A CAD System for Modeling Voltage Drop and Electromigration in VLSI Metallization Patterns

THE currents present in VLSI conductor patterns are responsible for two major problems: Electromigration of metal atoms and voltage drop. Electromigration places a limit on the useful lifetime of the circuit, and, if not properly accounted for, can present a serious reliability problem.

As a phenomenon that happens over an extended period of time, electromigration depends on all possible current waveforms; any estimation of the electromigration rate must be independent of individual input patterns.

Excessive voltage drop in conductor patterns can produce marginal, or unacceptable, circuit performance and require extensive redesign, greatly increasing circuit design and development costs. Voltage drop is an instantaneous phenomenon; it depends on the peak current at a particular moment.

The need for CAD programs, capable of ensuring adequate current-carrying capacity in VLSI circuits, is becoming increasingly apparent as density increases and line widths shrink. These programs are required in the initial design phase and during layout verification. They must be capable of ensuring that both electromigration and voltage drop are within acceptable limits.

One approach to the problem would be to use circuit analysis to determine current loads on the conductor pattern at various locations and solve a two-dimensional Poisson equation, using finite element or finite difference analysis, to determine voltage and current density at every point in the conductor.

This method would use an enormous amount of computer resources for a VLSI circuit, and it would do little to aid the circuit designer during the initial design phase. Furthermore, it would not address the pattern independence needs of electromigration analysis.

We have developed a simpler approach which automates many of the tasks previously performed by the circuit designer. The result is a system of programs which can be used at any level of circuit design.

Two programs are used to make up this system. SPIDER[1] uses SPICE[2] simulation of extracted resistances and capacitances to determine the current density and voltage at each point of the extracted metal structures. Regular sections of the metal structures are extracted as two-terminal resistive elements, while difficult sections are extracted as collections of three-terminal finite elements.

To model the effects of the circuit on the extracted RC network, another program, CREST,[3] is used to provide current sources to load the network during SPICE simulations. CREST is capable of supplying a statistically meaningful, expected current waveform for use in analyzing electromigration rate, or, if the user is willing to supply explicit input waveforms, CREST can estimate the current waveform for these inputs, which can then be used to determine voltage drop.

While SPIDER can be used for any metal line on any design, CREST is intended to provide current sources for CMOS connections to power and ground. As conditions warrant, CREST will be expanded for use in other technologies.

System Flow

The steps necessary for checking the current density in an existing layout are detailed in the following description and shown in an abbreviated form in Figure 1.

1. The user specifies, on a separate ICE mask layer, blocks of circuitry that can be considered as drawing power from the same points, on the power and ground buses, and the connection points for each block.
2. Resistances and capacitances of the conductor pattern are determined using a parasitic element extractor. The connection points specified for the blocks are used as the primary electrical nodes for the extracted elements.
3. The MOSFETs and parasitic capacitances of each block are extracted. All power (or ground) connections of devices in the block are to the node which corresponds to the power (or ground) connection specified by the user.
4. The user specifies either a single logical waveform for each circuit input, or a large number of input waveforms and the probability that each input occurs. CREST utilizes these inputs during its current waveform estimations.
5. CREST is used to simulate the extracted MOSFETs and provide transient current waveforms for each power and

Richard Burch, Joseph Hall, Farid Najm, Dale Hocevar, Ping Yang, and Michael McGraw

ground connection node. These current waveforms are converted into SPICE current sources and included with the extracted SPICE deck.

6. SPICE is used to determine the waveform of the current flowing through each of the resistances.

7. The current waveforms are converted to "continuous dc electromigration equivalent values,"[4] and revised line widths are computed.

8. The circuit analysis and adjustment steps (6 and 7) are repeated as necessary to produce the final set of recommended line widths. This iterative procedure is essential, since there may be more than one current path. Current in a particular section of a line can be a function of line width, and changes in width can change the distribution of current among the various paths.

9. The recommended line width changes are made in the layout data-base by the designer. Designer action, rather than automatic layout modification, is necessary to prevent the changes from causing design rule violations.

When this system is used in the design mode, no layout is available for extraction. The designer specifies proposed line lengths and widths, and the program checks current density and voltage drop. The designer must also specify the blocks of circuitry that are connected to various points of the proposed power buses. Connectivity information between the blocks, along with input waveforms for the circuit inputs or input waveforms for each block of circuitry, must be included. When used in the design phase, the result is a list of recommended line widths.

SPIDER is used interactively and iteratively to provide guidance to the designer during the initial design and layout phases. Use of the program in the design phase can ensure that any changes required by the final check, described above, will be minor and will not require extensive layout modification.

A feature of SPIDER which is par-

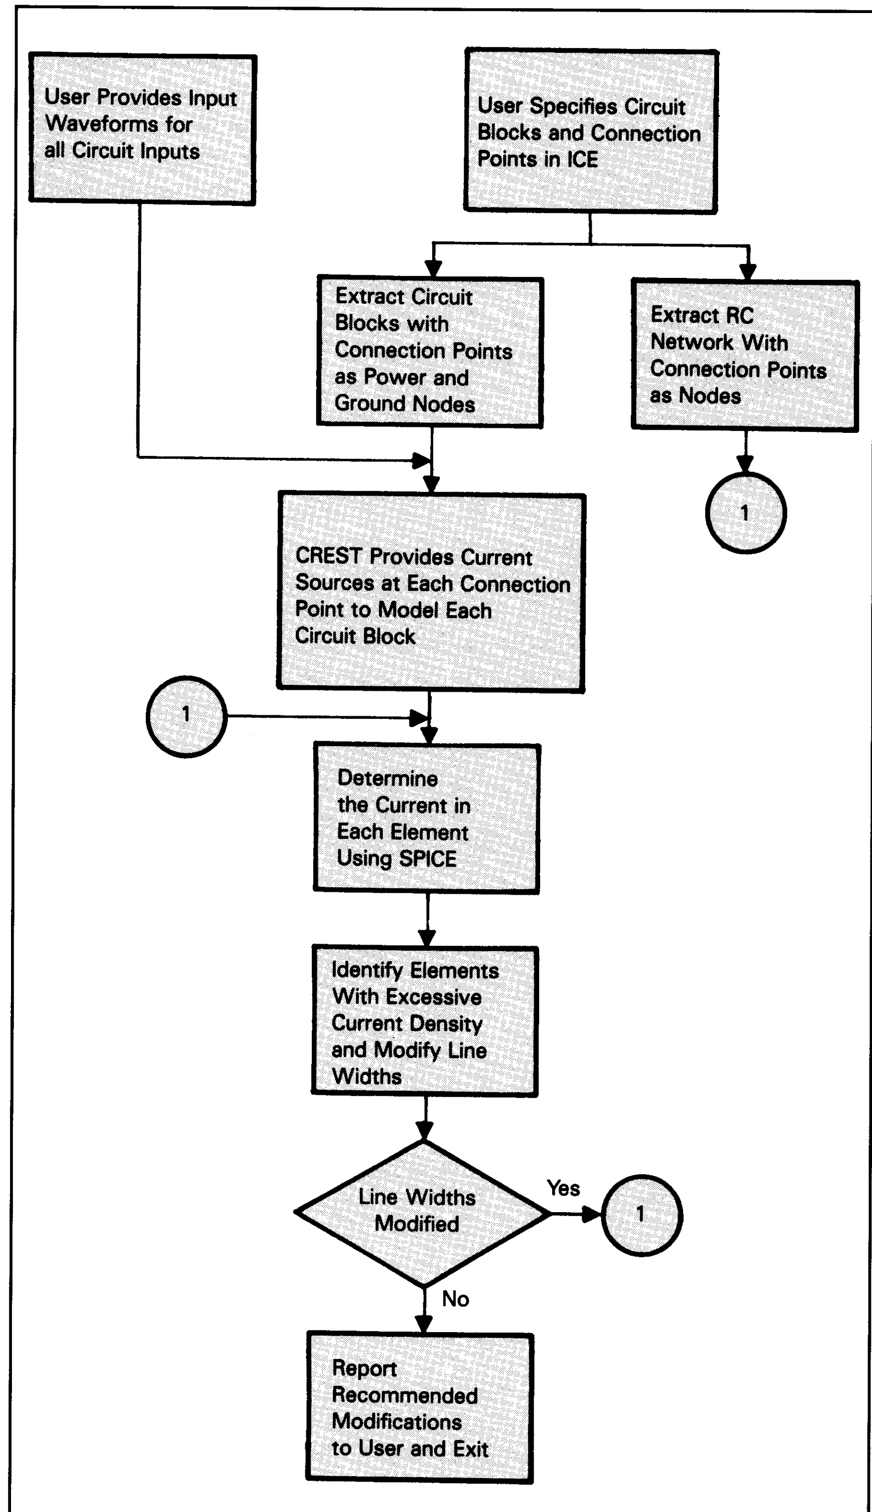

Figure 1. Flow Process of CAD System

ticularly useful in the design mode is the acceptance of any valid SPICE input. This feature allows the designer to include all or part of the actual circuitry, instead of the current waveforms provided by CREST. If this is done for CMOS, the effect of changing line resistance on current waveforms is automatically taken into account during SPIDER iterations.

Extraction and Simulation Approach

It is possible to implement a completely general finite element system designed to calculate the current density in the metal interconnect layer of an integrated circuit; however, this approach would be expensive and inefficient. An alternative method is to model the interconnect in terms of three-terminal finite elements and two-terminal resistive elements. Circuit simulation with SPICE can then be used to perform the electrical analysis. This approach is efficient and uses existing tools that designers are familiar with.

Most of the layout regions to be treated in the computation of current densities are regular in shape. Metal interconnect is primarily composed of long rectangular regions. Only a few regions are irregular, and these are usually at the juncture of regularly shaped rectangular interconnect regions.

Two-terminal elements are used to represent rectangular regions in which the direction of current flow can be determined from the geometry, as shown in Figure 2. Three-terminal triangular elements are extracted and used to represent irregular regions where the current flow direction is unknown. Each triangular region is equivalent to a three-resistor delta configuration (Figure 2), where each resistance is computed from the geometry of the triangle.

With these elements, the entire geometrical layout can be converted to a resistor network. Capacitance is accounted for by inserting discrete capacitors at the corners of the triangular elements and at the end points of the two-terminal elements. The result is an RC network suitable for analysis by SPICE.

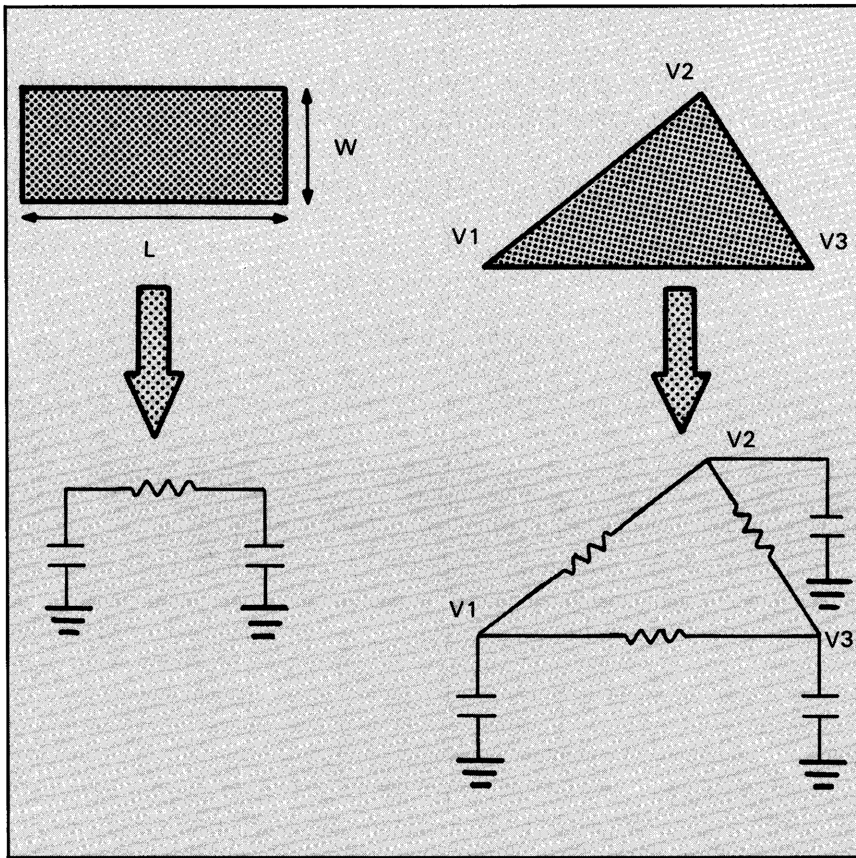

Figure 2. Circuit Representation of Rectangular and Triangular Elements

Determination of Current Waveforms

SPIDER requires transient current waveforms to load the extracted network at various points. A process has been developed to provide these waveforms. First, the designer indicates which groups of circuitry can be treated as a single block and the connection points of each block to the metal structures. Since it is usually not practical to include complete circuit blocks in the extracted SPICE deck, simplified models must be used to replace them.

Although any legal SPICE elements can be used, the most common model is a piece-wise linear current source that draws the same current that the circuitry would. A new program, CREST, is being developed to efficiently provide appropriate piece-wise linear current sources for CMOS circuitry. The next few sections briefly describe CREST; more detailed information is available from the authors and in Reference 3.

Basic Concepts

Current sources are required to model circuit operation under normal conditions (for electromigration) or worst-case conditions (for voltage drop). It is very difficult for the user to specify input waveforms for either condition without examining a variety of inputs. Unfor-

tunately, the number of input patterns is exponentially dependent on the number of inputs.

A SPICE analysis of the current drawn by a large circuit for a given input is feasible; however, performing this analysis for all possible inputs is prohibitively expensive. Any approach that avoids this pattern-dependence problem is called pattern-independent.

We present a new technique for deriving a pattern-independent estimate of the power and ground currents drawn by a CMOS circuit. To achieve pattern-independence, accurate information about the current is sacrificed and only statistical information is derived. Our goal is to produce an expected current waveform whose value at each time point is the expected or mean value of all values that the actual current can take at that time.

The expected current waveform is not the same as a time-average of the current; it is a waveform which, if the standard deviation of the actual current values is small, will approximately match the real current waveform(s). Under this condition, it can be used to estimate electromigration failure rates and voltage drop.

If the standard deviation is large, this waveform remains satisfactory for estimating electromigration median time to failure (MTF). Using the electromigration model in Reference 4, it can be shown that MTF is related either to the expected waveform of J (the current density) or that of J^2 (which we estimate, having found the expected waveform of J). This is the motivation and justification of our approach.

To derive the expected current waveform, E[i(t)], we build on the concept of signal probabilities,[5] which has recently become popular in the testing field.[6] Given that different input patterns may occur and that every pattern has a certain (user defined) probability of occurrence, then every input signal acquires a certain probability of being high, P_h.

Internal circuit node probabilities can be derived from the input probabilities, given the connectivity of the circuit. This concept has recently been used to estimate the power consumption of CMOS circuits.[7] Unfortunately, this approach is severely limited in application since all circuit inputs are assumed to switch randomly and at the same frequency. We make no such assumptions; we extend the signal probabilties concept to include transistion probabilities. The transition probability of a signal i at time t is the probability that it was low just before t and high just after t. Transitions probabilities at internal nodes can be derived from the transition probabilities at the input nodes, in a manner similar to the approach used with signal probabilities.

Signal probabilities are combined with transition probabilities to form probability waveforms. Probability waveforms are a series of transition edges with known transition probabilites at distinct points in time and constant signal probabilities between edges.

Each input to the circuit requires a probability waveform, which can be propagated throughout the circuit to obtain probability waveforms for each internal node. Probability waveforms can be generated from a series of logical input waveforms (Figure 3); this allows the designer to think in familiar terms.

The second innovation that allowed us to develop an efficient current estimation program was the concept of probabilistic device models. The simplest way to model a transistor is as a switch in series with a constant conductor; the transistor forms a conducting path between its source and drain when the switch is on.

Our probabilistic model is a modification of this. We represent a transistor as a conductor between source and drain. The value of the conductor is the chan-

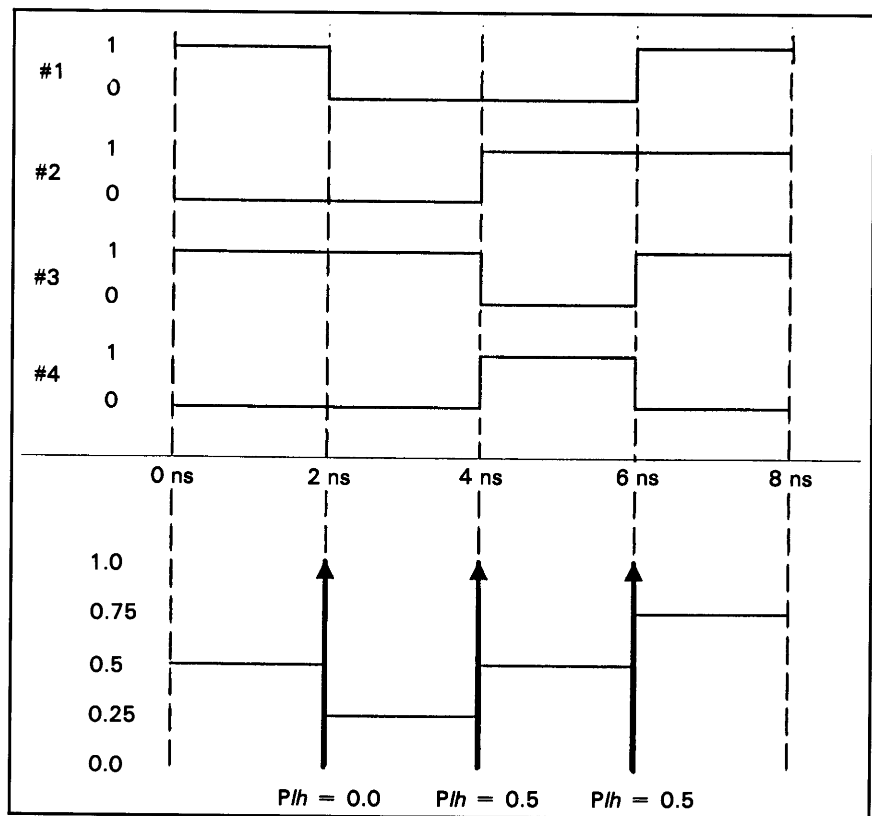

Figure 3. Four Waveforms with Equal Probability

nel conductance of the transistor. This model allows us to accurately estimate the expected current waveform drawn by the circuit.

For a standard logical gate, the current pulse is easily calculated. The expected conductance between the power (or ground) and output nodes of a gate can be determined by a simple combination of the expected conductances of all paths between them. This can be used to determine the expected peak output current, $E[pki_{out}]$, when the gate transitions. The expected charge at the output node, $E[q_{out}]$, is the capacitance times the probability of transition at the node.

If a simple triangular current pulse is chosen with area $E[q_{out}]$ and height $E[pki_{out}]$, then a width (delay) τ can be estimated for the pulse. The final current pulse that we use to approximate the total expected current drawn by the gate, $E[i_{tot}]$, is a pulse with area $E[q_{tot}]$ and width τ (Figure 4), where $E[q_{tot}]$ is the expected charge drawn by all nodes in the gate.

Probabilistic Simulation for Current Estimation

These concepts form the basis for CREST, an efficient simulator for estimating the expected current waveform drawn by CMOS circuitry under normal operating conditions. In this section, we describe CREST's input requirement, our overall simulation approach, and the modifications to the basic simulator required to accurately handle spatial dependency and feedback.

CREST needs three types of input. First, CREST requires a SPICE deck of the circuit whose current waveforms are to be estimated. Second, a list of all nodes that are connections to power and a list of all nodes that are connections to ground must be provided. Finally, CREST requires input waveforms for all circuit inputs. If more than one input waveform is present for a node, then each waveform must be accompanied by the probability that will occur. This allows CREST to create a statistically meaningful expected current waveform.

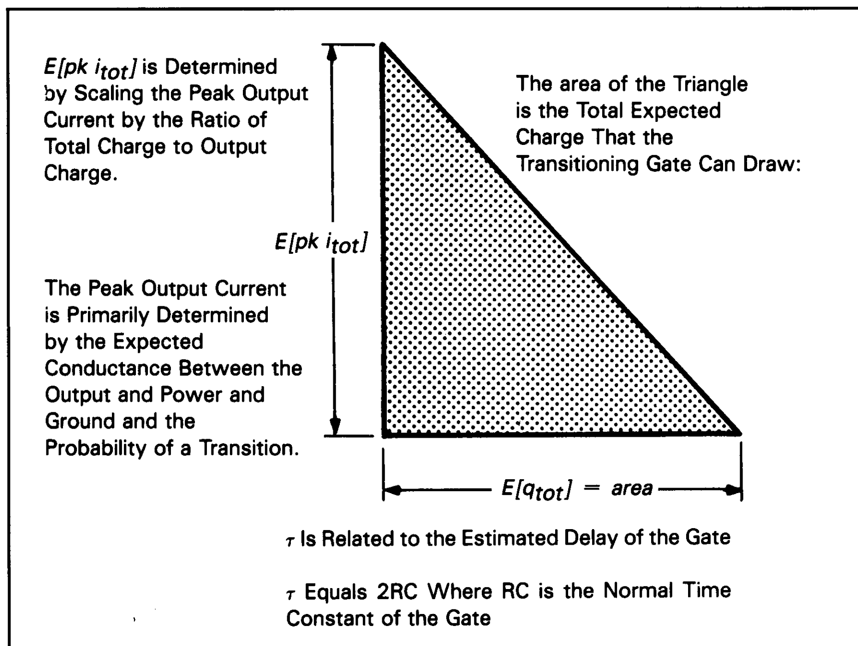

Figure 4. For Simplicity, a Simple Triangular Pulse is Being Used to Estimate the Current Pulse Drawn by a Gate when It Transitions

CREST is now ready to estimate the current waveforms of the circuit.

First, data structures are created from the SPICE input. Each MOSFET is represented with a probabilistic device model. A node structure is created for each electrical node which contains the connectivity and capacitance information necessary for current estimation.

Next, a simple partitioner joins together all elements with common source/drain nodes that are not power of ground nodes, to form gates. The circuit is now represented by a network of gates. Standard logic gates are identified and marked. Gates that are not standard are broken up into standard gates and pass transistor gates.

Once the partitioned circuit has been constructed, it must be simulated over a time interval specified by the user to estimate the expected current waveform. Current estimation, by the methods we propose, is similar to logic and timing simulation, and we borrow the event-driven simulation approach from there.

Each time there is a non-zero change of a transition on a node, an event is created. These events are kept in an event queue, and store the transition time, the probability of transition, and a pointer to the affected node. When an event occurs, all gates that have the event's node as an input are examined. If there is any probability of transition on the output of the gate due to the event, the new transition and signal probabilities are calulated by formulas that vary depending on the gate type.

If the gate is a standard gate, a current pulse is also generated in the manner described above. A new event is then added to the event queue at the current time plus the gate delay, indicating the transition on the output node. The process is repeated until the event queue is empty of events that fall within the specified simulation interval.

This approach assumes that the inputs to gates are spatially independent (the signal probability at each input is unrelated to the signal probability at all

others). Reconvergent fanout and feedback can cause these inputs to be spatially dependent. Spatially dependent sections of the circuit must be split into supergates, whose inputs are independent. These supergates are then simulated for every allowed combination of transitions on the nodes that cause the spatial dependency.

Although this can be simplified somewhat, the complexity is still exponential. Various heuristics are being investigated in search of an acceptable compromise between accuracy and simulation simplicity for supergates.

Violation Evaluation and Correction

Current Density Evaluation and Correction

For dc current, the relationship between MTF and current density can be determined with reasonable accuracy by accelerated life testing at elevated temperature and high current density.

MTF is a highly nonlinear function of both temperature and current density. Methods of determining MTF under operating conditions by extrapolating from accelerated life test data have been summarized by McPherson and Ghate,[4] who have also derived a procedure for converting a transient current density waveform to a dc current density which is equivalent for electromigration purposes. The user of SPIDER specifies a dc current density which will produce the desired MTF, and the program uses the McPherson and Ghate forms, as calculated from SPICE current waveforms, with the specified value.

The comparison is performed for each extracted element. Following the approach of McPherson and Ghate, the piecewise-linear current density in each resistive element is converted to a series of pulses having amplitude equal to the short-time current density J_{rms} calculated for each time interval, assuming:

$$MTF \alpha (J_{rms})^n \quad (1)$$

where

$$n = \begin{cases} 1.0 & \text{for } J_{rms} \leq 10^5 A/cm^2 \\ 1.5 & \text{for } 10^5 < J_{rms} < 10^6 A/cm^2 \\ 2.0 & \text{for } J_{rms} \geq 10^6 A/cm^2. \end{cases} \quad (2)$$

The result is the expression shown in Equation Box A (Equation 3).

Rather than trying to estimate the actual MTF, SPIDER performs a complex comparison of the equivalent current density in the element to the "acceptable" current density supplied by the user. If an excesive current density is identified in a two-terminal rectangular region, the line width is widened to correct the violation.

No attempt is made to correct excessive current density in triangular elements, although the violation is reported to the user. If any line widths are modified, then the SPICE simulation is redone (Figure 1). Any modified line widths and uncorrected violations are reported to the user after all iterations are complete.

There are two options available for current density correction. In the first and most computationally efficient option, it is assumed that width is increased when current density exceeds the specified value, but no change is made to regions where current density is below the specified value.

In the second option, width is reduced, if required, to a minimum width dictated by layout rules. This second option is useful in the initial design phase to minimize metallization area, but is not recommended for analysis of large circuits, since it can cause a large number of SPICE iterations and, for some large circuits, has exhibited convergence problems.

Voltage Drop Evaluation and Correction

SPICE output provides the voltage at each node as a function of time. The program compares the maximum voltage difference between each node and a reference, such as the input bond pad. A list of all violations of user specifications is produced.

Correction of voltage drop violations is more complicated than current density correction. An approach was devised which would provide the designer with recommendations for altering the widths. This approach required the use of a circuit optimization program. This automatic correction option is not currently supported. It was decided that it was easier and more efficient to report violations to the designer and allow him to alter the design to meet voltage drop specifications.

Output Formats

Four different types of output are available from SPIDER. In the first form, a list of recommended line width changes is reported to the user (Figure 5) as well as a list of elements which violated current density restrictions, but did not fit the sizing algorithms of SPIDER.

In the second form, the extracted geometries are plotted on a color plotter. The color of each geometry indicates the current density in the element; both logarithmic and linear scales are available in this mode.

The third output form identifies which nodes exceed the voltage drop specifications; however, a more useful way to view voltage drop violations is with a color plot where each node is colored to indicate the magnitude of the voltage drop.

Status

All previous descriptions have been for the system as it will exist when completed. This project is still in the development phase; however, some sections are nearly completed and can be of practical use.

The system flow, presented in the second section, can be broken down into three parts: SPIDER, CREST, and interfaces. The first prototype of SPIDER is complete and functioning; however, significant reductions in the number of

three-terminal finite elements are required to make it viable for large bars. Several promising algorithms are now being evaluated.

The first prototype of CREST is also complete. Extensive testing is just beginning and the results of these tests will determine what extensions or changes might be required. The interfaces between SPICE, SPIDER, and CREST are primarily in the planning stage. Many important steps still require significant user effort.

Results

In this section, we will present results from SPIDER and CREST. Since they are being independently developed, no tests have been conducted on the overall system. However, results for both SPIDER and CREST include examples from real circuits contained in TI products.

Figure 5 shows a simplified test layout. The original line widths and the new line widths recommended by SPIDER

Equation Box A

$$\beta = \left[\left(\frac{10^5}{J_s}\right)^{1/2} \sum_k \left(\frac{J_k}{J_s}\right) \left(\frac{t_k}{t_0}\right) + \sum_l \left(\frac{J_l}{J_s}\right)^{1.5} \left(\frac{t_l}{t_0}\right) + \left(\frac{J_s}{10^6}\right)^{1/2} \sum_m \left(\frac{J_m}{J_s}\right)^2 \left(\frac{t_m}{t_0}\right) \right]^{-1} \quad (3)$$

where β is the ratio of the MTF which would be observed under transient conditions to the MTF which would be observed at the dc current density J_s.

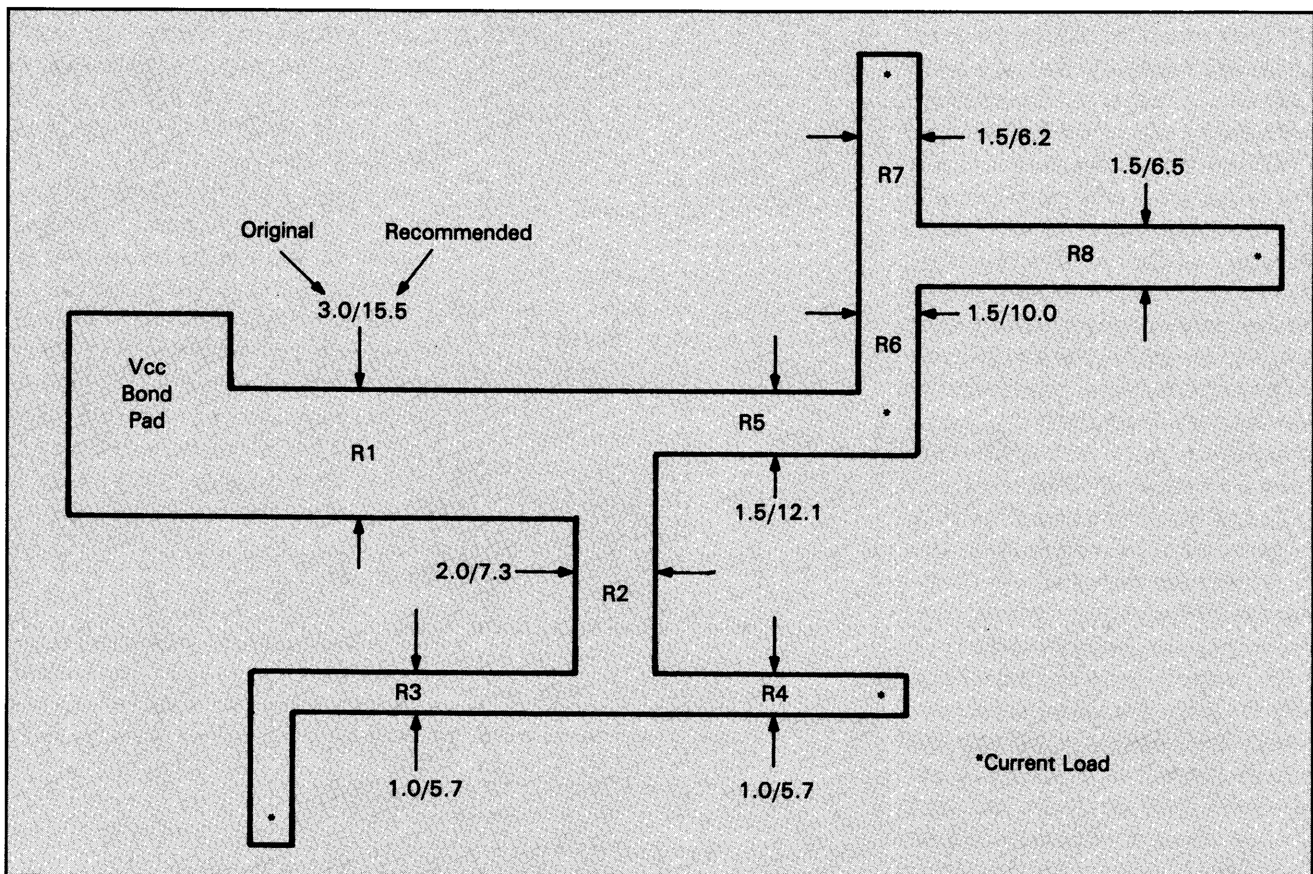

Figure 5. Simple SPIDER Example

to produce equivalent dc current densities of 2×10^5 A/cm^2 are shown in the figure.

As another test, a 300 node, 150 resistor network that represented the power distribution network of a 4-megabit DRAM design[8] was analyzed. When the current density specifications were set artificially low to produce violations, SPIDER converged to a solution in three iterations, provided the options were set for width increases only.

If width reductions were allowed, and the current density specification was set to a realistic value, all of the resistors were narrowed by a substantial amount. In this case SPIDER did not converge, even after ten iterations. This result is an indication that the width reduction option is appropriate only in the design mode.

An additional example is illustrated by Figure 6. The power and ground lines for a circuit were extracted using only two-terminal rectangular elements. The circles represent the current generator connection points specified by the designer. Current waveforms were generated from the various subcircuits using SPICE.

The circuit as shown did not exhibit any electromigration problems; however, when the electromigration specifications were reduced to one-tenth of normal, SPIDER recommended widening some of the lines by the factors shown. Further examination of the layout showed that these sections provided current to more than one output driver circuit, and therefore carry more current than other sections.

The first prototype of CREST has been completed. CREST's goal is to approximate the expected pulse drawn by the circuit for known signal and transition probabilities on the inputs of the circuit. For comparison, we generated the expected current pulse for a variety of examples by running SPICE on every input voltage waveform allowed by the probability vectors, weighting each pulse by the probability that the waveform producing it would occur, and summing the weighted pulses to produce the expected pulse.

In all examples tested the results were excellent. Peak currents were within 20 percent, average currents were within 15 percent, and timing estimates were within ten percent of SPICE. Figure 7 shows a typical example.

To be used to estimate current of today's large VLSI designs, this approach must be faster than SPICE. Table I compares CREST simulation speed to SPICE simulation speed for all vectors necessary to generate the expected pulse, and the SPICE simulation speed divided by the number of vectors.

In overall simulation speed CREST is dramatically faster than SPICE, due to the exponential number of vectors that must be simulated for multiple inputs. CREST simulation speed is still significantly faster than the SPICE analysis time per vector, and the trend we observed was for the speed gain of CREST over SPICE to grow with circuit complexity (Figure 8).

Circuits containing spatial dependency were also examined and yielded similar results; however, space constraints prevent their presentation. Some of the larger examples cannot be analyzed for probabilistic inputs since their supergates are so large. Several promising heuristics are being evaluated which offer a solution to this problem.

As mentioned in the introduction, CREST will run perfectly well for logical waveform inputs, rather than probabilistic inputs. Although nothing has been done to optimize for logical inputs (and there is a potential for dramatic increases), CREST is much faster on large examples with significant spatial dependency. This is evidenced by the last entry in Table I.

Conclusion

The main focus of this system will be to diagnose problems, and in some cases, suggest solutions. This system depends heavily on the designer for correct input. The most important decisions by the designer are how to block up the circuitry and which input waveforms should be considered.

An efficient, accurate, and easy to use system of programs is being developed to allow designers to ensure that their designs will meet the voltage drop requirements and current density guidelines to assure the device reliability goals (e.g., 100K op hours with a cumulative failure rate of less than 1 percent). Although the system is not completely ready for use, most of the major technical hurdles have been solved, and the prospects of having a system that fulfills our needs is excellent. The system is scheduled for completion by the end of 1988.

References

1. J.E. Hall, et al., "SPIDER — A CAD System for Modeling VLSI Metallization Patterns," *IEEE Transactions on Computer Aided Design*, Vol. CAD-6, November 1987, pp. 1023-1031.
2. L.W. Nagel, "SPICE2: A Computer Program to Simulate Semiconductor Circuits," *Ph.D. Thesis*, Department of Electrical Engineering, University of California, Berkeley, 1975.
3. R. Burch, et al., "Pattern-Independent Current Estimation for Reliability Analysis of CMOS Circuits," *Proceedings of the 25th Design Automation Conference*, Anaheim, Calif., June 1988.
4. J.W. McPherson and P.B. Ghate, *Proc. Symp. on Electromigration of Metals and First Int. Symp. on Multilevel Metallization and Packaging*, the Electrochemical Soc. Inc. 85.6, 64, 1985.
5. K.P. Parket and E.J. McCluskey, "Probabilistic Treatment of General Combinational Networks," *IEEE Transactions on Computers*, *June 1975, pp. 668-670.
6. S.C. Seth, et al., "PREDICT — Probabilistic Estimation of Digital Circuit Testability," *IEEE 15th Annual International Symposium on Fault-Tolerant Computing*, Ann Arbor, Mich., June 1985, pp. 220-225.
7. M.A. Cirit, "Estimating Dynamic Power Consumption of CMOS Circuits," *IEEE International Conference on Computer Aided Design*, Santa Clara, Calif., November 1987, pp. 534-537.
8. A.H. Shah, et al., "A 4Mb DRAM with Cross-Point Trench Transistor Cell," *IEEE Transactions on Computer-Aided Design*, Vol. CAD-6, No. 3, May 1987, pp. 337-345.

Table I. Execution Time Comparisons of CREST and SPICE. All times in seconds of CPU time.

Circuit		CREST		SPICE		
Name	Size #MOSFETs	Total Execution Time	Analysis Time	Total Execution Time	Analysis Time	Analysis Time per Vector
Decode 1	20	1.1	0.26	642	625	39
Decode 2	28	1.4	0.35	989	965	60
Decode 3	36	1.9	0.50	1588	1565	98
Invt10	20	1.0	0.22	221	209	52
Invt40	80	3.3	0.92	2094	2057	514
Bridge	10	0.8	0.20	23237	22049	22
34bit alu*	1839	147	73	83691	82042	82042

* The 34bit alu could not be tested for probabilistic waveforms since its supergates were too large for exact analysis. It was tested for a single set of input vectors.

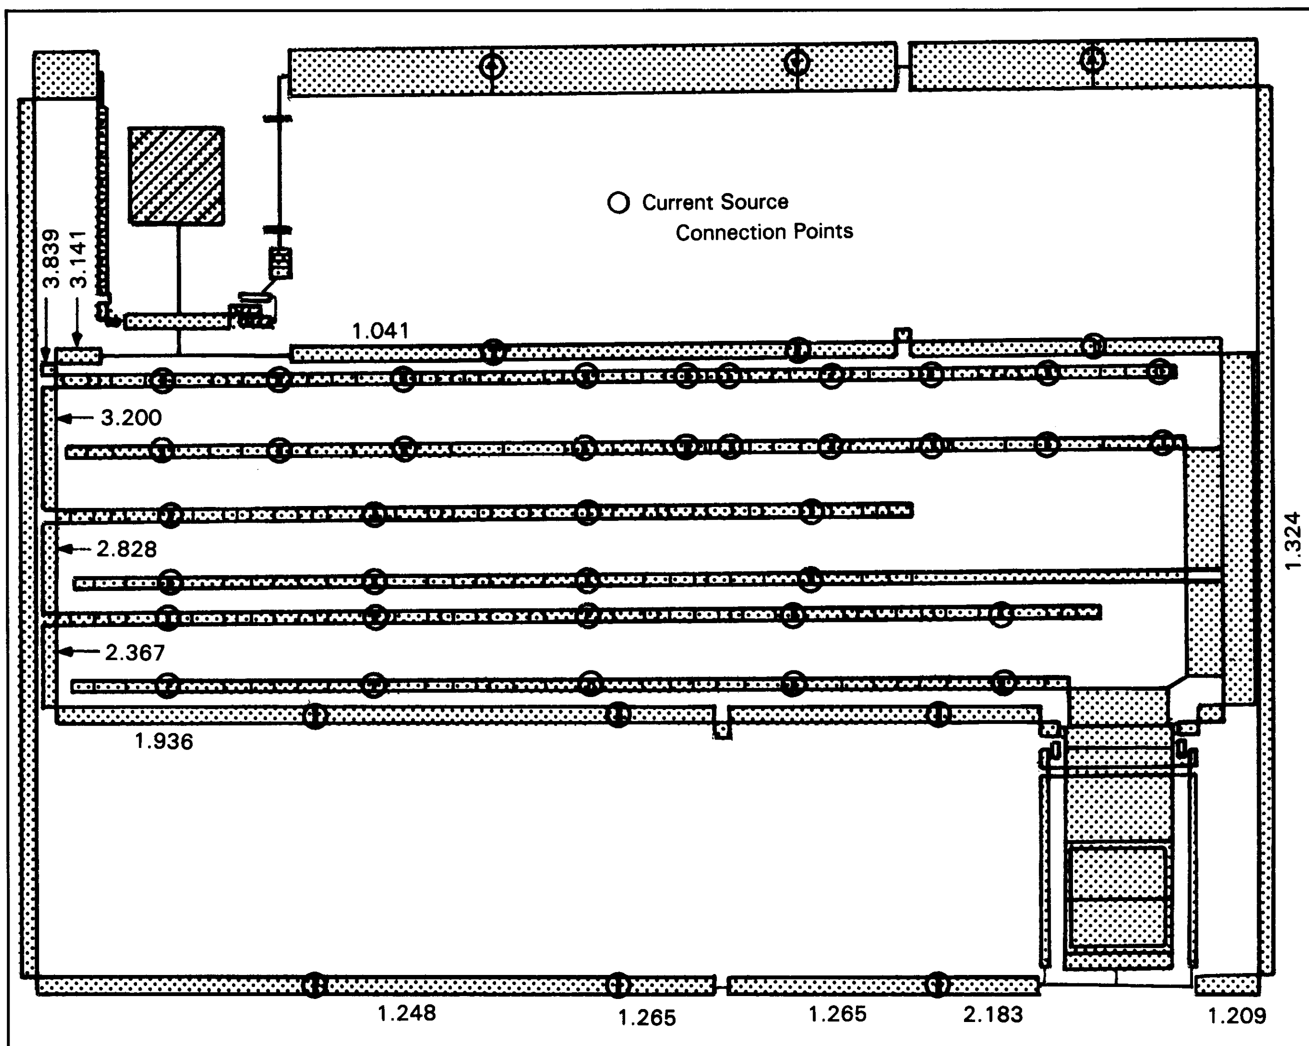

Figure 6. CREST Gain Over SPICE Analysis/Vector

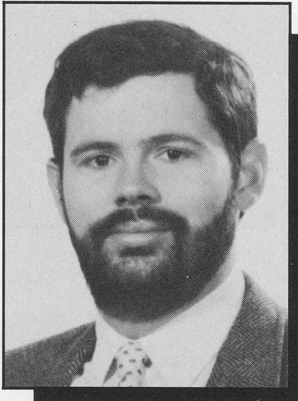

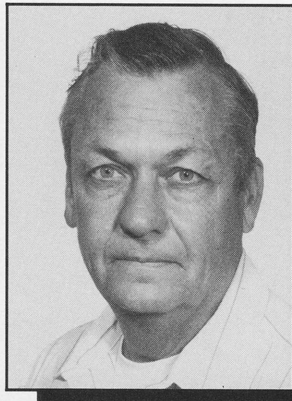

Richard Burch

Richard Burch received his B.S. in physics and his M.S.E.E. from the University of Oklahoma in 1984 and 1986, respectively. He joined TI as a summer student in 1984, and when he was not in school, he was involved in circuit simulations research at the Semiconductor Process and Design Center (SPDC).

After receiving his M.S., Richard joined SPDC as a Member of the Technical Staff, where he continued work on circuit simulation. In June 1987, Richard was given responsibility for SPDC's future contributions to the SPIDER project, including the development of CREST.

Joseph Hall

Joe Hall received his B.S.E.E., M.S.E.E., and Ph.D. degrees from Southern Methodist University. He has been with Texas Instruments since 1957, and he is currently a Senior Member of the Technical Staff in the Design Automation Division.

His most recent work includes modeling and analysis of latch-up and electrostatic discharge phenomena in CMOS integrated circuits, and development of a computerized system for current density evaluation in integrated circuit metallization.

From 1981 to 1983, Joe was an associate professor of electrical engineering at Texas A&M University. His research there included integrated circuit design as well as development of solid-state oxygen sensors. Joe is a Registered Professional Engineer in Texas.

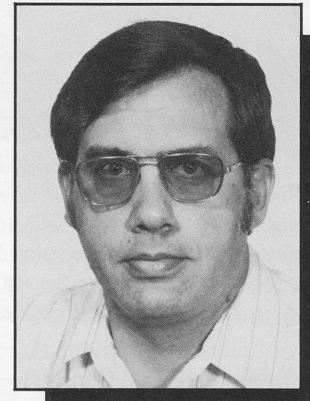

Michael McGraw

Michael McGraw received his A.B. degree in mathematics from the University of California at Berkeley in 1969.

He served in the United States Navy from 1971 to 1974, stationed in Vietnam, Long Beach, and San Diego. He joined the Design Automation Department of Texas Instruments in 1976, and has worked mainly in the areas of design verification and extraction of schematics and parameters from the layout.

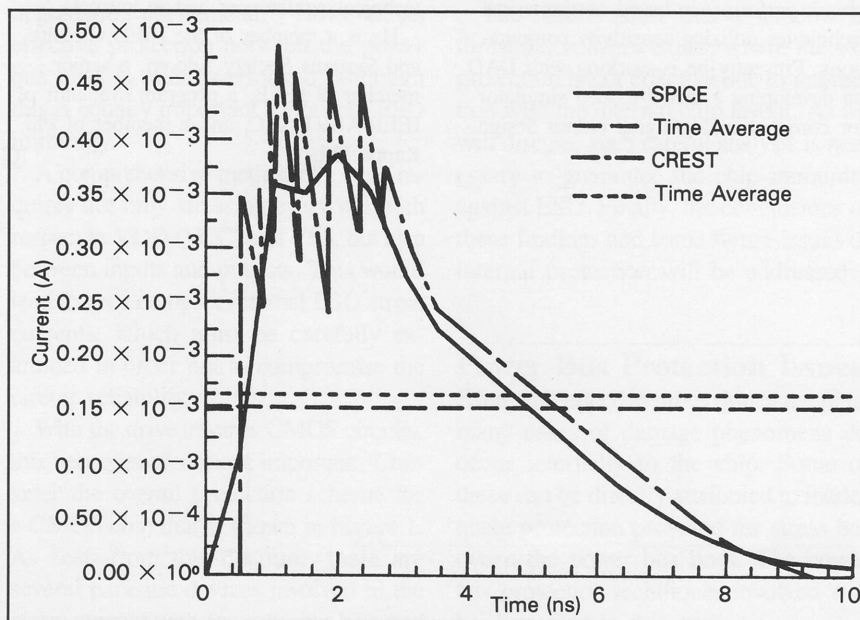

Figure 7. CREST Over SPICE Analysis/Vector

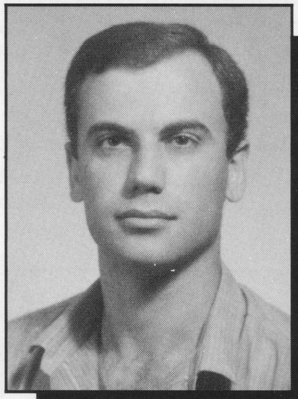

Farid Najm

Farid N. Najm received his B.S. in electrical engineering from the American University in Beirut with distinction in 1983, and his M.S.E.E. from the University of Illinois at Urbana-Champagne in 1986.

He worked as a Member of the Technical Staff at TI from 1987 to 1988. Currently, he is working towards his Ph.D. in electrical engineering at the University of Illinois.

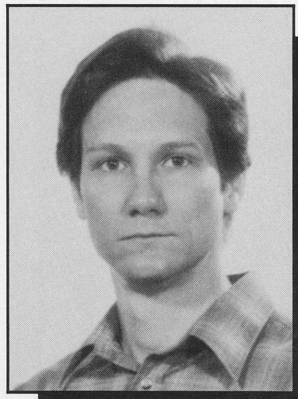

Dale Hocevar

Dale Hocevar received his B.S. degree in electrical engineering from the University of Tulsa in 1978, and his M.S. and Ph.D. degrees from the University of Illinois in 1980 and 1982, respectively.

In 1983, he joined the VLSI Design Laboratory of the Semiconductor Process and Design Center as a Member of the Technical Staff. He has worked on statistical optimization techniques for MOS circuit performance, developed algorithms for computation of transient sensitivities in SPICE, worked on accelerating SPICE via parallel algorithms and special hardware, and was involved in developing OASYS, a tool for optimizing circuit performance based on numerical techniques utilizing sensitivity computations. Presently he is working with DAD on developing a mixed-moded simulator for complex analog/digital circuit design.

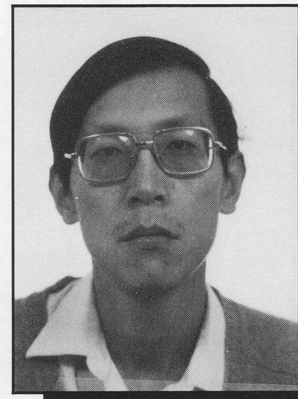

Ping Yang

Ping Yang received his B.S. degree in phyics in 1974 from the National Taiwan University, and his M.S. and Ph.D. degrees in electrical engineering from the University of Illinois at Champagne-Urbana in 1978 and 1980, respectively.

He joined the Central Research Laboratory at TI in August 1980. He is currently a TI Fellow and manager of the Modeling, Analysis, and Simulation branch in the VLSI Design Laboratory of the Semiconductor Process and Design Center.

Ping has co-authored a book, and has presented or published more than 80 articles at international scientific and technical conferences and in journals.

He is a member of the IEEE Circuits and Systems Society Adcom, a senior member of IEEE, a program evaluator of IEEE-ABET/AEC, and a member of Phi Kappa Phi.