Substrate-Aware Mixed-Signal Macrocell Placement in WRIGHT

Sujoy Mitra, *Student Member, IEEE,* Rob A. Rutenbar, *Senior Member, IEEE,*
L. Richard Carley, *Senior Member, IEEE,* and David J. Allstot, *Fellow, IEEE*

Abstract— **We describe a set of placement algorithms for handling substrate-coupled switching noise. A typical mixed-signal IC has both sensitive analog and noisy digital circuits, and the common substrate parasitically couples digital switching transients into the sensitive analog regions of the chip. To preserve the integrity of sensitive analog signals, it is thus necessary to electrically isolate the analog and digital. We argue that optimal area utilization requires such isolation be designed into the system during first-cut chip-level placement. We present algorithms that incorporate commonly used isolation techniques within an automatic placement framework. Our substrate-noise evaluation mechanism uses a simplified substrate model and simple electrical representations for the noisy digital macrocells. The digital/analog interactions determined through these models are incorporated into a simulated annealing macrocell placement framework. Automatic placement results indicate these substrate-aware algorithms allow efficient mixed-signal placement optimization.**

I. Introduction

THE ability to implement mixed-signal systems by integrating both digital and analog portions on the same silicon substrate [1], [2] has paved the way for a new generation of reliable, cost-effective single-chip solutions. However, this integration has also introduced a host of new problems that were previously solved by simply isolating the analog and digital circuity in separate packages. Most of these problems are a result of parasitic interactions—in particular, substrate interactions—between the digital and analog subsystems [3]. CAD tools for mixed-signal system design must model and manage such parasitic interactions to be of practical use to the design community. Recent efforts here have included new noise models [4], the first practical, large-scale substrate simulation techniques [5]–[7], and physical design techniques [5]–[11]. Our focus is on physical design, in particular chip-level macrocell placement. We argue that gross coupling problems are best addressed as early as possible in the physical design of a mixed-signal IC, i.e., during actual placement of noisy and sensitive blocks. Given a good placement, downstream layout/verification tools such as noise-sensitive routers [10], [11], power grid synthesis [7]–[9], and full-chip substrate simulators [5], [7] can be used to maximum advantage.

Manuscript received August 1, 1994; revised October 20, 1994. This work was supported by the Semiconductor Research Corporation under Contract DC-068-94.

The authors are with the Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213 USA.

IEEE Log Number 9408751.

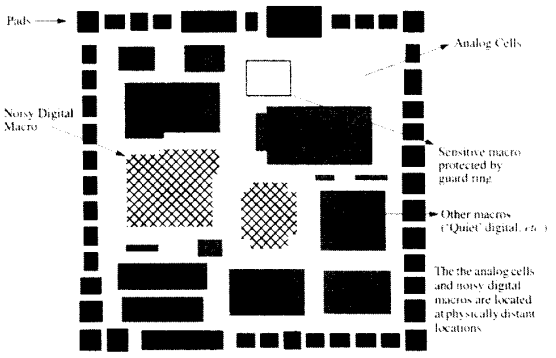

Fig. 1. Typical isolation mechanisms in mixed-signal systems.

This paper describes algorithms for automatic placement that are specifically targeted toward mitigating the effects of noise coupling through the substrate. The key innovation is the inclusion of efficient models for the chip substrate, noise sources and receivers on the macrocells, and mitigation measures such as guard rings *inside* the placement algorithm. We refer to this ability to make placement choices in light of their consequences for substrate noise coupling as *substrate-aware* placement.

II. Background

Prior work in this area has focused on chip-level layout methodologies, modeling and simulation for substrate coupling, and physical design algorithms.

Most methodologies for handling substrate coupling advocate the physical separation [12] of analog and digital circuitry in clearly delineated parts of the substrate. This is adequate in most cases, but some especially sensitive macros require additional isolation. This is achieved by inserting *guard rings* around these blocks (Fig. 1). These guard rings are usually low-impedance ties to a quiet potential and hence they have the effect of creating (ideally) a zero potential ring around the sensitive macro, thereby electrically isolating it from the rest of the chip. Unfortunately, most existing design methodologies remain *ad hoc*, and make no attempt either to quantify the level of interaction or formalize a clear methodology for separation or isolation. Often enough, it is possible for the experienced designer to manually manage the digital–analog interaction. But the growing complexity of these chips is making this task progressively more difficult and time consuming.

Modeling and simulation techniques for predicting substrate noise have been addressed in [4]–[6] and [13]. In [4], Su *et al.* address the problem of modeling the substrate. The active devices (transistors) are resistively coupled to a substrate represented by a single node. This model is adequate for epitaxial processes but rather inaccurate for the more common bulk processes. (Epitaxial processes compare poorly against bulk processes from the economic standpoint.) In the bulk case, the substrate is represented as a mesh or grid of appropriate linear elements, as in [5]. Noise couples into and out of the substrate through well diffusions and contacts modeled through appropriate linear approximations. The advantage of the single-node model is its computational simplicity for simulation; the advantage of the more general finite element-style substrate mesh is its generality and greater accuracy. The first practical techniques to simulate noise with complex substrate meshes appeared in [5], which introduced the idea of using Asymptotic Waveform Evaluation techniques (AWE [14]) to reduce the large linear substrate model to a small, tractable modal approximation.

CAD tools for physical design have traditionally ignored the substrate. Only recently has this begun to change. For example, the power distribution synthesis tool, RAIL [7], [8], considers noise coupling into power busses through the substrate. However, to our knowledge, the effects of substrate coupling have never been considered quantitatively in the context of automatic placement.

III. Strategy

We outline here the key ideas behind our basic placement strategy.

- *Macrocell Formulation:* We target placement problems composed of rectilinear analog and digital macrocells. We assume that all individual analog cells are placeable, and that many digital cells will already be grouped into larger functional blocks (e.g., via row-based placement inside the blocks). We are mostly concerned with placement problems as they appear during chip-level floorplanning when major functional blocks have already been laid out.
- *Simulated Annealing Optimization:* We adopt the basic iterative-improvement placement formulation of [15], and extensively refined in [16] and [17]. Cells are iteratively and arbitrarily relocated across the chip, without concern about whether they overlap or not. A cost function is evaluated after each placement perturbation; the goal of placement is to minimize this function. The cost function has two types of terms: *objective* terms which should be minimized (e.g., area and wirelength) and *penalty* terms which must be driven to zero to obtain a legal placement. Hard constraints are mapped into these penalty terms. For example, illegal overlaps among modules are evaluated and transformed into one such penalty term. In our substrate-aware placer, the key innovation is the formulation of penalty terms for hard electrical constraints, e.g., substrate-coupled switching noise above acceptable limits at any designer-specified location. This is the same overall strategy used successfully in [7] and [8].
- *Simplified Substrate Modeling:* We update a model of the substrate noise profile after *each* annealing placement perturbation. Since an annealer must visit many placement configurations (10^5 to 10^6) we need extremely fast, though necessarily approximate noise evaluation. The appropriate analogy here is to wirelength estimators used in the inner loop of a placer: we can afford simple estimators such as bounding-box length or simple spanning tree estimators [18], but we cannot afford to invoke a real router. Similarly, we can afford a simplified noise estimator, but not a full-chip simulator such as [5], [7], [19]. We rely on a coarse, lumped substrate grid, preprocessing of the grid to a smaller n-port model, and adaptive thresholding of the n-port matrix to render it sparse (and thus faster to solve) during the early phases of the placement.
- *Simplified Noise Injection and Noise Receiver Modeling:* Just as with the substrate model, we need only an extremely coarse model for substrate noise injectors and sensitive analog circuits. Because substrate coupling occurs through both capacitors and resistors, it is frequency dependent. However, in order to simplify the analysis which must occur in the inner loop of an annealer, we assume that there is a single user selected frequency (e.g., the clock frequency) that will provide a good first-order estimation for signal coupling. Further, all of the resistive and capacitive couplings to the substrate within a module are collected together into one or a few geometrically localized patches (the number would be determined by the cell size compared with the substrate grid size). Then, a magnitude must be assigned to the sinusoidal noise current injected across the admittance at each of these patches at the fundamental frequency selected. For example, this magnitude can be derived from circuit simulations of functional blocks with representative or worst-case input waveforms, or from the power estimates of the individual macrocells [20]. Finally, the designer must provide constraints on the magnitude of substrate voltage variations and on the gradients in substrate voltage that will ensure correct operation of sensitive analog cells.

Our basic goal is first-cut cell placement sensitive to significant substrate effects. We focus on an early phase of the layout, prior to signal and power grid routing, and thus prior to the availability of full-chip geometry for full-chip substrate simulation. The accuracy available in these tools is unaffordable within an iterative placement framework that must visit thousands of candidate placement configurations. Our central argument is that this accuracy is also unnecessary: simpler models can produce reasonable results more reliably than trial-and-error manual techniques. These automatic layouts can adequately serve as the starting point for further layout optimizations with more sophisticated noise coupling models.

IV. Annealing-Based Macrocell Placement

In this section we will briefly describe our use of simulated-annealing for substrate-aware placement. We chose this op-

timization method because it has been successful for many different classes of placement problems [15]–[17] and because its characteristics allow seamless incorporation of the features necessary for substrate-awareness. Annealing is a global optimization method whose principal advantage is a controlled hill-climbing mechanism by which many local optima can be avoided. Annealing is itself a stochastic iterative improvement strategy: the perturbations (called *moves*) that evolve the solution, and their acceptance/rejection are random, although controlled by a parameter usually called *temperature*. Analogous to physical annealing, the problem starts at a high temperature in which large uphill moves are highly probable; in this high-temperature regime we see aggressive global search of the problem space. As annealing progresses, the temperature cools and uphill moves decrease both in size and in likelihood. Eventually, only minor local improvements are possible, and the problem is said to be *frozen*.

Any annealing algorithm can be specified by describing its four key components: the representation for each *state* (in our case, placement configuration) visited during iterative improvement; the *move-set* which specifies the list of allowable perturbations to these states; the *cost-function* that measures the quality of the evolving solution; and the *cooling schedule* that controls the rate of hill-climbing as the solution evolves from randomness to a final "frozen" solution.

Our algorithm uses the so-called flat (nonhierarchical) Gelatt–Jepsen annealing formulation of [15], and extensively refined in [17] and [16]. Macrocells are arbitrarily relocated about the surface of the chip. Illegal overlaps are permitted, but penalized. Hence, the state representation is just the coordinate locations and instance selections (each cell can have a set of alternate instances) for each movable macrocell in the current placement. The size and shape of each macrocell is input to the placer in terms of its *footprint*. The location and shape of each terminal in the cell is indicated on this footprint. A list of alternate *instances*—if any—for each macro is also provided. The algorithm starts with a randomly generated placement and iteratively improves the solution by applying moves from the move-set. Moves include *translating* a cell to new location, *reorienting* a cell by rotating or reflecting it, and *reinstancing* a cell by substituting a different instance for the cell. Efficient, general-purpose adaptive cooling schedules are available to match the rate of cooling (and the choice of which move to make) to the problem [21]. See [22] for an extensive discussion of the implementation details for this basic placement strategy. This placement flow is summarized in Fig. 2.

The key innovations for substrate-aware placement are in the move-set and in the cost-function and its evaluation. Most macrocell placement algorithms allow the selection of alternate instances to handle cells available in a range of different shapes and sizes. Our algorithm also incorporates this feature, but uses it not only for geometric packing but also for electrical optimization. We model the insertion of a guard ring around a cell as an "alternative instance" of that cell, i.e., we simply add to the palette of alternatives for this macrocell a version of this cell with a guard ring. This allows the placer to select a move that inserts or deletes a guard ring. The advantage is that the placer now sees directly how the guard ring changes the shape and size of the cell, and also its new demands for routing area in its vicinity. We refer to this as *active mitigation*, in the sense that it is the placer's task to determine the appropriate set of mitigation measures required to satisfy the specified constraints.

```
Algorithm Placer(Netlist, Constraints, Macro definitions) {
    Create random placement instance, P;
    Initialize annealing temperature T (see White's Algorithm [32]);
    do {
        Perturb P to P + ΔP using move from move-set;
        Evaluate cost-function for placement instance P + ΔP;
        Use Metropolis Criterion to determine acceptance or rejection of instance P + ΔP;
        If ( Accept ) P := P + ΔP;
        Use cooling schedule (e.g. [21]) to modulate T;
    } while ( not frozen );
    Output placement instance, P, as solution;
} /* Placer */
```

Fig. 2. A simulated annealing-based macrocell-placement algorithm.

Absent any concern for substrate coupling, the classical cost function includes estimated area (the smallest bounding rectangle of the cells) and wirelength (sum of half-perimeter of bounding box for each net's pins) terms to minimize. The penalty term measures illegal cell overlaps (area of overlapping footprints) and is driven to zero. To incorporate substrate-awareness, we allow the user to specify the maximum allowable switching noise at any location on any sensitive analog macro. Violations give rise to new penalty terms that must also be driven to zero in the cost function. Thus we need to model the substrate and the noise sources that inject noise into the substrate. We update an approximate substrate noise profile *after each placement move*, and then use this profile to determine the noise at any designer-specified locations on each sensitive macro. Each possible constraint violation is then checked and penalized as appropriate. Note that in contrast to simulation applications, we must update this substrate noise profile thousands of times as the movable cells converge to their final locations. Hence, our real problem is fast noise evaluation with "just enough" accuracy.

V. Modeling the Substrate

In this section, we develop a method for estimating substrate-coupled switching noise. Various tradeoffs among accuracy, generality, and efficiency are possible. For example, one can determine the substrate noise-profile by evaluating a two-dimensional Fourier series expression derived by solving the Laplace equation for noise in an epitaxial substrate using the approach described in [23]; such a model is very accurate, but computationally expensive, and inapplicable to the commonly used bulk case. For generality, most substrate simulators model the substrate in terms of an electrical equivalent: a grid of linear electrical elements derived using box-integration [24] techniques. For example, [6] and [8] use a three-dimensional resistive grid for each substrate layer; [25] uses a grid of resistances and capacitances for higher accuracy; and [4] represents the substrate as a single node (i.e., a degenerate grid) for the special case of an epitaxial process.

In the following, we develop a highly simplified gridded substrate model with accuracy sufficient for our placement application.

Despite its generality, the gridded substrate model has one major drawback: its accuracy is adversely affected by the coarseness of the grid. To improve accuracy, it is necessary to refine the grid in all three dimensions. Many substrate simulators try to overcome the efficiency issue by modulating the coarseness of the grid across the surface. Thus, a fine grid is used around the "interesting" areas of the substrate; for example, near transistors specified to be sensitive by the user. Unfortunately, this approach cannot be used in our substrate model because the location of the interesting objects—the noise receivers—*varies* after each annealing move, and it is computationally untenable to redefine the grid after *every* placement perturbation. Thus, we must use a static substrate grid with an adequate degree of refinement in each dimension.

Let us assume that the substrate is discretized to have $N = N_x \times N_y \times N_z$ grid elements. It has been observed empirically that the complexity of solving an electrical system with N nodes is about $N^{1.3}$ when sparse matrix techniques are used. Such substrate models has been used previously [5], [7], [8], but we believe that a more abstract model is desirable in this case. For example, the substrate can be fairly complex: it may have backside contacts and the lead-frame will certainly affect the noise profile. We describe such an abstract model subsequently. Such models have been used previously with some success [26].

To begin, we make the observation that only one part of the substrate is relevant during placement: the *active* area or the *face* of the substrate where placement is allowed. The placeable objects—noise sources and receivers—interact with the substrate through the grid points located on this surface. Therefore, we can characterize the entire substrate in terms of its *n-port admittance* (*y-parameter*) equivalent for the n grid points accessible on the surface of the substrate. Of course, any substrate, even the epitaxial type, can generally be represented as a grid and characterized as an n-port. The electrical equivalents of the macrocells can then be *plugged into* the nodes at the surface of the substrate. Each move of a single macrocell now entails disconnecting it from its nearest substrate nodes, relocating it, then reconnecting to the nearest nodes. Once this is done, the entire model (cells and substrate) can be solved for the substrate noise profile using common matrix solution techniques. This profile can then be used to determine the noise levels at the noise receivers.

An n-port equivalent is a multiport transfer function (where each element is a function of frequency) allowing evaluations of responses to transient stimuli. For example, [26] renders the substrate as a multiport s-domain transfer function using AWE [14] before computing thermal transients. But this degree of accuracy is not affordable in the context of automatic placement. We need merely characterize the n-port in terms of its admittance parameters at one fundamental frequency. In addition, the complexity of the substrate model has decreased in that the admittance model has many fewer nodes. The downside, however, is that the n-port admittance matrix is not sparse.

These issues are best examined in the context of a simple example. Consider a single-layer substrate and a user specification of $N = N_x \times N_y \times N_z$ grid points in the three dimensions. For a resistive, homogeneous substrate, the gridded equivalent circuit will have N nodes. A substrate simulator would convert this gridded equivalent and the circuitry on the chip to an admittance matrix of size N or more (reflecting the nodes in the circuitry on chip). This admittance matrix will be banded and sparse allowing the solution complexity to be marginally superlinear. The matrix is sparse because the elements are located only between nodes that are physically adjacent.

Assuming that the macros are placed on the $z = 0$ plane, the n-port equivalent circuit will have $N' = N_x \times N_y$ nodes and hence is a much smaller matrix. The diagonal terms in this equivalent matrix represent the equivalent admittance from the appropriate node to the reference node in the circuit. The off-diagonal terms represent the equivalent admittance between the appropriate nodes in the original network. This observation implies that this n-port equivalent will not be sparse because there is an electrical path between every pair of ports in the original circuit. Thus, it is likely that what we gained by reducing the size of the substrate model from N to N' was lost to reduced sparsity. This is not unexpected because the circuits are exact electrical equivalents of each other.

However, we can actually overcome this problem by exploiting the structure of the annealing process. In any annealing optimizer, convergence toward a final solution (a placement) is gradual and incremental. Early in the process (at hotter temperatures), large, uphill moves are allowed, which vigorously explore the solution space. In these temperature regimes, great accuracy in the solution of the substrate noise profile is unnecessary. What we seek instead, for efficiency, is a variable accuracy model of the substrate which can increase in detail, accuracy—and computational expense—as we converge to a final solution.

Fortunately, there is a computationally simple method to achieve this variable accuracy for an n-port substrate model: we *threshold* the nondiagonal terms in the matrix. The idea is that by removing small elements from the matrix, we render it sparse and thus fast to solve. At the beginning of annealing, only the diagonal terms are present in the matrix; as annealing proceeds toward a solution, more and more off-diagonal terms appear. Near a final solution, the entire substrate model is employed. This can be achieved by establishing a threshold and monotonically reducing it to zero across the optimization process. If $\boldsymbol{Y_S} = [y_{ij}]$ is the n-port matrix, a suitable formulation is to set $y_{ij} = 0$ if it satisfies

$$ty^{max} + (1 - t)y^{min} \geq |y_{ij}|, \qquad t = [1, 0]. \tag{1}$$

Here, y^{max} and y^{min} are the maximum and minimum nondiagonal terms in the matrix. Parameter t varies from 1 to 0 across the annealing process; at the start, $t = 1.0$ and all nondiagonal terms are excluded.

We illustrate the value of thresholding through a series of experiments summarized in Fig. 3. A placement instance is shown in Fig. 3(a), and its noise contours (without thresholding) are shown in Fig. 3(b). Note that the noise level outside of the noise-source regions is about 200 mV. The effects of

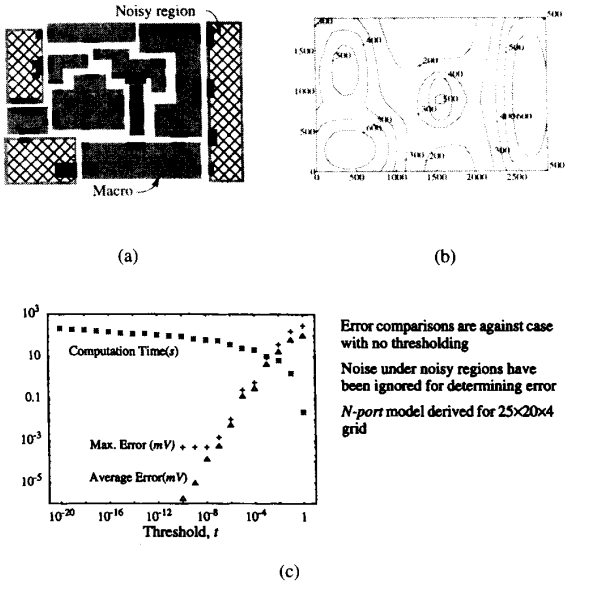

Fig. 3. Effect of thresholding on error and evaluation time. (a) Placement instance. (b) Noise contours (mV), no thresholding. (c) Effect of thresholding.

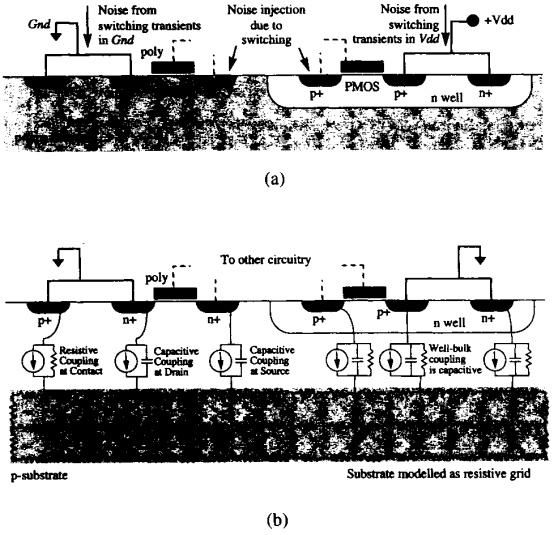

Fig. 4. Noise injection in CMOS logic circuits. (a) Noise generation mechanisms in CMOS. (b) Norton equivalents for substrate coupling.

thresholding are depicted in Fig. 3(c). Naturally, we expect the error to be the highest for $t = 1.0$. For this case, the average and the maximum errors are about 100 and 300 mV, respectively. These large errors decrease to about 4 and 10 mV, respectively, when the threshold is down to $t = 10^{-3}$. Within the same threshold range, the time per evaluation increases from 20 ms to 20 s. Thus, great improvements in efficiency can be obtained by thresholding the model when the error in the noise levels can be tolerated. A sparse matrix solver [27] running on a DECStation 5000^{TM}/200 was used in this experiment.

Note, however, that the times involved are still rather large given that repetitive evaluations would be required. One of the reasons is that the grid is quite fine—25 × 20—which implies that the *n-port* model can have as many as one-quarter of a million elements. The sparse matrix solver requires as much as 250 s per evaluation when the threshold is very low. Thus, the solver should employ a dense matrix technique when the matrix is no longer sparse.

VI. Modeling Substrate Interaction

As mentioned in Section II, the interaction of the macros and the substrate can be abstracted into a simple model comprised of the noise sources, the substrate, and the noise receivers. In this section, we will define simple models for noise sources and receivers, and indicate how they can be combined with the substrate model for evaluating switching noise.

A. Noise Sources

As mentioned in Section III, a simplified model is used for noise generation. Because substrate coupling occurs through both capacitors and resistors, it is frequency dependent. However, in order to simplify the analysis which must occur in the inner loop of an annealer, we assume that there is a single user-selected frequency (e.g., the clock frequency) that will provide a good first-order estimation for signal coupling. Although the assumption that all the noise in the system is at one (the fundamental) frequency only is simplistic, it has the advantage that it greatly simplifies the inputs that must be provided by the designer concerning the nature of the noise sources. To characterize a more complex time domain noise source might be extremely difficult. More accurate models would necessitate transient simulation and temporal description of the noise waveforms, both of which would be infeasible and expensive within an iterative placement framework, and, we argue, unnecessary at the very early stages of the layout process which we target.

In addition, all of the resistive and capacitive couplings to the substrate within a module are collected together into one or a few geometrically localized patches (the number would be determined by the cell size compared with the substrate grid size). These Norton equivalent *regions* are characterized by a complex current source and an admittance. Many such regions may be present on a digital macrocell. Then, a magnitude must be assigned to the sinusoidal noise current injected across the admittance at each of these patches at the fundamental frequency selected. For example, this magnitude can be derived from circuit simulations of functional blocks with representative or worst-case input waveforms. Finally, the designer must provide constraints on the magnitude of substrate voltage variations and on the gradients in substrate voltage that will ensure correct operation of sensitive analog cells.

How does the designer determine the admittance and current source values for a noise generating cell? Generally, there are three mechanisms that inject noise into the substrate. These are shown in Fig. 4(a). Power buses couple noise arising out of switching transients into the substrate through ohmic contacts;

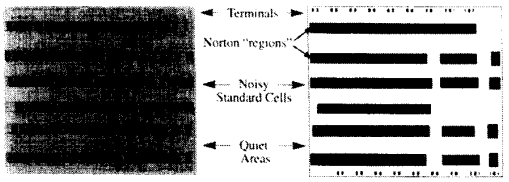

(a) (b)

Fig. 5. Modeling a digital macro. (a) Digital macrocell. (b) Footprint.

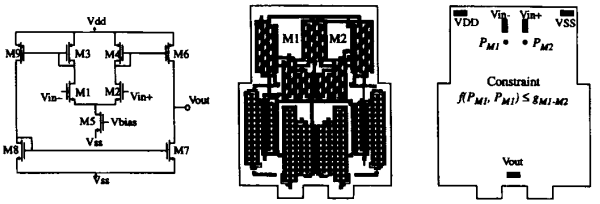

(a) (b) (c)

Fig. 6. The model for a two-stage op-amp. (a) Schematic. (b) Layout. (c) Footprint.

wells couple switching transients capacitively through reverse biased bulk/well junctions; and switching devices (transistors) couple noise capacitively through source/drain diffusions. The noise sources at the contacts can be modeled with an equivalent current source across a contact resistance. Current sources with shunting capacitances can also be used to model the effect of the charge injected into the substrate at the wells and the source/drain diffusions of the switching devices. A simplified model of such substrate interactions is shown in Fig. 4(b): a more accurate model of such substrate interactions (for example, one including the series resistance of the reverse biased source/drain diffusion region "diodes") would involve more elements to model some of the second-order effects. In a circuit simulator, the admittance between the circuit and the substrate can be determined at the selected frequency. Next, a transient circuit simulation of one or more functional blocks with representative or worst-case input waveforms can be carried out. Then, by measuring the currents flowing into the substrate from all elements in each patch (e.g., body contact of transistors) during the transient simulation, a magnitude and approximate sinusoidal frequency can be chosen for the current waveform.

For our placer, the geometry of a digital macrocell is specified in terms of its *footprint*. An example of a digital macro and its footprint appears in Fig. 5. Besides the outline and the terminal description, areas generating noise are indicated on the footprint and tagged with their Norton equivalent. For example, the noisy areas would correspond to the location of the noisy standard cells in a row-based functional block.

B. Noise Receivers

Substrate simulators usually couple sensitive (high impedance) signals or transistors to the substrate through parasitic capacitors and resistors. However, we require a much simpler coupling model. Thus, we will assume that the designer is able to provide *constraints* on the noise levels and the gradients in noise level at sensitive locations on analog macros. This scenario is best explained through an example.

Consider the two-stage operational amplifier (op-amp) depicted in Fig. 6. The schematic is given in Fig. 6(a) (some biasing circuitry is omitted for clarity). Fig. 6(b) is the layout for this macro and Fig. 6(c) shows its footprint. The footprint encodes the outline of the cell and the location, shape, and layers of its terminals. In addition, the cell may now be tagged with a substrate-noise constraint. In practice, substrate noise-induced mismatches between any pair of matched devices would disturb the output, but here we focus on the mismatches of the input devices only for reasons of clarity and because their effect is the greatest. For example, the impact of the mismatches between $M3$ and $M4$ in Fig. 6(a) is reduced by the gain across the input devices $M1$ and $M2$ when referred to the input.

The input devices are susceptible to substrate noise via two different mechanisms. First, due to manufacturing mismatches between their body effect coefficients ($\Delta\gamma$), variations in substrate voltage will be manifested as variations in input offset voltage. Second, even if the body effect coefficients match perfectly, a gradient in substrate potential will still result in a noise-dependent input offset voltage. Thus, if $M1$ is located at $P1$ and $M2$ is located at $P2$, the user may specify a constraint (2) aimed at keeping the potential gradient within a limit, $g_{M1\text{-}M2}$:

$$f(P_{M1}, P_{M2}) = \frac{|\text{Noise}\,(P_{M1}) - \text{Noise}\,(P_{M2})|}{|P_{M1} - P_{M2}|} \leq g_{\text{M1-M2}}. \tag{2}$$

Such gradients can easily be determined from potential profiles in the substrate. The placer will attempt to locate this macro such that this constraint is satisfied. The op-amp designer is usually aware of such substrate sensitivities, and is therefore capable of tagging such constraints onto the footprint of the cell shown in Fig. 6(c). Extremely sensitive analog macros are often surrounded by guard rings. Such guard rings can also be modeled by our Norton equivalent (i, y). In this case, i is zero and the admittance y is entirely resistive.

A simple example of this substrate-interaction model is shown in Fig. 7. A placement instance is shown in Fig. 7(a). The corresponding noise profile on a 10×10 substrate grid is shown in Fig. 7(b). This profile is determined from the electrical equivalent of substrate interaction depicted in Fig. 7(c). A bulk substrate without a backside contact is assumed. (This is typical for inexpensive packages like the plastic quad flat pack.) Noise couples in through direct coupled current sources at the grid points located under the noise source indicated in Fig. 7(a). (For clarity of illustration, we make the simplifying assumption that the substrate contacts are uniformly distributed across the substrate and hence characterized into the substrate model.) A guard-ring and its electrical equivalent are also indicated. Note that the profile in Fig. 7(b) flattens out in the region surrounded by the guard ring. Naturally, the guard ring serves as a low impedance path (typically, the series combination of the diffusion and the

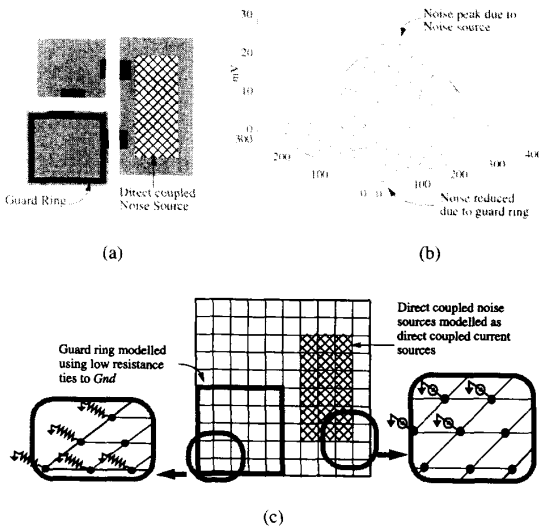

Fig. 7. Substrate interactions for a given placement instance. (a) Placement instance. (b) Noise profile. (c) Electrical equivalent of placement instance.

TABLE I
SUMMARY OF EXPERIMENTS WITH EXAMPLE $C1$

Expt.	Constraint (V)	Normalized		Time (min)
		Area	WireLength	
1	-	1	1	4
2	$V_n(A) \leq 0.6$	1	2.33	139
3	$V_n(A) \leq 0.1$	1.1	1.166	200

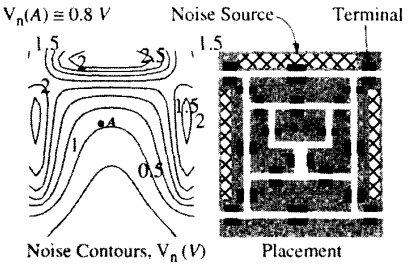

Fig. 8. Placement without considering substrate noise.

wiring) to a quiet potential, thereby flattening out the noise profile by "intercepting" the substrate noise before it reaches the protected area. Usually, this ring would be connected to the low-noise analog ground supplying the analog cells on the chip.

C. Evaluating Substrate Coupled Switching Noise

A placement instance combined with the models described above can be easily solved to determine the substrate noise profile. Note that noise sources and receivers have been modeled such that they do not add any extra nodes to the electrical equivalent. Let the n-port admittance (y-parameter) model of the substrate be $\boldsymbol{Y_S}$. This is a matrix of size $N' \times N'$, where N' is the number of grid points on the active area of the substrate. The electrical interaction of the placement instance and the substrate is represented by a diagonal matrix, $\boldsymbol{dY}$ of size $N' \times N'$. A diagonal element in this matrix is nonzero if there is a macro with a nonzero Norton admittance located at the substrate node corresponding to that element. Let $\boldsymbol{J}$ be a vector of size N' such that an element therein is nonzero if there is a macro with a nonzero Norton current source located at the substrate node corresponding to that element. Also, let $\boldsymbol{V}$ be the vector (of size N') representing the potential profile across the substrate measured against an arbitrarily selected reference potential for the entire system. Then from basic circuit theory [28]

$$[\boldsymbol{Y_S} + \boldsymbol{dY}]\boldsymbol{V} = \boldsymbol{J}. \tag{3}$$

In general, this equation would be complex and dense. However, it may be rendered sparse by applying the thresholding methods described in Section V to the substrate model $\boldsymbol{Y_S}$. Also, the designer may choose to use a simplified substrate interaction model without any capacitances. In this case, (3) is realvalued and a solver tuned to solving real matrices can be used. The matrix $[\boldsymbol{Y_S} + \boldsymbol{dY}]$ is also diagonally dominant, implying that efficient iterative solution schemes like the *Incomplete Choleski Conjugate Gradient* [29] can be used to achieve further gains in speed.

VII. EXPERIMENTAL RESULTS

We will now describe and analyze a small set of experiments that demonstrate the viability of the placement techniques presented here. The algorithms described in the preceding sections have been implemented in a prototype substrate-aware mixed-signal placement tool called WRIGHT. This program is about 14 000 lines of C code. A program to generate n-port substrate models using a circuit simulator [30] and simple structural and electrical descriptions of the substrate was also implemented. WRIGHT incorporates a sparse [27] and a dense [31] matrix package. All examples decribed here were executed on a DECStation 5000^{TM}/200 workstation.

To begin, we ran a set of controlled experiments on a simple synthetic example, $C1$. This has 9 blocks, 20 nets, 40 terminals, 3 noisy macros, and 1 sensitive macro. This example is essentially a jigsaw puzzle: the macros here are rectilinear slices out of a rectangle. A placement that finds the global optimum with respect to area and wire length reassembles the pieces back into the original solid rectangle. These experiments are summarized in Table I, and the results are depicted in Figs. 8–10. The noise contours and the final placement are given for each. In these figures, each gray box is a macro, the darker boxes are terminals, and the hatched boxes are noise sources. A bulk substrate abstracted from a 10 × 10 × 2 grid is used.

In experiment 1, we place to minimize area and wire length *without* any regard to substrate noise. The solution (Fig. 8) has both minimum wire length and area. As expected, WRIGHT

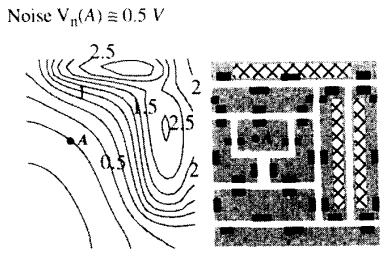

Fig. 9. Constrain $V_n(A) \leq 0.6$ V.

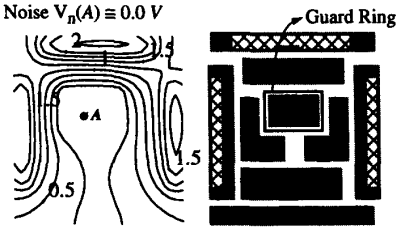

Fig. 10. Constrain $V_n(A) \leq 0.1$ V.

just reassembles the puzzle when ignoring the substrate noise. Note that the noise level at location A is predicted as 0.8 V.

In experiment 2, we impose a constraint on the acceptable substrate noise level at A, a location on the sensitive macro. The new placement is shown in Fig. 9. Here WRIGHT attempts to satisfy this constraint by moving the sensitive macro away from the noisy ones. WRIGHT satisfies the constraints by appropriately choosing the relative locations of sensitive and noisy macros.

Finally, we tighten the noise constraint further in experiment 3. The placer now decides that the area penalty involved in further separating the macro containing A is higher compared to putting a guard ring around that macro (Fig. 10). This is an example of the placer making a noise/area tradeoff in favor of actively mitigating a noise constraint violation.

From the results given in Table I, we can draw several conclusions. Not unexpectedly, the CPU time increases when noise constraints are activated. Note, however, only one matrix solve *per* placement iteration is required to evaluate the noise at all receivers: the number of noise constraints or noise receivers does not affect the efficiency of this evaluation. However, like all other annealing placement algorithms, the time required to complete placement is a function of the *size* of the problem as quantified by the number of macros.

In Fig. 11, we place a larger example, $C2$, which is a simplified version of the mixed-signal routing benchmark from [10] and [8]. This example has 25 macros, 381 terminals, 163 nets, 3 noisy digital macros and 12 sensitive analog macros with receivers located at their centers. The analog macros have constraints similar to example $C1$. The constraints and the results of running WRIGHT are tabulated in Table II. We note that WRIGHT satisfies a greater number of constraints when the consideration of substrate noise is enabled. Moreover, the

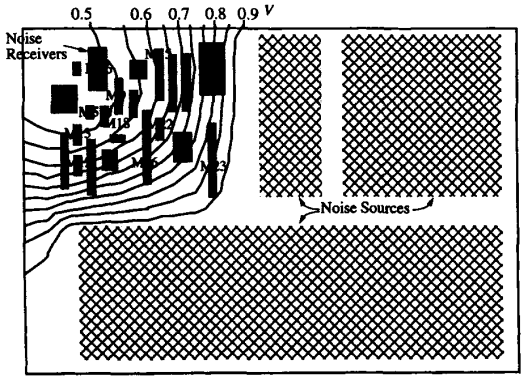

Fig. 11. Noise contours for example $C2$ placed with 12 noise constraints.

TABLE II
PLACING A LARGE EXAMPLE $C2$

Receiver, i	Constraints		Substrate Noise			
	$V_i^{noise} \leq V_i^{max}$		Ignored		Considered	
	$V^{max}{}_i(V)$	Weight, w_i	$V_i^{noise}(V)$	Violation, v_i	$V_i^{noise}(V)$	Violation, v_i
M1	0.9	1	0.58		0.76	
M5	0.2	100	0.88	0.68	0.46	0.26
M12	0.8	1	0.57		0.69	
M13	0.8	1	0.61		0.73	
M15	0.8	1	0.80		0.69	
M9	0.2	10	0.69	0.49	0.51	0.31
M18	0.7	1	0.66		0.56	
M16	0.9	1	0.63		0.48	
M22	0.9	1	0.75		0.89	
M26	0.6	1	0.89	0.28	0.71	0.11
M23	0.7	1	0.92	0.22	0.84	0.14
M24	0.7	1	0.83	0.13	0.48	

Placement metrics		
$\sum w_i v_i$	73.53	29.35
Wire-Length (μm)	252449	282461
Time (*min*)	2.87	433
Area (μm^2)	3.6×10^7	3.6×10^7

degree of constraint violation is also minimized. WRIGHT uses a simple weighted sum of constraint violations (noise voltages), $\sum_i w_i v_i$, as the penalty term to be driven toward zero here. By design here, the noise constraints for M5 and M9 cannot be met, so WRIGHT attempts to minimize the constraint violation for each macro. However, the constraint on M9 has greater weight attached to it, so the placer attempts to reduce the violation for M9 more aggressively. So, as Fig. 11 shows, M9 is placed at a quieter location a greater distance away from the noisy digital macros compared to M5, thereby reducing the substrate noise constraint penalty term. Overall, we see that wire length is traded off to achieve these gains. (The area does not change because the bounding box of the chip is fixed.) This more complex example requires about 7 hours to place when using a substrate model derived from a rather fine $15 \times 20 \times 3$ grid. Improvements in time efficiency can be expected with a coarser grid. This example was run with thresholding enabled.

VIII. Conclusions

In this paper we presented a strategy for substrate-aware mixed-signal macrocell placement. This approach incorporates simplified switching noise estimation into a simulated anealing placement algorithm. Preliminary results suggest our algorithms generate placements that are intuitively correct. This satisfies our goal of generating reasonable first-cut placements that consider the most significant effects. One of the basic principles underlying the selection and design of our substrate and substrate interaction models has been the tradeoff between efficiency and accuracy. This is necessary because the models are meant to be used during a part of the design cycle when little information about the completed chip is available and efficient evaluation is critical. We expect that our placements will be incrementally tuned based on results from downstream layout steps (e.g., power grid and signal routing) and full-chip substrate simulation. We also expect that further work on estimation of power dissipation and switching activity [20] for digital circuitry will be directly applicable to our simple noise injection models.

Our current work focuses on improving the speed and the accuracy of our models and on extending this method for handling thermal layout problems [26].

Acknowledgment

The authors would like to thank W. Getchell of Bell-Northern Research, B. Stanisic of IBM, G. Gad-Elkarim of Cadence Design Systems, and N. K. Verghese of Carnegie Mellon University for their comments and the highly stimulating discussions during the initial phases of this project.

References

[1] D. A. Hodges, P. R. Grey, and R. W. Brodersen, "Potential for MOS technologies for analog integrated circuits," *IEEE J. Solid-State Circuits,* vol. SC-13, pp. 285–294, June 1978.

[2] J. R. Lineback, "ASIC houses rush to add analog functions to libraries," *Electron.* vol. 60, no. 22, p. 31, Oct. 29, 1987.

[3] J. A. Olmstead and S. Vulih, "Noise problems in mixed analog/digital integrated circuits," in *Proc. Custom Integrated Circuits Conf.,* 1987, pp. 659–662.

[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, pp. 420–430, Apr. 1993.

[5] N. K. Verghese, D. J. Allstot, and S. Masui, "Rapid simulation of substrate coupling effects in mixed-mode IC's," in *Proc. IEEE Custom Integrated Circuits Conf.,* May 1993, pp. 18.3.1–18.3.4.

[6] N. K. Verghese, T. J. Schmerbeck, and D. J. Allstot, *Simulation Techniques and Solutions for Mixed-Signal Coupling in Integrated Circuits.* Norwell, MA: Kluwer Academic, 1994.

[7] B. R. Stanisic, N. K. Verghese, R. A. Rutenbar, L. R. Carley, and D. J. Allstot, "Addressing substrate coupling in mixed-mode IC's: Simulation and power distribution synthesis," *IEEE J. Solid-State Circuits,* vol. 29, pp. 226–238, Dec. 1994.

[8] B. R. Stanisic, R. A. Rutenbar, and L. R. Carley, "Power distribution synthesis for analog and mixed-signal ASIC's in RAIL," in *Proc. IEEE Custom Integrated Circuits Conf.,* May 1993, pp. 17.4.1–17.4.5.

[9] B. R. Stanisic, R. A. Rutenbar, and L. R. Carley, "Mixed-signal noise-decoupling via simultaneous power distribution design and cell customization in RAIL," in *Proc. IEEE Custom Integrated Circuits Conf.,* May 1994, pp. 24.3.1–24.3.4.

[10] S. Mitra, S. K. Nag, R. A. Rutenbar, and L. R. Carley, "System-level routing of mixed-signal ASIC's in WREN," in *Proc. IEEE Int. Conf. Comput.-Aid. Design,* 1992, pp. 394–399.

[11] U. Chowdhury and A. Sangiovanni-Vincentelli, "Constraint-based channel routing for analog and analog/digital circuits," in *Proc. IEEE Int. Conf. Computer-Aided Design,* Nov. 1990, pp. 198–201.

[12] A. E. Dunlop, G. F. Gross, C. D. Kimble, M. Y. Luong, K. J. Stern, and E. J. Swanson, "Features in LTX2 for analog layout," in *Proc. ISCAS-85,* 1985, pp. 21–23.

[13] T. J. Schmerbeck, R. A. Richetta, and L. D. Smith, "A 27MHz mixed analog/digital magnetic recording channel DSP using partial response signalling with maximum likelihood detection," in *Proc. IEEE Int. Solid-State Circuits Conf.,* 1991, pp. 136–137.

[14] L. T. Pillage and R. A. Rohrer, "Asymptotic waveform evaluation for timing analysis," *IEEE Trans. Computer-Aided Design,* vol. 9, pp. 352–366, Apr. 1990.

[15] D. W. Jepsen and C. D. Gelatt, Jr., "Macro placement by Monte Carlo annealing," in *Proc. 1984 IEEE Int. Conf. Comput. Design,* Nov. 1984, pp. 495–498.

[16] J. M. Cohn, D. J. Garrod, R. A. Rutenbar, and L. R. Carley, "KOAN/ANAGRAM II: New tools for device-level analog placement and routing," *IEEE J. Solid-State Circuits,* vol. 26, pp. 330–342, Mar. 1991.

[17] C. Sechen and A. Sangiovanni-Vincentelli, "The TimberWolf placement and routing package," *IEEE J. Solid-State Circuits,* vol. SC-20, pp. 510–522, Apr. 1985.

[18] B. T. Preas and M. J. Lorenzetti, Eds., *Physical Design Automation of VLSI Systems.* Menlo Park, CA: Benjamin/Cummins, 1988.

[19] S. Masui, "Simulation of substrate-coupling in mixed-signal MOS Circuits," in *Proc. VLSI Circuits Symp.,* 1992, pp. 42–43.

[20] A. Yang and K. Mayaram, "Estimating power dissipation in VLSI circuits," *IEEE Circuits Devices,* pp. 11–19, July 1994.

[21] M. D. Huang, F. Romeo, and A. Sangiovanni-Vincentelli, "An efficient general cooling schedule for simulated annealing," in *Proc. IEEE Int. Conf. Computer-Aided Design,* 1986, pp. 381–384.

[22] J. M. Cohn, D. J. Garrod, R. A. Rutenbar, and L. R. Carley, *Analog Device-Level Layout Automation.* Norwell, MA: Kluwer Academic, 1994.

[23] A. G. Kokkas, "Thermal analysis of multiple-layer structures," *IEEE Trans. Electron. Devices,* vol. ED-21, pp. 674–681, Nov. 1974.

[24] K. Fukahori and P. R. Grey, "Computer simulation of integrated circuits in the presence of electrothermal interaction," *IEEE J. Solid-State Circuits,* vol. SC-11, pp. 834–846, Dec. 1976.

[25] N. K. Verghese, S.-S. Lee, and D. J. Allstot, "A unified approach to simulating electrical and thermal substrate coupling interactions in IC's," in *Proc. Int. Conf. Comput.-Aid. Design,* Nov. 1993.

[26] S.-S. Lee and D. J. Allstot, "Electrothermal simulations of integrated circuits," *IEEE J. Solid-State Circuits,* vol. 28, pp. 1283–1293, Dec. 1993.

[27] K. S. Kundert and A. Sangiovanni-Vincentelli, *Sparse User's Guide,* Version 1.3a ed., Dep. Elec. Eng. Comput. Sci., Univ. California, Berkeley, 1988.

[28] J. Vlach and K. Singal, *Computer Methods for Circuit Analysis and Design.* New York: Van Nostrand Reinhold, 1983.

[29] 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.

[30] Meta-Software, *HSPICE User's Manual,* Meta-Software Inc., Campbell, CA, 1987.

[31] J. J. Dongarra *et al., LINPACK User's Guide,* Soc. Ind. Appl. Math., Philadelphia, 1979.

[32] S. R. White, "Concepts of scale in simulated annealing," in *Proc. IEEE Int. Conf. Comput. Design,* 1984, pp. 646–651.

Sujoy Mitra (S'94) received the B.Tech. degree in electronics and electrical communication engineering from the Indian Institute of Technology, Kharagpur, India, in 1987, and the M.S. degree in electrical and computer engineering from Carnegie Mellon University in 1992, where he is currently working toward the Ph.D. degree.

He has held software and design engineering positions at Texas Instruments, Bangalore, India, during 1987–1990. His research interest include physical design algorithms and design automation of mixed-signal systems.

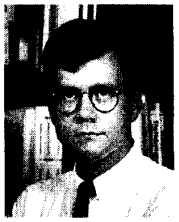

Rob A. Rutenbar (S'77–M'78–SM'90) received the Ph.D. degree in computer engineering (CICE) from the University of Michigan, Ann Arbor, in 1984.

He then joined the Faculty of Carnegie Mellon University, Pittsburgh, PA, where he is currently Professor of Electrical and Computer Engineering, and of Computer Science, and Director of the Semiconductor Research Corporation-CMU Center of Excellence in CAD and IC's. His research interests focus on circuit and layout synthesis for mixed-signal ASIC's, for high-performance digital IC's, and for FPGA's.

Dr. Rutenbar received a Presidential Young Investigator Award from the National Science Foundation in 1987. 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, and Program Committees for the ACM/IEEE Design Automation Conference and European Design & Test Conference. He is on the Editorial Board of IEEE SPECTRUM, and chairs the Analog Technical Advisory Board for Cadence Design Systems. He is a member of ACM, Eta Kappa Nu, Sigma Xi, and AAAS.

L. Richard Carley (S'77–M-81–SM'90) received the S.B. degree from the Massachusetts Institute of Technology, Cambridge, MA, in 1976 and was awarded the Guillemin Prize for the best EE Undergraduate Thesis; he received the M.S. degree in 1978 and the Ph.D. degree in 1984, also from MIT.

He is currenly 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 received 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.

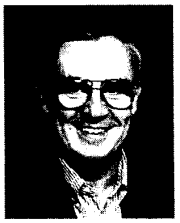

David J. Allstot (S'72–M'72–SM'83–F'92) received the B.S. degree in engineering science from the University of Portland, Portland, OR, the M.S. degree in electrical and computer engineering from Oregon State University, Corvallis, and the Ph.D. degree in electrical engineering and computer science from the University of California at Berkeley in 1979. His Ph.D. work dealt with the analysis, design, and implementation of switched-capacitor filters.

He has held industrial positions with Tektronix, Texas Instruments, and MOSTEK, and academic positions with UC Berkeley (Acting Assistant Professor and Visiting MacKay Lecturer), Southern Methodist University, Dallas, TX (Adjunct Assistant Professor), and Oregon State University (Associate Professor). 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. He has advised about 40 M.S. and Ph.D. students and has published more than 100 papers with students and colleagues.

Dr. Allstot was a co-recipient of the 1980 IEEE W.R.G. Baker Prize Paper Award, and has received excellence in teaching awards from SMU, OSU, and CMU. He has served as Associate Editor of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS, Guest-Editor for the IEEE JOURNAL SOLID-STATE CIRCUITS, Technical Program Committee member of the IEEE Custom IC Conference, the IEEE International Symposium on Circuits and Systems, the 1994 Symposium on Low Power Electronics, and 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 SIGNAL PROCESSING and a member of the Technical Program Committee of the IEEE International Solid-State Circuits Conference. He is member of Eta Kappa Nu and Sigma Xi.