ij} submatrices are all sparse. The A_i matrices are typically dense, but will be sparsified using a technique described in Section 3. This is a system of $(n_0 + m_1 + m_2 + \dots + m_k)$ linear equations, where n_0 is the number of global nodes and m_i is the number of ports in each partition.

From the above reduction scheme, the voltages and currents in the entire power grid can be solved in the following steps:

- Obtain global grid voltages by solving equation (4).
- For each partition, obtain \mathbf{I} from equation (3) using the port voltages
- Solve equation (2) for each partition using \mathbf{I} on the right hand side, to obtain voltages at the internal nodes of partitions.

The flow of the macromodel approach is illustrated in Figure 3.

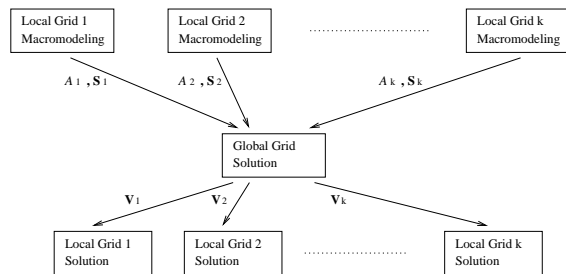


Fig. 3. Flow of the macromodeling approach

D. Macromodeling

Macromodeling is the procedure of deriving Equation (3) from the modified nodal equations of the partition. When

only resistances and capacitances in the local grid are considered, the linear format of the modified nodal equation, Equation (2), could be expressed in the form

$$\tilde{G} \cdot \mathbf{U} = \mathbf{J}, \quad \tilde{G} \in \mathbf{R}^{n \times n}, \mathbf{U} \in \mathbf{R}^n, \mathbf{J} \in \mathbf{R}^n \quad (5)$$

where n is the number of nodes in the local grid, \tilde{G} is the conductance coefficient matrix, \mathbf{U} is the voltage vector of the nodes of the local grid, and \mathbf{J} is the vector of currents that flow out of each node in the local grid. Here, \tilde{G} and \mathbf{J} are equivalent to $\mathbf{G} + \mathbf{C}/h$, $\mathbf{b}(t) + \mathbf{C}/h \cdot \mathbf{x}(t-h)$ in Equation (2), respectively. With a fixed time step h , \tilde{G} is stable for all the time steps while \mathbf{J} needs to be updated at each time step. By splitting the vector \mathbf{U} into the vectors at the internal nodes and ports, Equation (5) can be written as

$$\begin{bmatrix} G_{11} & G_{12} \\ G_{12}^T & G_{22} \end{bmatrix} \begin{bmatrix} \mathbf{U}_1 \\ \mathbf{V} \end{bmatrix} = \begin{bmatrix} \mathbf{J}_1 \\ \mathbf{J}_2 + \mathbf{I} \end{bmatrix} \quad (6)$$

where

- \mathbf{U}_1 and \mathbf{V} are vectors of voltages at the internal nodes and ports respectively
- \mathbf{J}_1 and \mathbf{J}_2 are vectors of current sources connected at the internal nodes and ports respectively
- \mathbf{I} is the vector of currents through the interface
- G_{12} is the admittance of links between the internal nodes and the ports
- G_{11} is the admittance matrix of internal nodes
- G_{22} is the admittance matrix of ports.

From (6), we may rewrite the first set of equations as

$$\mathbf{U}_1 = G_{11}^{-1}(\mathbf{J}_1 - G_{12}\mathbf{V}) \quad (7)$$

Substituting this value of \mathbf{U}_1 into the second equation of (6), we get

$$\mathbf{I} = (G_{22} - G_{12}^T G_{11}^{-1} G_{12})\mathbf{V} + (G_{12}^T G_{11}^{-1} \mathbf{J}_1 - \mathbf{J}_2) \quad (8)$$

Here, $G_{12}^T G_{11}^{-1} \mathbf{J}_1 - \mathbf{J}_2$ is the constant vector \mathbf{S} in Equation (3) and $G_{22} - G_{12}^T G_{11}^{-1} G_{12}$ is the port admittance matrix A in Equation (3).

It may be noted that the pre-multiplication and post-multiplication operations with G_{11}^{-1} can be carried out without explicitly inverting G_{11} , but through multiple invocation of the direct or iterative solver.

The above calculation can be made very efficient by using the fact that the coefficient matrix G is symmetric and positive definite. We show below how A and \mathbf{S} can be computed efficiently from the submatrices of the Cholesky factors, rather than the Cholesky factors themselves. Relating G_{11} , G_{12} , and G_{22} to the submatrices of the Cholesky factors of G , we have

$$\begin{aligned} \begin{bmatrix} G_{11} & G_{12} \\ G_{12}^T & G_{22} \end{bmatrix} &= \begin{bmatrix} L_{11} & 0 \\ L_{21} & L_{22} \end{bmatrix} \begin{bmatrix} L_{11}^T & L_{21}^T \\ 0 & L_{22}^T \end{bmatrix} \\ &= \begin{bmatrix} L_{11}L_{11}^T & L_{11}L_{21}^T \\ L_{21}L_{11}^T & L_{21}L_{21}^T + L_{22}L_{22}^T \end{bmatrix} \end{aligned} \quad (9)$$

Computing A in terms of submatrices of factors, we get

$$\begin{aligned} A &= G_{22} - G_{12}^T G_{11}^{-1} G_{12} \\ &= L_{21}L_{21}^T + L_{22}L_{22}^T - L_{21}L_{11}^T(L_{11}L_{11}^T)^{-1}L_{11}L_{21}^T \\ &= L_{22}L_{22}^T \end{aligned} \quad (10)$$

Similarly, vector \mathbf{S} is given by

$$\begin{aligned}\mathbf{S} &= G_{12}^T G_{11}^{-1} \mathbf{J}_1 - \mathbf{J}_2 \\ &= L_{21} L_{11}^T (L_{11}^T)^{-1} L_{11}^{-1} \mathbf{J}_1 - \mathbf{J}_2 \\ &= L_{21} L_{11}^{-1} \mathbf{J}_1 - \mathbf{J}_2\end{aligned}\quad (11)$$

The above simplified technique reduces the computation cost dramatically compared to direct computation using equation (8) where the most time consuming operation is computing $G_{11}^{-1} G_{12}$. This involves performing m solves of the following system $G_{11} \mathbf{x} = G_{12}$. In the second proposed approach, the timing consuming operation is the single factorization of the G matrix to obtain the partial factors L_{11} , L_{12} and L_{22} using equation (9). Since these factors are triangular, computation of (10) and (11) is comparatively inexpensive. Therefore, this proposed macromodel generation approach has a significant advantage over the direct approach using (8).

E. Partitioning technique

The main difficulty in macromodeling is that the model is often fully dense even though the partition from which it is created itself may be very sparse. Note that the entries of matrix A in equation (8) are admittances of paths between pairs of ports. Thus, a nonzero entry at position (i, j) results if there is a conducting path in the partition between these ports, even though there may not be a direct link between these ports. As a result, the number of nonzero entries in A is $\mathcal{O}(m^2)$, where m is the number of ports, unless the grid inside the partition is heavily fragmented. Nevertheless, there is a substantial win if m^2 is much smaller than the number of nodes in the partition that are abstracted away by this model. In addition, as we can see from Formula (10), (11), the computation cost of macromodel (A, \mathbf{S}) is proportional to the dimension of m . Thus, the key idea in the partitioning strategy is to identify a subnetwork and an interface boundary such that the number of internal nodes is much larger than the square of the number of nodes at the interface. Here, we suggest two heuristic approaches to accomplish this requirement.

First, in some designs, the natural hierarchical boundaries of circuit blocks meet the above criteria. For instance, a large memory array with 3 local metal layers may have several millions of internal nodes, but it may have very few (hundreds of) nodes interfacing with the upper layer of the global grid, and almost none with other circuit blocks. In this case, the partitions could be easily identified manually based on the hierarchy of the design.

However, in some designs, the connection of power grid between the different partitions are very tight and partitioning along the block boundaries may cause a large number of ports. In this case, an automatic partitioning tool is used. The objective is to partition the graph such that the number of ports is minimized, subjected to the constraint that the maximal number of nodes a partition is smaller than a specified number. Furthermore, it is advantageous to partition the graph so that the number of nodes in all local partitions are balanced. In this case, the execution time can be minimized when run parallelly.

Although such a partitioning problem is a well studied, the difficulty with existing solutions is that they have not been designed to handle large graphs with tens of millions of nodes. Therefore, we propose the following methodology, where the graph partitioning is performed in three phases:

- Step 1: Reduce the size of the graph by clustering the tightly-connected neighboring nodes into one node. This significantly reduces the size of the graph.
- Step 2: Partition the reduced graph with an established partitioning tool, such as the state-of-the-art tool, hMETIS [17], [18].
- Step 3: Expand the clustered nodes in partitioned graph by replacing each node with the constituent nodes from its cluster.

Since Step 1 demands only the examination of a node and its immediate neighboring nodes, the graph is not required to be held in memory in its entirety at any point in time. Therefore, portions of the graph can be load and clustered successively, relieving the memory constraint. After the clustering is completed, the graph is sufficiently reduced to hold the entire graph in memory and to execute standard graph partitioning algorithms on it.

F. Analysis of the Computation Cost

In this section we present the computational advantage of macromodeling over the flat model analysis approach.

Suppose the cost of factorizing a matrix is $C_1(l)$, and the cost of one forward and one backward substitution is $C_2(l)$, where l is the size of the matrix, and $C_2(l) \ll C_1(l)$. Let N be the number of nodes in the entire power network.

If no macromodels are used for the power grid analysis, the computation cost of the first run is $C_1(N)$ and the computation cost of a subsequent run is $C_2(N)$. In the macromodeling approach, the computation cost of the first run can be expressed as

$$C_1(n_1) + C_1(n_2) + \dots + C_1(n_k) + C_1(n_0 + m_0 + m_1 + \dots + m_k) + C_2(n_1) + C_2(n_2) + \dots + C_2(n_k) \quad (12)$$

Here, $n_i, i \in [1, k]$ is number of nodes in each partition, n_0 and m_i are defined in Section II-C, and $n_0 + n_1 + n_2 + \dots + n_k = N$. The computation cost from macromodeling is given by $C_1(n_1) + C_1(n_2) + \dots + C_1(n_k)$ by using the simplified macromodeling method described in Section II-D. The cost of finding the solution to the global network is $C_1(n_0 + m_0 + m_1 + \dots + m_k)$ and the cost of solving the local grids is $C_2(n_1) + C_2(n_2) + \dots + C_2(n_k)$, since the factors obtained from macromodeling can be used for solving the local grid.

The computation cost of each subsequent run can therefore be approximated as

$$C_2(n_0 + m_0 + m_1 + \dots + m_k) + 2C_2(n_1) + 2C_2(n_2) + \dots + 2C_2(n_k) \quad (13)$$

where the computation cost of macromodeling is $C_2(n_1) + C_2(n_2) + \dots + C_2(n_k)$ since the A_i 's are unchanged and only the \mathbf{S}_i 's must be recalculated during the subsequent run in macromodeling.

Expressions (12) and (13) provide a rough estimate of computation costs based on the size of the network and its

partitions. In reality, the density of a matrix is an important factor that influences the solution speed. Generally, the conductance matrices for partitions are denser than the conductance matrix of the entire network, and thus the conductance matrix in equation (8) used for the global solution is a dense matrix.

Typically, Cholesky factorization requires $n^3/6$ multiplications and substitution requires $n^2/2$ multiplications. However, the sparsity of the conductance matrix, combined with efficient reordering, enables the computation cost to be less than quadratic with the dimension of the matrix in practice. Since the computation cost for the flat approach remains greater than linear, the computation cost for the macromodel approach will be lower than that for the flat analysis even with the overheads associated with partitioning, if the design is sufficiently large.

Most important of all, the *divide and conquer* procedure applied to the power network makes parallel execution of power network simulation possible. During parallel execution, the execution time of the first run is given by

$$\max(C_1(n_1), C_1(n_2), \dots, C_1(n_k)) + C_1(n_0 + m_0 + m_1 + \dots + m_k) \\ + \max(C_2(n_1), C_2(n_2), \dots, C_2(n_k))$$

where $C_1(n_0 + m_0 + m_1 + \dots + m_k)$ is the global solution time, $\max(C_1(n_1), C_1(n_2), \dots, C_1(n_k))$ represents the maximum execution time among macromodeling of partitions and $\max(C_2(n_1), C_2(n_2), \dots, C_2(n_k))$ represents the maximal execution time out of partition solutions. Similarly, the execution time of the subsequent run is given by

$$C_2(n_0 + m_0 + \dots + m_k) + 2 \times \max(C_2(n_1), C_2(n_2), \dots, C_2(n_k))$$

Moreover, the memory requirement with macromodels is the maximum memory required for solving any partition, rather than the sum of memory requirement of each partition.

Besides run time and memory advantage, macromodeling provides a certain flexibility to a design/analysis situation so that significant analysis effort can be saved. Given below are few examples of design/analysis situations when such flexibility is useful.

Example-1: When a designer is interested in the detailed analysis of only a specific circuit block, significant design time is saved by not simulating the other partitions, while accounting accurately for the effect of these other blocks on the block of interest.

Example-2: A designer knows *a priori* in which circuit block or blocks the worst drop is to be expected, and the objective of the analysis is to only to find the worst IR drop estimate for the design. Then, it will be necessary to simulate only a few blocks (partitions) in the last step of the macromodel approach.

Example-3: The process of fixing problems in a power grid is usually an iterative one. The process consists of detecting an error, making local changes to the grid to correct the problem, and re-running the analysis. In this case, only the macromodeling of the partition whose grid was changed needs to be recalculated. The speed-up in analysis due to this makes it possible for the designer to fix the problems interactively with the analysis tool.

III. SPARSIFICATION OF MACROMODELS

In Section II-E, we pointed out that the number of entries in the macromodel has $\mathcal{O}(m^2)$ complexity for model size m . Although the macromodels reduce the size of the system to the smaller system described in equation (4), the density of the coefficient matrix of equation (4) increases considerably due to the density of the A_i submatrices. For an iterative solver, this is undesirable as the number of floating point operations (FLOPs) to solve the system increases. For a direct solver, this affects both the number of required FLOPs, as well as the memory required for factorization. The additional memory demand is caused by excessive fills created by the dense parts during factoring. Therefore, to derive the most benefit from the macromodeling approach, it is important that the coefficient matrix in equation (4) is kept sparse. While the partitioning strategy explained in Section II-E is a natural way of achieving this, other sparsification techniques in conjunction with good partitioning schemes are useful for making the macromodeling approach effective. In this section, we present a novel technique to sparsify the port admittance matrices of the macromodels.

Our sparsification method is motivated by the observation that although the matrix A is dense, it consists of a large number of values that are numerically small and will have little influence on the results if approximated to zero. We provide an algorithm to sparsify the coefficient matrix A by dropping some of its entries, while keeping the error introduced by this process below a specified value. The proposed sparsification technique also preserves the symmetry and the positive definite property of the matrix. Note that the sparsification procedure needs to be performed only once (during the first run).

A. Problem Definition

The problem is stated as follows: Given the transfer characteristic equation of each partition

$$\begin{bmatrix} i_1 \\ i_2 \\ i_3 \\ \vdots \\ i_m \end{bmatrix} = \begin{bmatrix} a_{1,1} & a_{1,2} & a_{1,3} & \dots & a_{1,m} \\ a_{2,1} & a_{2,2} & a_{2,3} & \dots & a_{2,m} \\ a_{3,1} & a_{3,2} & a_{3,3} & \dots & a_{3,m} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ a_{m,1} & a_{m,2} & a_{m,3} & \dots & a_{m,m} \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \\ v_3 \\ \vdots \\ v_m \end{bmatrix} + \begin{bmatrix} s_1 \\ s_2 \\ s_3 \\ \vdots \\ s_m \end{bmatrix}, \quad (14)$$

and given a nominal voltage value of v_j ($j \in [1, m]$): B , and the error threshold of i_j ($j \in [1, m]$): e_j , transform equation (14) into

$$\begin{bmatrix} i'_1 \\ i'_2 \\ i'_3 \\ \vdots \\ i'_m \end{bmatrix} = \begin{bmatrix} a'_{1,1} & a'_{1,2} & a'_{1,3} & \dots & a'_{1,m} \\ a'_{2,1} & a'_{2,2} & a'_{2,3} & \dots & a'_{2,m} \\ a'_{3,1} & a'_{3,2} & a'_{3,3} & \dots & a'_{3,m} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ a'_{m,1} & a'_{m,2} & a'_{m,3} & \dots & a'_{m,m} \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \\ v_3 \\ \vdots \\ v_m \end{bmatrix} + \begin{bmatrix} s_1 \\ s_2 \\ s_3 \\ \vdots \\ s_m \end{bmatrix}$$

such that the number of $a'_{j,k}$ ($j \neq k, a'_{j,k} = 0$) is maximized, subject to $|i'_j - i_j| \leq e_j$ ($j \in [1, m]$) and $a'_{j,k} = a'_{k,j}$ (maintaining matrix symmetry).

Here, the error threshold e_j can be defined as $|s_j \times x\%|$, where s_j , $1 \leq j \leq m$, is as defined in Equation (14), and $x\%$ is the user-defined error limit. As we can see from Equations (4), voltage and current follow a linear relationship.

If each current in the system is within the error bound, $x\%$, the error bound for voltage is $x\%$. Therefore, from the specification that e_j is no more than $|s_j \times x\%|$, we can tell that the error bound of $x\%$ is guaranteed.

B. Problem Formulation

This problem can be formulated into a 0-1 multidimensional knapsack problem [19], [20]. In this section, we describe the transformation from the above problem to the knapsack problem.

The task here involves zeroing out off-diagonal elements of the matrix A . It is easy to show that these sparsification operations maintain the positive definite property of the matrix. To see this, we note first that the partition can be thought of as being purely resistive (for example, any capacitors are linearized). Given this "resistive" network, one may build an equivalent network of a set of equivalent resistances R_{jk} between each pair of ports j and k . The matrix A is then simply the conductance matrix for this network of R_{jk} 's, and is therefore diagonally dominant. This leads to three conclusions: (1) all off-diagonal elements must be non-positive, (2) zeroing out off-diagonal elements of A maintains the diagonal dominance of the matrix, and therefore its positive definite property, and (3) the resulting IR drop voltages after sparsification are conservative since some resistances are ignored.

The problem formulation is described as follows. First consider the maximum error that an element of matrix $a_{j,k}$ can cause if it is rounded off to 0. Since B is positive and $a_{j,k} \leq 0, j \neq k$, the maximum negative error caused by rounding off $a_{j,k}$, $en_{j,k}$, is given by

$$en_{j,k} = a_{j,k} * B, j \neq k$$

Let $X_{j,k}$ represent a Boolean value, 1 when element $a_{j,k}$ is rounded to zero, and 0 otherwise. The matrix sparsification problem can be formulated as 0-1 knapsack problem as follows.

$$\begin{aligned} & \text{Maximize} && z(x) = \sum_{k=1}^m \sum_{j=1}^k X_{j,k} \\ \text{subject to} && - \sum_{k=1}^{j-1} en_{j,k} * X_{j,k} - \sum_{k=j+1}^m en_{j,k} * X_{j,k} \leq e_j, j \in [1, m] \\ && X_{j,k} \in \{0, 1\} \text{ for all } X_{j,k}, j < k \end{aligned}$$

In the above formulation, the indices of the variables $X_{j,k}$ are required to satisfy the relation $j < k$, so that $X_{j,k} = 1$ indicates the rounding-off of both $a_{j,k}$ and $a_{k,j}$ to maintain the symmetricity of A . Therefore, the resulting sparsified matrix is symmetric and positive definite.

The 0-1 knapsack problem can be solved optimally either by dynamic programming or using an integer programming solver. In our implementation, we use the latter, but with some modifications for speed considerations. First, we relax the integer requirement and solve the fractional knapsack problem using a linear programming solver [21]. Next, the fractional x_{jk} 's are sorted, and the corresponding entries in the matrix are successively set to zero until the maximum error in $|i_j|$ reaches the specified limit, e_j . The run time of sparsification is very small compared to the simulation time of the linear solver. Therefore, it impacts the speedup very slightly.

IV. EXPERIMENTAL RESULTS

The hierarchical analysis method using macromodels was implemented using C and embedded in an existing industrial power analysis tool [1]. An efficient direct linear solver based on Cholesky factors was used in all the experiments. The extracted power grids of six high performance general purpose/DSP microprocessor chips were used to benchmark the performance of macromodeling (Tables I and II) and sparsification (Table III) techniques. Chips 1, 2 and 4 are DSP and communication chips whose power grids are implemented in 3 layers of metal. Chips 3, 5 and 6 are high performance microprocessor chips using 5 and 6 metal layers respectively. The analyses were carried out on the Sun workstations whose clock frequencies range from 300MHz to 460MHz and whose memories are around 1 - 4 Gigabyte. The run time measures used for comparison are based on the actual time required to complete the task.

A. Performance of Macromodeling Technique

Table I compares the performance of the proposed hierarchical approach using macromodels with that of the nonhierarchical approach. Two metrics are compared: the peak memory demand and the total run-time. The number of nodes, in millions, for the entire power network is given in Column 2, the number of partitions used by the macromodeling algorithm are listed in Column 3, and the number of nodes in the largest partition is given in Column 4. Columns 5 and 6 show the total time, in minutes, required by the hierarchical approach, while Column 8 shows the total time taken for completing the analysis on the flat model. The run-time in Column 5 corresponds to the cases when the macromodels for the various partitions were generated serially on a single computer, whereas Column 6 is for the cases when these computations are performed in parallel. The run-time reported in this table is the time taken for analyzing the power network at the first time point in a sequence of simulations. Columns 4 and 9 show the peak memory demand, in Gigabytes, during the analysis with macromodels and without macromodels, respectively. Chip-4, 5, 6 can not be solved without macromodeling due to their large size.

It is evident from the above table that the problem size tackled with the proposed approach is substantially reduced from the original problem. This is the primary goal of the proposed approach so that a chip-level analysis of very large designs is made possible. Based on the benchmarks, it can be seen that the size of the linear system that needs to be solved with the new approach is about 10X smaller than the traditional approach.

The effect of problem size reduction is clearly reflected in the peak memory requirements of the different approaches shown in the table. Again, a 10X to 20X reduction in memory requirement is seen possible with the hierarchical approach. This implies that, given sufficient computing resources, the new method enables the analysis of much larger designs that are now common. For the design such as Chip 4, 5 and 6, flat analysis is infeasible due to limited memory resources. Therefore, our proposed macromodel-

TABLE I
RUN-TIME AND MEMORY COMPARISON FOR THE FIRST SIMULATION

Chip	#nodes (millions)	With macromodel					Without macromodel	
		# part	#nodes(max) (millions)	Total Run-time		Peak Memory (GB)	Run-time (minutes)	Peak Memory (GB)
				Serial(min)	Parallel(min)			
Chip-1	3.9	12	0.40	43	7	0.2	93	1.5
Chip-2	2.7	9	0.58	25	6	0.3	57	1.2
Chip-3	7.5	11	0.79	136	26	0.4	629	2.6
Chip-4	20.0	7	3.5	444	152	1.3	-	-
Chip-5	41.5	20	4.1	995	93	1.8	-	-
Chip-6	63.5	30	3.3	1796	82	1.6	-	-

ing approach is the only approach for analyzing these large networks in detail. From the results, we can see that without macromodels the run time can be several hours (e.g., 10.5 hours for Chip 3) for a supply network with millions of nodes. As a result of reducing the size complexity, the run-time is reduced by a factor of 2X to 5X even when the macromodels are computed one after another on a single computer. The run-time is dramatically reduced by 10X to 23X, if the parallelism created by the macromodel is utilized. It is noteworthy that the speedups improve with the size of the circuit under consideration. Also, note that given additional computing resources, the proposed approach allows the analysis of large designs (such as Chip 6) with the same run time as smaller design. Therefore, our approach provides a scalable solution that is practical in a parallel or a network of workstations environment.

TABLE II
COMPARISON OF RUN TIME FOR 1000 SUBSEQUENT SIMULATIONS

Chip	Without macromodel run-time (hours)	With macromodel run-time	
		Total(hours)	Parallel(hours)
Chip-1	8.4	28.0	4.0
Chip-2	8.0	22.5	4.4
Chip-3	33.7	43.5	6.6

Table II compares the performance of the two approaches based on the time required to perform simulations of 1000 successive time points, after the first time point. Thus, the shown run-times are independent of the time taken to generate the macromodels. Column 2 shows the run time without macromodels. For the hierarchical approach, run-times for both serial (Column 3) and parallel (Column 4) execution are shown. Since the memory requirement of these runs is less than that of the first run, these figures are omitted in Table II.

The hierarchical approach executed in serial mode recorded unfavorable run-times for the benchmarks in the cases where it was feasible to obtain the results for the non-hierarchical approach using the available computational resources. However, the disparity in run-times between the nonhierarchical and hierarchical approach (in serial mode) diminishes as the size of the original network becomes larger, as evidenced from the results for Chip-3, which has 7.5 million nodes. This behavior is not unexpected, and can be explained by the fact that the overhead associated with computing the \mathbf{S} vector for each partition at every

time step, and back-solving each partition again in the final step of the solution, is a dominant factor. This behavior is exhibited for networks up-to a certain size, where the original matrix and the reduced matrix do not differ greatly in terms of the time required for a back-solve. However, as the network becomes larger, the difference in problem sizes with and without macromodels are significantly different, and the overhead cost of handling the partitions becomes negligible in the overall cost. As a result, the hierarchical approach becomes favorable for large networks even in the serial execution mode.

The run-time advantage of parallel execution mode is very clear from Table II. Results show that the parallel execution utilizing hierarchy is 1.8 – 5.1 times faster than the nonhierarchical approach. As designers would like to simulate the power grid with long traces of current signatures in order to obtain good coverage of the IR-drop situations, efficiency of simulation in this phase is crucial. The parallel execution mode, as well as the flexibility in the hierarchical analysis and its ability to analyze grids with tens of millions of nodes, make the hierarchical analysis approach extremely attractive.

B. Performance of the Sparsification Technique

The sparsification procedure described in Section III reduces the number of nonzero elements while maintaining an acceptable level of accuracy. In our implementation, the specified error e_j is defined as $e_j = \max(const, |s_j \times x\%|)$, where $const$ is a small positive constant, s_j , $1 \leq j \leq m$, is as defined in Equation (14), and $x\%$ is the user-defined error limit, which is typically 0% – 10%. The sparsification technique was implemented using a linear programming solver *lp_solve_2.3*[21].

Table III reports the sparsity and run-time improvements achieved for two benchmark examples, analyzed at different levels of accuracy. The second column in the table shows the voltage value of the clean power supply and the value of the maximum voltage drop observed in the circuit. Columns 3 reports the number of nodes and the total number of ports respectively in the global grid. The number of nonzero elements in the coefficient matrix of equation (4) are shown in Column 4. Column 5 shows the maximum voltage error caused by the sparsification procedure. The ratio of maximum observed error in voltage to the maximum voltage drop is shown in column 6. Finally, column 7 reports the time required to solve for the global node volt-

ages with Equation (4). In Column 7, the execution time for LP solver is not included, since generally it completes in seconds.

For each benchmark, the proposed sparsification technique was tested at four levels of accuracy. The benchmark Chip-7 is a 6-layer, mesh type, power grid. Its power grid is much denser than the other examples, and this example also has some partitions with large number of ports. As a result, the coefficient matrix obtained for this example could not be solved with the available computing resources without sparsification.

The results clearly show that the sparsity of the coefficient matrix is improved by as much as 11X, incurring only 2.6% error in the final results. The improved sparsity improved the run-time for the dense example, Chip-7 significantly, besides greatly reducing the memory requirement.

TABLE III
THE EFFECT OF SPARSIFICATION


Chip	Vdd/ Max- dp(v)	#node/ #port	#non- zeroes	Max-err (v)	Error %	Run- time (secs)
Chip -3	2.0/ 0.12	23261/ 379	106909	0.0	0.0%	38.3
			98505	0.000043	0.04%	36.0
			98253	0.00058	0.5%	32.4
			98005	0.0013	1%	30.6
Chip -7	1.8/ 0.02	2932/ 2849	2067572	0.0	0.0%	-
			366612	0.000045	0.2%	36.5
			249588	0.00021	1.0%	24.4
			197868	0.00051	2.6%	20.1

V. CONCLUSIONS AND FUTURE WORK

In this paper, we have presented a hierarchical power network analysis method using novel macromodeling and matrix sparsification techniques, where the macromodeling method gives the exact same solution as the flat method, and the matrix sparsification technique guarantees a conservative approximation. The proposed techniques were shown to gain significant memory and run-time advantages over the traditional approach of analyzing the power network without using the hierarchy. The experimental results based on analyzing the entire power network of six high performance microprocessor designs confirmed these claims. The hierarchical analysis approach shows excellent promise as a viable alternative to the traditional nonhierarchical analysis method, capable of handling the increasing size of power grids in modern microprocessors.


REFERENCES

- [1] A. Dharchoudhury, R. Panda, D. Blaauw, R. Vaidyanathan, B. Tutuianu, and D. Bearden, "Design and analysis of power distribution networks in PowerPC microprocessors," in *Proceedings of the ACM/IEEE Design Automation Conference*, pp. 738–743, 1998.
- [2] G. Steele, D. Overhauser, S. Rochel, and Z. Hussain, "Full-chip verification methods for DSM power distribution systems," in *Proceedings of the ACM/IEEE Design Automation Conference*, pp. 744–749, 1998.
- [3] H. Chen and D. Ling, "Power supply noise analysis methodology for deep-submicron VLSI chip design," in *Proceedings of the ACM/IEEE Design Automation Conference*, pp. 638–643, 1997.
- [4] S. Taylor, "The challenge of designing global signals in UDSM CMOS," in *Proceedings of the IEEE Custom Integrated Circuits Conference*, pp. 429–435, 1999.
- [5] S. Zhao, K. Roy, and C.-K. Koh, "Frequency domain analysis of switching noise on power supply network," in *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design*, pp. 487–492, 2000.
- [6] G. Bai, S. Bobba, and I. N. Hajj, "Simulation and optimization of the power distribution network in VLSI circuits," in *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design*, pp. 481–486, 2000.
- [7] 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," in *Proceedings of the International Symposium of Low Power Electronics and Design*, pp. 179–184, 2000.
- [8] S. R. Nassif and J. N. Kozhaya, "Fast power grid simulation," in *Proceedings of the ACM/IEEE Design Automation Conference*, pp. 156–161, 2000.
- [9] D. Blaauw, R. Panda, and R. Chaudhry, "Chapter 24: Design and analysis of power distribution networks," in *Design of High-Performance Microprocessor Circuits*, IEEE Press, New York, pp. 499–520, 2001.
- [10] H. Kriplani, F. Najm, and I. Hajj, "Pattern independent minimum current estimation in power and ground buses of CMOS VLSI circuits," *IEEE Transactions on Computer-Aided Design*, vol. 14, no. 8, pp. 998–1012, 1995.
- [11] A. Krstic and K. Cheng, "Vector generation for maximum instantaneous current through supply lines for CMOS circuits," in *Proceedings of the ACM/IEEE Design Automation Conference*, pp. 383–388, 1997.
- [12] Y.-M. Jiang, T. Young, and K. Cheng, "VIP – an input pattern generator for identifying critical voltage drop for deep sub-micron designs," in *Proceedings of the International Symposium of Low Power Electronics and Design*, pp. 156–161, 1999.
- [13] R. Chaudhry, D. Blaauw, R. Panda, and T. Edwards, "Current signature compression for ir-drop analysis," in *Proceedings of the ACM/IEEE Design Automation Conference*, pp. 162–167, 2000.
- [14] S. Bobba and I. N. Hajj, "Maximum voltage variation in the power distribution network of VLSI circuits with RLC models," in *Proceedings of the International Symposium of Low Power Electronics and Design*, 2001.
- [15] G. H. Golub and C. F. V. Loan, *Matrix Computations*. Baltimore, MD: The Johns Hopkins University Press, 1984.
- [16] C. Ho, A. Ruehli, and P. Brennan, "The modified nodal approach to network analysis," *IEEE Trans. Circuits and Systems*, vol. CAS-22, no. 6, pp. 504–509, 1975.
- [17] G. Karypis and V. Kumar, *hMETIS: A Hypergraph Partitioning Package*. University of Minnesota, 1998.
- [18] G. Karpis, R. Aggarwal, V. Kumar, and S. Shekhar, "Multi-level hypergraph partitioning: Application in VLSI domain," in *Proceedings of the ACM/IEEE Design Automation Conference*, pp. 526–529, 1997.
- [19] T. H. Cormen, C. E. Leiserson, and R. L. Rivest, *Introduction to Algorithms*. New York, NY: McGraw-Hill, 1990.
- [20] K. G. Murty, *Operations Research Deterministic Optimization Models*. Englewood Cliffs, NJ: Prentice Hall, 1995.
- [21] M. R. C. M. Berkelaar, "LP SOLVE 2.3 Users' Manual," 1998.




Min Zhao received her B.S. in computer and its application, M.S. in electrical transmission and its automation, from Dalian Maritime University, Dalian, China, in 1993, 1996 and the Ph.D. degree in electrical engineering from University of Minnesota in 1999. She is currently an Electrical Engineer in Motorola, Austin, Texas. Her research interest is VLSI CAD for domino logic, technology mapping, power grid analysis and inductance analysis.

and member of the Executive Committee the ACM/IEEE Design Automation Conference 2001. He was the Technical Program Chair for the International Symposium on Low Power Electronic and Design in 1999 and the General Chair for this symposium in 2000.




Rajendran Panda received the B.E. (Honors) degree from Madurai University in 1978, the LL.B. degree from Bangalore University in 1984, and the Ph.D. degree from the University of Illinois at Urbana-Champaign in 1996. From 1978 to 1991, he was with the Bharat Heavy Electrical Industries, Bangalore. He is currently with Motorola, Austin, managing the High Performance Design Technology group. His work has focussed on low power design, power supply and signal integrity issues, and

circuit optimization. He has published extensively and holds four patents in these areas.



Sachin Sapatnekar received the B.Tech. degree from the Indian Institute of Technology, Bombay in 1987, the M.S. degree from Syracuse University in 1989, and the Ph.D. degree from the University of Illinois at Urbana-Champaign in 1992. From 1992 to 1997, he was an assistant professor in the Department of Electrical and Computer Engineering at Iowa State University. He is currently an Associate Professor in the Department of Electrical and Computer Engineering at the University of Minnesota.

He has coauthored two books, "Timing Analysis and Optimization of Sequential Circuits," and "Design Automation for Timing-Driven Layout Synthesis," and is a co-editor of a forthcoming volume, "Layout Optimizations in VLSI Designs," all published by Kluwer. He has been an Associate Editor for the IEEE Transactions on VLSI Systems and the IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing, has served on the Technical Program Committee for various conferences, and as Technical Program and General Chair for Tau and ISPD. He is a recipient of the NSF Career Award and best paper awards at DAC 1997, ICCD 1998, and DAC 2001.



David Blaauw received his B.S. in Physics and Computer Science from Duke University in 1986, his M.S. in Computer Science from the University of Illinois, Urbana, in 1988 and his Ph.D. in Computer Science from the University of Illinois, Urbana, in 1991. He worked at the Engineering Accelerator Technology Division, IBM Corporation, Endicott, as a Development Staff Member, until August 1993. From 1993 till August 2001, he has worked for Motorola, Inc. in Austin, TX, where he was the manager

of the High Performance Design Technology group. His work has focussed on VLSI design and CAD with particular emphasis on circuit analysis and optimization problems for high performance microprocessor designs and low power DSPs. Since August 2001, he has been on the faculty at the University of Michigan as an Associate Professor.

David Blaauw has published extensively and holds 13 issued patents. He received the Distinguished Innovator Award at Motorola in 2000 and the best paper award at the ACM/IEEE Design Automation Conference in 2000. He has given a number of invited talks and tutorials and is a member of the technical program committee of a number of conferences. He is the Technical Program Co-Chair