

Addressing Substrate Coupling in Mixed-Mode IC's: Simulation and Power Distribution Synthesis

Balsha R. Stanisic, *Member, IEEE*, Nishath K. Verghese, *Student Member, IEEE*, Rob A. Rutenbar, *Senior Member, IEEE*, L. Richard Carley, *Senior Member, IEEE*, and David J. Allstot, *Fellow, IEEE*

Abstract—This paper describes new techniques for the simulation and power distribution synthesis of mixed analog/digital integrated circuits considering the parasitic coupling of noise through the common substrate. By spatially discretizing a simplified form of Maxwell's equations, a three-dimensional linear mesh model of the substrate is developed. For simulation, a macromodel of the *fine* substrate mesh is formulated and a modified version of SPICE3 is used to simulate the electrical circuit coupled with the macromodel. For synthesis, a *coarse* substrate mesh, and interconnect models are used to couple linear macromodels of circuit functional blocks. Asymptotic Waveform Evaluation (AWE) is used to evaluate the electrical behavior of the network at every iteration in the synthesis process. Macromodel simulations are significantly faster than device level simulations and compare accurately to measured results. Synthesis results demonstrate the critical need to constrain substrate noise and simultaneously optimize power bus geometry and pad assignment to meet performance targets.

I. INTRODUCTION

THE push for reduced cost, more compact circuit boards, and added customer features has provided incentives for the inclusion of analog functions on primarily digital MOS integrated circuits. This has prompted the development of a new generation of electrically cognizant analysis and synthesis tools. To date however, these tools have largely ignored the critical effect of the common chip substrate.

The common substrate couples noise between the on-chip digital and analog circuits that corrupts low-level analog signals, impairing the performance of mixed-signal IC's. The potential crosstalk problem of noise finding its way into sensitive analog circuitry has traditionally been handled effectively by judicious use of multiple power bus layouts and desensitized analog circuitry [1]. However, in the future, digital speeds will increase, additional analog circuitry will be included, chips will become more densely packed, interconnect layers will be added and analog resolution will be increased. In consequence, the noise crosstalk problem will worsen and designer skills will be severely taxed.

As an illustration of the significance of the problem, consider the performance of single-bit sigma-delta D/A converters which have traditionally been limited to a resolution of 16

b. Recently [2], 18-b resolution was achieved with the same topology by separating the analog and digital functional blocks onto two different IC's, with the aim of alleviating the substrate coupling problem. As technologies continue to scale and compactness of design becomes more important, analog designers will not have the leeway to fabricate their functional blocks on independent IC's. Performance degradation due to substrate noise will become difficult to control and even more difficult to predict. The need for high performance simulation and synthesis tools to identify and help avoid the problem is increasingly evident in the industry today [3].

Accurate simulation of substrate coupling has only recently begun to receive attention [4]–[6]. A device simulation program has been used to study the mechanism of substrate coupling noise [4], [5]; however, its applicability for simulating integrated circuits is limited by the long simulation times required. A single node substrate model has also previously been used to simulate substrate coupling [3], [4]. This approach, however, is applicable only to technologies employing a lightly doped epitaxial layer on a heavily doped substrate. To date, we are unaware of any process-independent tool efficient enough to simulate substrate coupling effects in large integrated circuits, or any tool able to synthesize power distribution cognizant of these same substrate effects.

In this paper we present such a process-independent simulation strategy for substrate coupling effects that is not only valid for all silicon IC technologies but also fast and accurate. We also present a power distribution synthesis strategy that incorporates substrate coupling effects and simultaneously optimizes power bus geometry and power I/O pad assignment. Section II introduces the substrate coupling problem and develops the basic substrate model. Section III describes the simulation techniques used: macromodeling and circuit simulation with macromodels. Section IV compares simulation results to those from a device simulation program and also compares simulation results with reported measurements on a test chip [4]. Section V introduces the power distribution synthesis strategy, Section VI describes the simultaneous power I/O pad assignment and power bus synthesis formulation, and Section VII describes the optimization approach. Sections VIII and IX describe the electrical modeling and evaluation strategies in the synthesis process, respectively. Section X shows synthesis results for several mixed-signal examples, and compares synthesis and simulation results to those reported in [4].

Manuscript received July 28, 1993; revised October 3, 1993. This work was supported by International Business Machines and the Semiconductor Research Corporation under Contract 93-DC-068.

The authors are with the Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213.
IEEE Log Number 9214786.




Fig. 1. The substrate coupling problem.

II. MODELING THE SUBSTRATE

The general substrate coupling problem is illustrated in Fig. 1 in which digital switching nodes are capacitively coupled to the substrate through junction capacitances and interconnect/bonding pad capacitances, causing fluctuations in the underlying voltage. As a result of these fluctuations, a substrate current pulse flows between the switching node and the surrounding substrate contacts. The switching-induced current flow causes the substrate potential underlying critical transistors in the path of this flow to change. As a result of the body effect and the junction capacitances of a sensitive transistor, changes in the backgate voltage induce noise spikes in its drain current and consequently its drain voltage. To simulate this phenomenon, a necessary first step is the development of a suitable model for the substrate.

Outside the diffusion/active areas and contact areas, the substrate can be treated as consisting of layers of uniformly doped semiconductor material of varying doping densities. In these layers, a simplified form of Maxwell's equations can be formulated, ignoring the influence of magnetic fields and using the identity $\nabla \cdot (\nabla \times a) = 0$:

$$\epsilon \cdot \frac{\partial}{\partial t} (\nabla \cdot E) + \frac{1}{\rho} \nabla \cdot E = 0 \quad (1)$$

where E is the electric field intensity vector, and ρ and ϵ are the sheet resistivity and dielectric constant of the semiconductor, respectively.

There are two general approaches to solve (1): *analytical* and *numerical*. The *analytical* approach involves a search for an exact solution for structures with mathematically tractable (usually rectangular) geometries. Although it is possible to obtain a closed-form analytical solution of (1) for a simple structure, analytical solutions of complicated geometries do not generally exist. Hence *numerical* techniques must be used to solve the geometries encountered in typical integrated circuits. To employ numerical techniques, the substrate is spatially discretized using a simple box integration technique [7].

From Gauss' law,

$$\nabla \cdot E = k \quad (2)$$

where E is the electric field intensity and $k = \rho'/\epsilon$ where ρ' is the charge density of the material. Integrating $\nabla \cdot E$ over a volume Ω_i surrounding node i as shown in Fig. 2,

$$\int_{\Omega_i} \nabla \cdot E d\Omega = \int_{\Omega_i} k d\Omega. \quad (3)$$




Fig. 2. A control volume in the box integration technique.




Fig. 3. Resistances and capacitances around a mesh node in the electrical substrate mesh.

From the divergence theorem,

$$\int_{S_i} E dS = \int_{\Omega_i} k d\Omega \quad (4)$$

where S_i is the surface area of the cube shown in Fig. 2. The integral on the left side of (4) can be approximated as

$$\sum_j E_{ij} \cdot S_{ij} = \sum_j E_{ij} \cdot w_{ij} d_{ij} = k \cdot \Omega_i \quad (5)$$

and hence,

$$\nabla \cdot E = k = \frac{1}{\Omega_i} \sum_j E_{ij} \cdot w_{ij} d_{ij}. \quad (6)$$

Using (6) in (1) and noting that $E_{ij} = V_i - V_j/h_{ij}$ in Fig. 2, (1) reduces to

$$\sum_j \left[\frac{(V_i - V_j)}{R_{ij}} + C_{ij} \left(\frac{\partial V_i}{\partial t} - \frac{\partial V_j}{\partial t} \right) \right] = 0 \quad (7)$$

where $R_{ij} = \rho h_{ij}/w_{ij} d_{ij}$ and $C_{ij} = \epsilon w_{ij} d_{ij}/h_{ij}$ as modeled with lumped circuit elements in Fig. 3.

Since the relaxation time of the substrate (outside of active areas and well diffusions) given by $\tau = \rho\epsilon$ is of the order of $10^{-11} s$ (with $\rho = 15\Omega - cm$ and $\epsilon_0 = 11.9$), it is reasonable to neglect intrinsic substrate capacitances for operating speeds of up to a few GHz and switching times of the order of 0.1 ns. Moreover, if the capacitances to the substrate introduced




Fig. 4. The peak-peak noise voltage at the sensitive node as a function of the number of grids in the z direction (with the grid density in the x and y directions fixed) for different substrate thicknesses.

by the depletion regions of well diffusions and interconnects overlying field oxide can be accurately modeled as lumped circuit elements outside the mesh, the substrate can be modeled as a purely resistive mesh. (Note that junction capacitances of active devices are already modeled outside the mesh as lumped capacitances.)

Although the electric field varies nonlinearly as a function of distance, the box integration method approximates this variation as a piecewise constant function. In regions where the gradient of the electric field is high, it is necessary to use fine grids to accurately approximate the nonlinearity of the electric field. Elsewhere, coarse grids can be used to reduce the overall number of grids. However, since the field intensity cannot be determined before discretization, the density of grids needed is not known *a priori*.

We determine the density of grids for a substrate with either contacts/diffusions on the surface or a backside contact using the setup of Fig. 4. It consists of a noisy node and a sensitive node separated by a fixed lateral distance. With a fixed number of uniformly distributed grids in the x and y directions between these nodes, the peak-peak noise voltage at the sensitive node is determined as a function of the grid density in the z direction for different substrate thicknesses. For the geometry shown in Fig. 4, noise-coupling is a strong function of grid density in the z direction as expected, since the gradient of the electric field is high in that direction. The number of grids needed (in the z direction) is also a strong function of the substrate thickness.

Fig. 5 shows the effects of grid density in the x and y directions with a fixed grid density in the z direction. Since the fixed boundaries are in the xy plane, the noise coupling is not as sensitive to the number of grids in the x and y directions; consequently, coarse grids are used in these directions. Thus, we empirically determine a grid density based on total substrate thickness from the results of Figs. 4 and 5. While a large body of research on the subject of automatic mesh generation exists [28], we have avoided it in our work since we have found that in the case of a linear substrate model with contacts/diffusions at the surface or a backside contact, it is relatively easy to determine the density of grids required *a priori* using the technique mentioned above.




Fig. 5. The peak-peak noise voltage at the sensitive node as a function of the number of grids in the x and y directions (with the grid density in the z direction fixed) for different amounts of lateral separation.

III. SIMULATION TECHNIQUES

Once the substrate is discretized into a 3-D mesh it is necessary to solve it in conjunction with the electrical circuit. However, solving a 3-D mesh using traditional variable time-step trapezoidal integration techniques is prohibitive in terms of cpu time and memory since the substrate mesh has in general many more nodes than the rest of the electrical circuit. By macromodeling the mesh and performing transient analysis with the macromodel, computations for the mesh need to be performed only on those nodes that physically connect the electrical circuit to the substrate. Since the objective is to use SPICE3 [8] as a basic framework that is modified to perform substrate coupling simulations, an admittance (Y -parameter) macromodel of the substrate mesh is formulated. In computing the admittance macromodel, the matrix solution technique utilized plays a crucial role in determining the cpu time requirement. To optimize the simulation time required, specialized matrix solution techniques are used as described below.

A. Matrix Solution

The matrix solution of a large 3-D mesh network using the traditional direct solution method is known to require a large amount of cpu time and memory for LU factorization even when reordering and sparse matrix techniques are used. Since the matrix resulting from the substrate mesh is both strictly diagonally dominant and symmetric (with a nodal analysis formulation), the incomplete Choleski conjugate gradient (ICCG) iterative method [13] with an ILU (1) (incomplete LU factorization with diagonal correction only) preconditioner has been adopted as the matrix solver. Unlike other iterative methods, the ICCG method always converges to the exact solution in at most n iterations where n is the order of the matrix. Experience with ICCG shows typical iteration counts to be far fewer because of preconditioning.

B. AWE Admittance Macromodel

If the substrate is modeled as an RC mesh, Asymptotic Waveform Evaluation (AWE) can be used to determine its admittance macromodel. AWE is a technique to approximate the

behavior of a linear(ized) circuit using a few dominant poles and residues in either the time or the frequency domains [9]. The AWE technique involves the computation of the circuit moments in an efficient and recursive manner. A reduced-order pole-residue model of the circuit transfer function is determined from the circuit moments using a form of Pade approximation [10]. Similarly, AWE approximations to the admittance parameters of an RC mesh can be determined in a simple manner. The admittance parameter macromodel can be simulated together with the nonlinear portions of the circuit in the time domain using the inverse Laplace transform symbolically, on a term-by-term basis [11].

C. DC Admittance Macromodel

If the substrate is modeled as a purely resistive mesh, the macromodel consists of only the steady-state/DC values of the admittance parameters, the higher-order mesh moments being zero. In contrast to the AWE macromodel where $2q + 1$ matrix inversions (using the ICCG method) have to be performed for every port in the mesh to determine the q th order approximation to the admittance parameters, the DC macromodel requires only one matrix inversion per port for the computation of its admittance parameters. Moreover, transient simulation of the DC macromodel along with the nonlinear portions of the circuit is trivial and requires only the introduction of each admittance parameter into its corresponding location in the global admittance matrix generated by SPICE3 at every time-point in the simulation run.

IV. SIMULATION RESULTS

To validate the macromodeling strategy, simulation results have been compared both to results from the device simulation program MEDICI [12], and also to results from measurements reported on an experimental chip [4].

A. Comparisons to Device Simulation

The device simulation program models two-dimensional distributions of potential and carrier concentrations in a device to predict its electrical characteristics for any bias condition. It solves Poisson's equation and both the electron and hole continuity equations using numerical simulation techniques to analyze devices such as diodes, BJT's, MOSFET's, etc. for dc, steady-state, or transient operating conditions. Since mixed-mode IC's are generally fabricated in processes with either a heavily doped bulk with an epitaxial layer or a lightly doped substrate (without an epitaxial layer), circuits representative of both processes have been verified using the simulation tool. The doping profile of a $2 \mu\text{m}$ BiCMOS technology [5] is used for both simulations with MEDICI and the macromodel.

Fig. 6 shows the experimental setup used to simulate substrate coupling in a heavily doped substrate with a lightly doped epitaxial layer [5]. It consists of a diffused region equivalent to the drain of a switching transistor and a single NMOS transistor current source considered to be part of a sensitive analog circuit separated by a distance of $30 \mu\text{m}$. Several established shielding techniques to reduce the coupling from the switching node to the sensitive transistor,




Fig. 6. Circuit schematic/layout profile for simulations with a heavily doped bulk and a lightly doped epitaxial layer.




Fig. 7. Effect of various shielding techniques on peak-peak noise voltage at the sensitive node in Fig. 6.

including increased separation, an n-well diffusion, a p+ ring and a p+ contact strapping the backside of the substrate to ground potential [4] have been tested with our macromodeling technique. Figs. 7 and 8 compare the results obtained with the macromodeling technique to those obtained with the device simulation program and plots the drain noise voltage (peak-to-peak) and settling time behavior of the sensitive transistor as functions of the shielding technique used. The separation between the switching and sensitive nodes is $30 \mu\text{m}$ unless otherwise specified. The substrate contact on the far right of Fig. 6 is present in all the cases of Fig. 7 and Fig. 8. In case A there is no guard ring or backside contact between the switching and sensitive nodes. Case C shows the effect of simply increasing the separation between the two nodes to $200 \mu\text{m}$. In case D a backside contact is used to additionally bias the substrate. In case B an $8 \mu\text{m}$ wide n-well is placed midway between the switching and sensitive nodes while in case E an additional p+ contact is placed between the nodes.

Fig. 9 is the setup used for simulations with a uniformly lightly doped substrate (without an epitaxial layer) [5] and is otherwise identical to Fig. 6. Fig. 10 shows the peak-peak voltage and Fig. 11 the settling time behavior of the result-




Fig. 8. Effect of various shielding techniques on settling time of the noise voltage at the sensitive node in Fig. 6.

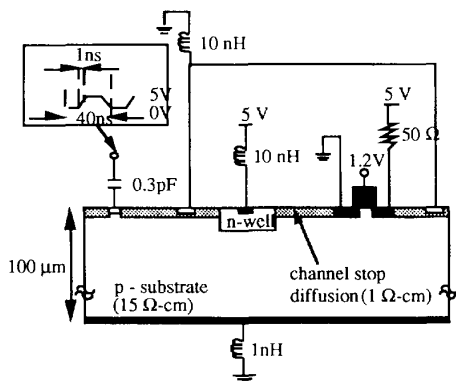


Fig. 9. Circuit schematic/layout profile for simulations with a lightly doped substrate.

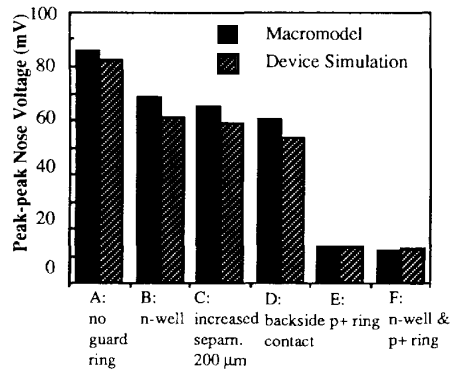


Fig. 10. Effect of various shielding techniques on peak-peak noise voltage at the sensitive node in Fig. 9.

ing noise voltage waveforms using both the macromodeling technique and the device simulation program.

Table I compares the cpu time requirements for the device simulation program and the AWE (second order) and DC macromodeling techniques. With the DC macromodeling technique, the n-well diffusions of Fig. 6 and Fig. 9 were

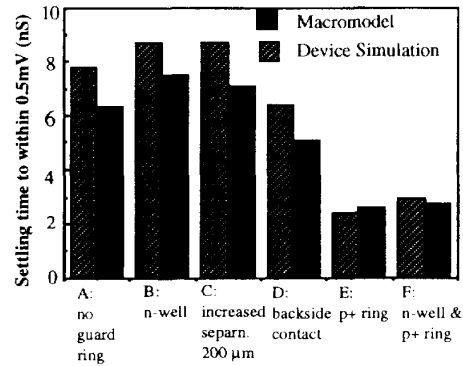


Fig. 11. Effect of various shielding techniques on settling time of the noise voltage at the sensitive node in Fig. 9.

TABLE I
RUN TIME (ON DECSTATION 5000) COMPARISON BETWEEN THE DEVICE SIMULATION PROGRAM AND THE MACROMODELING TECHNIQUES

Number of Mesh Nodes	Device Simulation cpu time (s)	AWE Macromodel cpu time(s)	dc Macromodel cpu time(s)
2940	4375	60.6	5.5
3716	7192.8	92.3	7.2
6605	16 882.2	202.2	20.3
8712	23 732.2	294.3	28.3

extracted as ports and connected to lumped capacitances equivalent to the depletion capacitance of the n-well. Both the AWE and DC macromodeling techniques produced identical results for the examples shown. As exemplified in Table I, the main advantage in using the macromodeling techniques is the significantly shorter cpu times required in their analysis as compared to the device simulation program.

B. Comparisons to Measured Results

The DC macromodeling technique has been used to simulate substrate coupling on the experimental chip reported in [4]. The 2 mm × 2 mm test chip realized in a 2 μm BiCMOS n-well process consists of transistors fabricated in a 15 μm lightly-doped epitaxial layer over a heavily-doped bulk. An on-chip ring oscillator drives a block of 12 CMOS inverters with each inverter output capacitively coupled to the substrate. The switching noise introduced into the substrate is measured by ten single transistor NMOS current sources distributed across the chip. The substrate is biased using a combination of several p+ contacts on the die surface. Seven of the current sources are shielded from the substrate noise using guard rings placed at varying distances (6 μm or 22 μm) from the sources and biased either with a dedicated package pin or connected to two large substrate contacts (one located at the chip center and one diffusion ring surrounding the chip).

The measured results [4] and simulated results are in good agreement as shown in Fig. 12 although only approximate values have been used for bonding pad and chip-to-package capacitance and bond-wire/package pin inductances. Simulations with the macromodel confirm the observation in [4] that is a substrate with a heavily doped bulk, lateral current flow in




Fig. 12. Effect of various guarding configurations on the drain noise voltage of an NMOS transistor on the test chip [4].

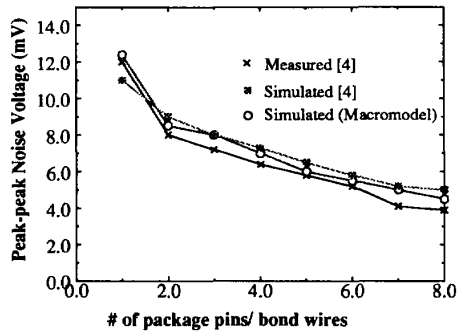


Fig. 13. Effect of multiple substrate bias package pins/bond wires on noise voltage.

the epitaxial layer is negligible as long as the devices and/or contacts are separated by more than 4 times the thickness of the epitaxial layer. A single node substrate model [3], [4] can be justified in these circumstances. When devices/contacts are closer than the critical distance, additional resistances must be introduced into the single node model to represent the resulting lateral current flow. The values of these resistances are easily obtained from the macromodel.

Fig. 13 shows the effect of reducing the power supply inductance on the noise voltage of a sensitive transistor. By increasing the number of package pins used to bias the substrate power supply, noise voltage is reduced dramatically. Fig. 13 compares the measured results [4], simulated results with a single node model [4], and simulated results with the macromodel. The settling time behavior of the noise voltage waveforms are shown in Fig. 22(b).

V. POWER DISTRIBUTION SYNTHESIS STRATEGY

As can be seen from the previous sections, efficient simulation tools allow designers to devise and analyze power distribution techniques to reduce the noise problem for key analog circuits. The next step is to create tools that automate the design of the power distribution network, and that accommodate such design techniques. However, existing power grid layout techniques focus mainly on the geometric problem—assuring connectivity to all power pins [14], [15]—while allowing only rather simple electrical constraints, e.g., pad-to-

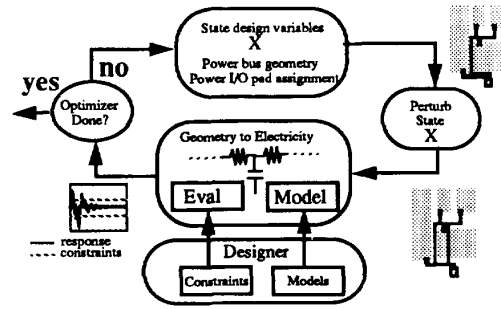


Fig. 14. Optimization-based power synthesis.

pin ohmic drop [17]–[19]. Yet to handle realistic mixed-signal design problems, we must be able simultaneously to optimize the topology of the power grid, the sizing of individual segments, and the choice of I/O pad number and location, under tight dc, ac, and transient electrical constraints arising from the interaction of the power grid with the rest of the IC—notably via substrate coupling effects.

We suggest a power distribution synthesis strategy based on combinatorial optimization techniques that allows us to attack all these concerns simultaneously. We employ an iterative improvement approach illustrated in Fig. 14.

To design a power grid, we begin with an initial (rather arbitrary) “state” for the power bus geometry and power I/O pad configuration. We then perturb this geometry to create a new candidate power grid or pad configuration and update the electrical models for the busses and I/O pads. We next combine those models with designer-supplied circuit macromodels for blocks being supplied by the power grid, and for the substrate. With this complete electrical model—power grid, blocks, pads, substrate—we evaluate the resulting electrical performance, and compare against designer constraints. For example, we might evaluate the coupled noise waveform at a sensitive node against a designer-supplied peak-to-peak noise amplitude constraint. Finally, the optimizer accepts or rejects the perturbation based on the result. We continue the iterative improvement loop until the optimizer determines no further improvement is possible.

Our optimization-based strategy comprises four major components which we describe in the following four sections. Section VI describes the geometric representation of the problem. Section VII describes the optimization formulation. Section VIII discusses necessary electrical modeling techniques. Finally, Section IX discusses efficient electrical evaluation of the performance of the evolving power grid.

VI. GEOMETRIC REPRESENTATION

The highest level physical design decision that affects the geometric design of the power grid is the overall design style selected for the IC. As with previous methods, we focus on the custom 2-D macrocell design style [14], [17], [18], which can appear in either a flat or a slicing representation. We select the slicing representation to avoid the channel identification preprocessing step associated with flat representations.

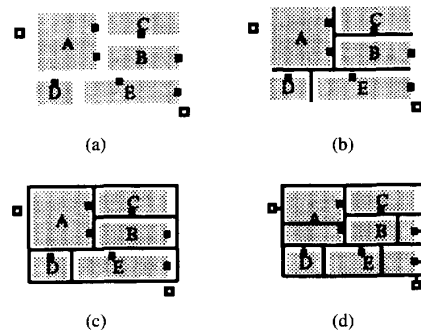


Fig. 15. Creating the general grid. (a) Slicing tree placement. (b) Power bus segments for slicing channels. (c) Power bus segments for I/O channels. (d) Power bus segments for macros (includes over-the-cell).

A. Power Busses

To represent power bus geometry, we offer a new formulation which we term a *general grid*. Unlike previous synthesis methods which made at most one single bus available for each power net in *many* channels [14]–[17], we allow a single bus, at most, for each power net in *every* channel and *every* designer-supplied over-the-cell power feedthrough. This formulation was inspired by work done for signal nets in gate-arrays [20] and was independently rediscovered for standard cell power grids in [19]. The idea is essentially subtractive: we allocate power grid segments everywhere, and formulate as an optimization problem which to resize, and which to remove. The idea is illustrated in Fig. 15.

Starting with a slicing tree macrocell placement in the layout [Fig. 15(a)] we generate power bus nets in all available channels [Fig. 15(b)]. Then, to connect to power I/O pads, we build power bus segments in the four I/O channels that surround the macrocells [Fig. 15(c)]. Finally, we introduce power bus segments perpendicular to these channel segments to handle over-the-cell feedthroughs and to connect to macros and pads [Fig. 15(d)]. For maximum flexibility, we provide power bus geometry for *all* power nets in *each* feedthrough, but later ensure that at most one net uses the resource.

This collection of segments constitutes the *general grid*, the master topology of which all final topologies are a subset. Each individual grid segment's width is an independent variable. These segments may be sized to one of several discrete widths ranging from the minimum width allowable in the technology to a designer-specified maximum. In addition to sizing the master topology, we add *zero* width to the sizing options to remove segments for topological selection. Now, both sizing and topological optimization can proceed in the guise of sizing a master topology. We avoid singularities in the associated electrical model by modeling the zero width segment (ideally an electrical "open") as a minimum conductance. Similarly, a maximum conductance is used to model a "short," as we shall see in the next subsection. This formulation has several advantages: it can handle over-the-cell power routing at designer discretion; it can discover power layouts in both tree and grid topologies; and it can partition quiet and noisy pins onto separate nets/pads.

B. Reconfigurable I/O Blocks

In prior work, the power busses have been the major focus of attention and optimization methods have synthesized *either* power bus topologies [16], [19] *or* power bus sizings [17], [18]. Our general grid formulation allows us to handle these concerns simultaneously. However, Fig. 13 clearly demonstrated the importance of proper power I/O pad assignment as another component of this problem. We can easily extend the idea of discretely-selected power bus segment widths to handle power I/O pad assignment. Each I/O pad is connected to each power bus via a resistor. We refer to this as a *reconfigurable* I/O block. By discretely selecting between on/off switch values for resistors inside each I/O pad, we can effectively reassign the I/O to the appropriate power net, or even no net.

VII. OPTIMIZATION

Most previous methods assume a fixed power I/O pad assignment, then determine power bus topologies, then determine the sizing. In contrast, we reformulate the problem to allow pad assignment, topology selection, and segment sizing to be optimized simultaneously. We argue that all these decisions must be handled simultaneously to make optimal tradeoffs under tight constraints. To attack this combined problem we employ simulated annealing [24], a general optimization strategy based on iterative improvement with controlled hill climbing.

Hill climbing allows annealing strategies to avoid local minima in a complex cost surface, and reach better global solutions. To characterize our power distribution synthesis strategy, we need to describe the four components of any annealing-based optimizer: 1) the *representations* for intermediate states of the power distribution visited during iterative improvement; 2) the *moves* that transform one intermediate state of the layout to the next; 3) the *cost function* used to evaluate the quality of each intermediate power distribution; and 4) the *cooling schedule* used to control hill climbing.

A. Representation

We represent power bus segments with the *general grid*, updating their widths as the power distribution evolves. Similarly, we represent reconfigurable I/O blocks with their user-defined switch configurations and associated geometries.

B. Move Set

Controlled random perturbations, or *moves*, in conjunction with the cost function, are responsible for generating a "good" final power distribution. Our move set supports two classes of perturbations: changes to paths of connected segments and changes to interdependent sets of reconfigurable I/O blocks.

For on-chip power nets, moves alter segment paths between macrocells and pads (see Fig. 16). We first compute dc currents in each power bus segment. These currents assign a direction to each segment, used for path finding between macros and pads. A move begins by selecting a random segment in a random power net. It then traverses a random number of segments, going either toward a pad or macro. The resulting




Fig. 16. Power bus moveset. (a) Segments with current directions. (b) Resize a path. (c) Reroute a path. (d) Swap two paths.

set of connected segments forms our starting point. The move then randomly selects from among three possible changes to make. A *resizing* move alters the width of segments on this path [Fig. 16(b)]. A *rerouting* move selects a second random path through zero-width segments, removes the first path, and distributes its area over the zero-width segments. [Fig. 16(c)]. A *swapping* move selects a second random path of similar area, and redistributes the segment widths of the first path over the second, and vice versa [Fig. 16(d)]. Note that the rerouting and swapping moves can alter the connectivity of the grid. Finally, we adaptively weight our move selection to maximize the probability of selecting successful moves [26].

For reconfigurable I/O blocks, the key is to maintain specified dependencies among blocks during perturbations. A move randomly selects an independent I/O block, connects it to either a power bus or the substrate, then randomly selects compatible connections for all dependent blocks. For example, a move which connects the package pin to VDD will also force its associated chip pad to connect to the chip VDD bus.

C. Cost Function

Our success in generating a “good” final power distribution is measured by a weighted-sum cost function, updated incrementally after each move:

$$Cost = w_1 Area + w_2 dc + w_3 ac + w_4 Tran \quad (8)$$

The weights are determined empirically. Our goal is to minimize power bus area consumed (the *Area* term) while satisfying interacting electrical constraints (the *dc*, *ac*, and *Tran* terms). The annealer searches among the power bus and pad configurations and selects one minimizing this function.

The *Area* term sums the area consumed by each power bus within its assigned routing area. Each channel is partitioned into signal and power net routing areas. Power bus area consists not only of the physical area of each segment, but the spacing halo around each segment dictated by design rules. Minimizing an accurately estimated *Area* term ensures we reserve maximal space for subsequent signal routing.

The remaining terms are all penalty terms, i.e., they increase with violations from a hard electrical constraint, but contribute

zero to the cost when the constraint is met. Each constraint is normalized and includes an absolute tolerance, *ATOL*, obtained as input, to handle zero specifications. All constraint costs are given in the following form:

$$Constraint\ cost_i = \max\left(\frac{performance_i - spec_i}{spec_i + ATOL}, 0\right) \quad (9)$$

The *dc* term penalizes violations of dc constraints. For electro-migration, we constrain power bus maximum current density and power pad maximum current. For maintaining local power supply levels at macros, we constrain voltage drops at the macro between power nets. To limit voltage offsets and indirectly control power dissipation and circuit bias, we constrain ground shift between macros and pads. To directly constrain chip power dissipation and macrocell bias currents, we constrain macrocell currents.

The *ac* term penalizes violations of ac constraints. Particularly, we focus on chip resonance frequencies and design power distribution to ensure they are well above clock frequencies present on chip. We also examine the driving point impedance at some key digital switching blocks to gauge the amplitude of noise generated.

The *Tran* term penalizes violations in the transient constraints. Designers typically characterize noise effects in terms of peak-to-peak noise seen at a sensitive analog node [3], [4]. Thus, we also constrain peak-to-peak noise amplitude at these nodes. Since the noise problem originates in the digital portion of the chip, we also can constrain that peak-to-peak noise amplitude.

We should also note here the interaction between the cost function, the move set, and the general grid representation. Because segment widths and pad assignments are discrete, the cost surface is highly discontinuous. Moreover, we permit infeasible intermediate configurations of the power grid, i.e., segments can be sized to zero width, disconnecting parts of the grid. Such a formulation poses severe problems for traditional gradient optimizers, but poses no problems for an annealer. Indeed, the advantage of this strategy is that it simplifies the design of the move set, permits simultaneous optimization of pad assignment, grid topology and segment sizes, and allows large-scale perturbations of the entire power distribution configuration which lead to more aggressive search among tradeoffs.

D. Cooling Schedule

The *cooling schedule* controls the hill climbing during annealing. The power distribution is “heated” to a high enough temperature to allow any perturbation in power distribution to be accepted. As the power distribution is carefully “cooled,” pad configurations solidify, power bus topologies crystallize, and power bus segment sizes fine-tune to generate the final power distribution. We employ an efficient, automatic cooling schedule [25].

VIII. ELECTRICAL MODELING

Good design practice provides near ideal power supplies on the card so we need only consider the electrical aspects of




Fig. 17. Substrate coupled macros and busses on-chip. (a) Coarse substrate mesh couples. (b) Macros connected to it via arbitrary circuit models.

the chip and package. We model the principal components of chip and package power distribution. This includes not only models for the chip macrocells and power busses, but models for the chip substrate and chip-to-package interconnect. With each annealing perturbation we map the attempted geometric changes to their electrical equivalents. In contrast to the simple resistor and dc current source modeling of previous methods [17]–[19], we support arbitrary independent and controlled sources, frequency and time-varying sources, and complete RLC models for all components of the power distribution network. We provide parameterizable default models, but to handle the expected variety of realistic models, we also support input of designer-specified models.

There are four modeling problems: 1) chip macrocells, 2) power bus segments, 3) chip-package interconnect, and 4) chip substrate. For macrocells we accept user input of a linear model arbitrarily interconnected between macro terminal pins and any additional substrate pins defined in the macro. For each power bus segment, we generate a parasitic RC π model. For chip-package interconnect, we estimate wire bond inductance. For the substrate, we either accept a user-supplied model, or automatically generate from technology data a *coarse* resistive mesh model with the designer-specified number of grids. For each substrate pin in macrocells and net segments, we assume connection to the nearest substrate mesh node. Fig. 17 illustrates the resulting model.

IX. ELECTRICAL EVALUATION

For optimization problems of this scale, an annealer will typically visit 10^4 – 10^5 design configurations. Since each move requires us to reevaluate the cost function, we require efficient techniques for the electrical evaluation of our power distribution models. Because we allow arbitrary power grid topologies, and complex macro, substrate, and chip-package models, the simple analytical formulations of early approaches [17], [18] cannot capture the behavior of the resulting networks. To deal




Fig. 18. Simultaneous switching of multiple sources. (a) Weighted impulses generate (b) weighted pulses.




Fig. 19. Mixed1 sample/hold, placement and power distribution.

with realistic design constraints, we must handle not only dc analysis (see, e.g., [19]), but also ac and transient analyses.

A. Evaluating dc Performance

To evaluate circuit behavior we use modified-nodal analysis. This method is general, and allows us to analyze the dc behavior of the arbitrary linear network resulting from each annealing move. It is also efficient, requiring 1 LU factorization and forward-and-back substitution. Further, matrix reformulation is never required and matrix reordering is rarely needed since each new topology visited by the annealer is a subset of the master topology created by the general grid and reconfigurable block formulations. The computed dc currents and voltages are used for cost function evaluation; the currents are saved for each power bus segment to provide the directions needed to trace paths in the next move.

B. Evaluating ac and Transient Performance

We again employ AWE [9] for efficient ac and transient evaluation. For this synthesis application, AWE is typically 100–1000 \times faster than SPICE. An ac response to an input is available directly from AWE's computed pole-zero approximation. A transient response for a typical input waveform comprised of steps and ramps is also easily generated. An interesting aspect of our use of AWE is in how we model simultaneous switching. Most noise analysis tools based on CMOS switching sources use trapezoidal, triangular, or sawtooth waveforms to model current switching [21]–[23], all of which can be described with a superposition of steps and ramps. The transient response from multiple input sources of varying strengths switching simultaneously can be efficiently estimated with one AWE evaluation. Using the linearity of the system and the constant amplitude scaling for ramps and steps, the multiple source response is obtained by weighting

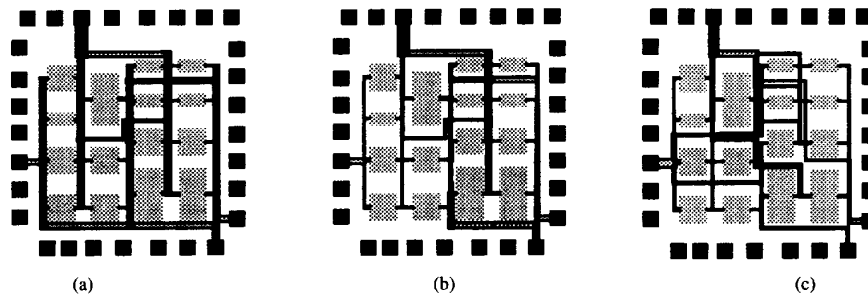


Fig. 20. Analog1 benchmark layouts for VDD and GND nets. (a) Manual layout. (b) Topology fixed and sized. (c) Topology selected and sized.




Fig. 21. Mixed3 benchmark layouts and simulated noise comparisons. (a) Layout for dc constraints. (b) Layout for dc and transient constraints. (c) SPICE GND noise comparison between layouts for macro 5 (highlighted).

TABLE II
SYNTHESIZED POWDER BUS AREA AND RUN
TIMES FOR BENCH MARK PROBLEM SIZES

Example	Power Bus Segments/Pin Blocks Configured	Circuit Nodes/Substrate Grid Size	Power Bus Area [μm^2]	CPU Time IBM RS6000/550 [hours]
Mixed 1	252/-	588/2 × 2 × 1	2.39E5	1.7
Analog 1(a)	142/-	—	4.74E5	manual
Analog 1(b)	142/-	330/6 × 6 × 1	3.82E5	2.2
Analog 1(c)	142/-	330/6 × 6 × 1	3.25E5	3.4
Mixed 3(a)	294/-	618/10 × 10 × 1	7.83E5	5.3
Mixed 3(b)	294/-	618/10 × 10 × 1	9.06E5	14.6
Noise test	132/9	560/2 × 2 × 1	6.6E4	1.3

the impulses in the original moment generation phase of the AWE algorithm as shown in Fig. 18. In other words, with one superposition of weighted impulses, AWE approximates the effect of multiple, independently switching sources.

X. SYNTHESIS RESULTS

We have implemented our ideas in a tool called RAIL. RAIL has generated power distribution for several analog and mixed-signal examples [27]. We have included three mixed-signal and one analog example in this section. The power bus area, circuit size, and CPU time for these power distribution synthesis examples are shown in Table II.

The first example in Fig. 19 is a sample and hold circuit with on-chip clock generation. Three sets of power nets (one analog, two digital) were synthesized by RAIL. The most interesting feature is the automatic tapering in the analog ground bus and the minimizing of the analog VDD bus since much of the current from the VDD rail flows through signal pads connected to NMOS drains.

The next example in Fig. 20 is an industrial analog bipolar IC. This example demonstrates the importance of simultaneous topology selection and sizing. Fig. 20(a) shows the conservative manual design for the power busses. Fig. 20(b) demonstrates the benefit from automatically sizing this given (manual) topology. Finally, Fig. 20(c) shows the further benefit of simultaneous power bus topology selection and sizing.

The third example is a larger mixed-signal chip which illustrates the importance of ac and transient noise constraints during power distribution synthesis. Fig. 21(a) shows the synthesized layout when ac and transient constraints are ignored. Short routes which meet dc electromigration and voltage constraints are created. However the noise is intolerable at 600 mV as shown in Fig. 21(c). After applying transient noise constraints on sensitive analog power net connections, the layout in Fig. 21(b) is generated and the noise is reduced to 75 mV.

Our last example demonstrates the accuracy obtainable through linear macromodeling, and the importance of simultaneous optimization of power busses and power I/O pad assignment. We again use the test chip of [4], for which detailed noise measurements have been published. From the




Fig. 22. Results for 68 pin PLCC package with ceramic decoupling capacitor and noise test chip [4]. (a) PLCC cavity with chip, decoupling capacitor, and wire bonds to lead frame. (b) Settling time at drain for number of wire bonds. (c) Noise voltage (peak-to-peak) at drain for number of wire bonds.

description of the chip-package environment, Fig. 22(a) shows our rendering of the chip-package physical design. The 68 pin PLCC cavity contains the 2×2 mm chip and a ceramic decoupling capacitor. First, we created linear macromodels of the analog current sources (noise receivers) and the oscillator-based digital switching logic (noise generators). Next, we synthesized power distribution with the constraint that the noise coupled at the drain of one current source must not exceed 4 mV. Further, only one substrate I/O pad was allowed. As expected, meeting both the dc constraints and the transient 4 mV noise constraint was unachievable. However, when RAIL was allowed to perform power pad I/O assignment simultaneously, it synthesized the power busses to meet dc constraints, and added five substrate contact pins to meet the transient noise constraint, thus completing the power distribution.

Fig. 22(b) and (c) show measured and simulated settling times and noise voltages for this experiment, respectively. We believe this result justifies the use of coarse mesh substrate approximations in heavily doped substrates.

XI. CONCLUSIONS

We have described simulation and power distribution synthesis techniques to address the parasitic coupling of noise through the substrate in mixed-signal integrated circuits. The simulation tool uses a modified version of SPICE3 which spatially discretizes the substrate from the layout information, develops its macromodel and performs circuit simulation with the macromodel. Simulation results indicate a significant reduction in cpu time while preserving the accuracy of device simulations. Results also compare favorably with measurements made on a test chip. The synthesis tool, RAIL, creates analog power distribution while minimizing noise coupling to sensitive analog circuits. RAIL simultaneously optimizes power bus topology, segment sizing, and power I/O pad assignment. Its circuit evaluation efficiency is derived from a synergy between our geometric formulations and AWE. Power distribution synthesis results demonstrate the critical need to

model the substrate and constrain the noisy ac and transient behavior it introduces. The simulator can be valuable as a verification tool for analog and mixed-mode IC designs and the complementary synthesis tool can be valuable as a power distribution floorplanner for these same designs.

ACKNOWLEDGMENT

The authors thank T. J. Schmerbeck of International Business Machines, Dr. S. Kumashiro of NEC, Dr. V. Raghavan, J. E. Bracken of Performance Signal Integrity, and S. Masui of Nippon Steel Corporation for several useful discussions, and T. Mukherjee of Carnegie Mellon University for first using RAIL. M. J. Loinaz, T. Blalack, D. K. Su, and Prof. B. A. Wooley of Stanford University provided useful information on their test chip.

REFERENCES

- [1] J. A. Olmstead and S. Vulih, "Noise problems in mixed analog/digital integrated circuits," in *Proc. Custom Integrated Circuits Conf.*, 1987, pp. 659-662.
- [2] B. M. J. Kup, E. C. Dijkmans, P. J. A. Naus, and J. Sneep, "A bit stream digital-to-analog converter with 18-b resolution," *IEEE J. Solid State Circuits*, vol. 26, no. 12, pp. 1757-1763, Dec. 1991.
- [3] T. Schmerbeck, R. Richetta, and L. Smith, "A 27 MHz mixed analog/digital magnetic recording channel DSP using partial response signalling with maximum likelihood detection," in *Tech. Dig. IEEE Int. Solid State Circuits Conf.*, Feb. 1991, pp. 136-137.
- [4] D. K. Su, M. J. Loinaz, S. Masui, and B. A. Wooley, "Experimental results and modeling techniques for substrate noise in mixed-signal integrated circuits," *IEEE J. Solid State Circuits*, vol. 28, no. 4, pp. 420-430, Apr. 1993.
- [5] S. Masui, "Simulation of substrate coupling in mixed-signal MOS circuits," in *Tech. Dig. Symp. VLSI Circuits*, 1992, pp. 42-43.
- [6] N. K. Verghese, D. J. Allstot, and S. Masui, "Rapid simulation of substrate coupling effects in mixed-mode ICs," in *Proc. IEEE Custom Integrated Circuits Conf.*, May 1993, pp. 18.3.1-18.3.4.
- [7] S. Kumashiro, "Transient simulation of passive and active VLSI devices using asymptotic waveform evaluation," Ph.D. dissertation, Dep. Elec. Comput. Eng., Carnegie Mellon Univ., Aug. 1992.
- [8] T. L. Quarles, "The SPICE3 implementation guide," Univ. California, Berkeley, Tech. Rep. ERL-M89/44, Apr. 1989.
- [9] L. T. Pillage and R. A. Rohrer, "Asymptotic waveform evaluation for timing analysis," *IEEE Trans. Comput.-Aided Design Integrated Circuits Syst.*, vol. 9, no. 4, pp. 352-366, Apr. 1990.

- [10] X. Huang, V. Raghavan, and R. A. Rohrer, "AWESim: A program for the efficient analysis of linear(ized) circuits," in *Proc. IEEE Int. Conf. Computer Aided Design*, 1990, pp. 534-537.
- [11] V. Raghavan, J. E. Bracken, and R. A. Rohrer, "AWESpice: A general tool for the accurate and efficient simulation of interconnect problems," in *Proc. 29th Design Automation Conf.*, 1992, pp. 87-92.
- [12] *TMA MEDICI: Two Dimensional Device Simulation Program*, Technology Modeling Associates, Inc., Vols. 1, 2, Mar. 1992.
- [13] J. A. Meijerink and H. A. Van der Vorst, "An iterative solution method for linear systems of which the coefficient matrix is a symmetric M-matrix," *Math. Computat.*, vol. 31, pp. 148-162, Jan. 1977.
- [14] H.-J. Rothermel and D. A. Mlynski, "Computation of power supply nets in VLSI layout," in *Proc. 18th Design Automation Conf.*, June 1981, pp. 37-42.
- [15] Z. A. Syed, A. El Gamal, and M. A. Breuer, "On routing for custom integrated circuits," in *Proc. 19th Design Automation Conf.*, June 1982, pp. 887-893.
- [16] A. S. Moulton, "Laying the power and ground wires on a VLSI chip," in *Proc. 20th Design Automation Conf.*, June 1983, pp. 745-755.
- [17] S. Chowdhury, "An automated design of minimum-area IC power/ground nets," in *Proc. 24th Design Automation Conf.*, June 1987, pp. 223-229.
- [18] R. Kolla, "A dynamic programming approach to the power supply net sizing problem," in *Proc. European Design Automation Conf.*, Mar. 1990, pp. 600-604.
- [19] T. Mitsuhashi and E. Kuh, "Power and ground network topology optimization for cell based VLSIs," in *Proc. 29th Design Automation Conf.*, June 1992, pp. 524-529.
- [20] A. M. Patel, N. L. Soong, and R. K. Korn, "Hierarchical VLSI routing—an approximate routing procedure," *IEEE Trans. Comput.-Aided Design*, vol. CAD-4, no. 2, Apr. 1985.
- [21] R. Burch, F. Najm, P. Yang, and D. Hocevar, "Pattern-independent current estimation for reliability analysis of CMOS circuits," in *Proc. 25th Design Automation Conf.*, June 1988, pp. 294-299.
- [22] S. Chowdhury and J. Barkatullah, "Current estimation in MOS IC logic circuits," in *Proc. IEEE Int. Conf. Computer-Aided Design*, Nov. 1988, pp. 212-215.
- [23] J. Hall, D. Hocevar, *et al.*, "SPIDER-A CAD system for modeling VLSI metallization patterns," *IEEE Trans. Comput.-Aided Design*, vol. CAD-6, no. 6, pp. 1023-1031, Nov. 1987.
- [24] S. Kirkpatrick, D. Gelatt, and M. Vecchi, "Optimization by simulated annealing," *Science*, vol. 220, 1983.
- [25] M. Huang, F. Romeo, and A. Sangiovani-Vincentelli, "An efficient general cooling schedule for simulated annealing," in *Proc. Int. Conf. Computer-Aided Design*, Nov. 1986, pp. 381-384.
- [26] S. Hustin and A. Sangiovani-Vincentelli, "TIM, a new standard cell placement program based on simulated annealing algorithm," presented at the IEEE Physical Design Workshop Placement and Floorplanning, Hilton Head, SC, Apr. 1987.
- [27] B. R. Stanisic, R. A. Rutenbar, and L. R. Carley, "Power distribution synthesis for analog and mixed-signal ASICs in RAIL," in *Proc. IEEE Custom Integrated Circuits Conf.*, May 1993, pp. 17.4.1-17.4.5.
- [28] Z. J. Cendes and D. N. Shenton, "Adaptive mesh refinement in the finite element computation of magnetic fields," *IEEE Trans. Magnetics*, vol. 21, pp. 1221-1228, 1985.



Balsha R. Stanisic (M'84) was born in Milwaukee, WI, in 1962. He received the B.S. and M.S. degrees in electrical and computer engineering from the University of Wisconsin at Madison in 1984 and 1987, respectively, and the Ph.D. degree from Carnegie Mellon University in 1993.

From 1984 to present he has worked at IBM's AS/400 division, Rochester, MN, in the area of analog circuit CAE and design. His research interests include VLSI layout algorithms and synthesis techniques. His current research is in the area of analog power distribution synthesis. He is the author of RAIL.



Nishath K. Verghese (S'91) received the B.Eng. (Hons.) degree from the Birla Institute of Technology and Science, Pilani, India, in 1990 and the M.S. degree from Carnegie Mellon University in 1993.

From 1990 to 1991 he was with the VLSI Design Laboratory, McGill University, where he worked on new circuit implementations of sigma-delta A/D converters and on estimation techniques for power bus current in CMOS logic circuits. He is currently working toward the Ph.D. degree at Carnegie Mellon University where his research focuses on the simulation of substrate coupling in mixed-mode IC's. His research interests include design and verification of analog and mixed-mode circuits.



Rob A. Rutenbar (S'78-M'84-SM'90) received the Ph.D. degree in computer engineering (CICE) from the University of Michigan, Ann Arbor, in 1984.

He subsequently joined the faculty of Carnegie Mellon University, Pittsburgh, PA. He is currently Professor of Electrical and Computer Engineering, and of Computer Science, and Director of the SRC-CMU CAD Center. His research interests focus on circuit and layout synthesis for mixed-signal ASIC's.

In 1987, Dr. Rutenbar received a Presidential Young Investigator Award from the National Science Foundation. At the 1987 IEEE-ACM Design Automation Conference, he received a Best Paper Award for work on analog circuit synthesis. He is currently on the Executive Committee of the IEEE International Conference on CAD, Program Committees for the ACM/IEEE Design Automation Conference and European Design and Test Conference, and a member of the Editorial Board of IEEE SPECTRUM. He is on the Strategic Advisory Council for Cadence Design Systems. He is a member of ACM, Eta Kappa Nu, Sigma Xi, and AAAS.



L. Richard Carley (S'77-M'84-SM'90) received the S.B., M.S., and Ph.D. degrees from the Massachusetts Institute of Technology, Cambridge, in 1976, 1978, and 1984, respectively.

He is a Professor of Electrical and Computer Engineering at Carnegie Mellon University. He has worked for MIT's Lincoln Laboratories and has acted as a consultant in the area of analog circuit design and design automation for Analog Devices and Hughes Aircraft among others. In 1984 he joined Carnegie Mellon, and in 1992 he was promoted to Full Professor. His current research interests include the development of CAD tools to support analog circuit design, the design of high performance analog signal processing IC's, and the design of low-power high-speed magnetic recording channels.

Dr. Carley was awarded the Guillemin Prize for the best EE Undergraduate thesis in 1976, a National Science Foundation Presidential Young Investigator Award in 1985, a Best Technical Paper Award at the 1987 Design Automation Conference, and a Distinguished Paper Mention at the 1991 International Conference on Computer-Aided Design.



David J. Allstot (S'72-M'78-SM'83-F'92) received the B.S.E.S. degree from the University of Portland, Portland, OR, the M.S.E.E. degree from Oregon State University, Corvallis, OR, and the Ph.D. degree in electrical engineering and computer science from the University of California, Berkeley, in 1979.

He has held industrial positions with Tektronix, Texas Instruments, and MOSTEK, and academic positions with the University of California, Berkeley, Southern Methodist University, Dallas, TX, and Oregon State University. He is currently Professor of Electrical and Computer Engineering at Carnegie-Mellon University Pittsburgh, PA, and Associate Director of the SRC-CMU Research Center for Computer-Aided Design.

Dr. Allstot's professional service positions have included: Associate Editor of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS, Guest Editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS, Editorial Board Member of the *Analog Integrated Circuits and Signal Processing Journal*, Technical Program Committee member of the IEEE Custom IC Conference and the IEEE International Symposium on Circuits and Systems, and elected member of the IEEE Circuits and Systems Society Board of Governors. He is currently Editor of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS II: ANALOG AND DIGITAL PROCESSING. He was a corecipient of the 1980 IEEE W. R. G. Baker Award, and has received excellence in teaching awards from SMU, OSU and CMU. He is a member of Eta Kappa Nu and Sigma Xi.