Chapter 6. Custom Integrated Circuits

Academic and Research Staff

Professor Jonathan Allen, Professor John L. Wyatt, Jr., Professor Srinivas Devadas, Professor Jacob White, Professor Dimitri A. Antoniadis, Professor Berthold K.P. Horn, Professor Hae-Seung Lee, Professor Bruce R. Musicus, Professor Tomaso Poggio, Professor Charles G. Sodini

Graduate Students

Robert C. Armstrong, Donald G. Baltus, Cyrus S. Bamji, Michael J. Bryan, Curtis S. Chen, Steven J. Decker, Ibrahim M. Elfadel, Mikko Hakkarainen, Craig Keast, Kevin Lam, Jennifer A. Lloyd, Andrew Lumsdaine, Steven P. McCormick, Ignacio McQuirk, Keith S. Nabors, Joel R. Phillips, Mark W. Reichelt, Mark N. Seidel, Luis M. Silviera, David L. Standley, Ricardo Telichevsky, Christopher B. Umminger, Filip Van Aelten, Woodward Yang

Technical and Support Staff

Dorothy A. Fleischer, Susan E. Nelson

6.1 Custom Integrated Circuits

Sponsors

Analog Devices, Inc.
International Business Machines Corporation
Joint Services Electronics Program
Contract DAAL03-89-C-0001
U.S. Air Force - Office of Scientific Research
Contract AFOSR 86-0164B

Project Staff

Professor Jonathan Allen, Robert C. Armstrong, Donald G. Baltus, Cyrus S. Bamji, Steven P. McCormick, Mark W. Reichelt, Filip Van Aelten

Through our VLSI CAD research, we seek the means to produce custom integrated circuits correctly, quickly, and economically. Traditionally, correctness has been checked at several representational levels such as layout (via design rule checking), circuit, and logic (both via simulation). The techniques for checking correctness are local to the particular representational level involved. While these techniques are important components of design testing, they do not attempt to provide alignment and consistency checks between the different levels of representation and an input behavioral specification. In addition, they do not characterize the set of possible designs at each representational level corresponding to the initial functional specification, while ranging over a variety of performance levels. For this reason, there is an increasing need to provide CAD tools that can serve as the framework for design exploration, providing the desired performance together with consistently aligned representations at all levels.

In this research group, we are studying a variety of research topics with an emphasis on performance directed synthesis of custom VLSI designs. Professor Allen has recently provided an overview of these issues,[1] in which he emphasizes the need for coordinating the design optimization process over the several levels of representation.

Since there is so much to be gained by aggressive circuit design, our research group has placed considerable emphasis on this level. In addition, the group has focused on optimized architectures for digital signal processing (1) because of the highly parallel nature of the algorithms involved, and (2) because of the need for a very high level of real time performance. We have also increased our emphasis on developing means for formally specifying designs that facilitate performance directed design exploration, as well as efficient verification of correctness of the coordinated set of design level represent-

[1] J. Allen, "Performance Directed Synthesis of VLSI Systems," *Proc. IEEE* 78(2): 366-355 (1990).

ations. All of these studies are taking on an increasingly theoretical approach, signifying the transition of digital system design from an art to a science.

The circuit level of design representation is the most abstract level at which performance can be characterized explicitly in terms of delay, and also where performance can be substantially manipulated through changes in circuit style and device sizing. It is an important level of design because optimization techniques which manipulate circuit representation can have great impact on resulting performance. More abstract representations can be transformed into a circuit representation, but the need to transform a circuit into layout in a general yet efficient way is needed for the rapid design of custom circuits. Baltus[2] has developed a very general yet efficient technique for compiling circuits into layout. The algorithm converts both NMOS and CMOS circuits in any style, and provides special means for the efficient layout of circuits with highly varying transistor width-to-length ratios. While initially focused on implementation technologies with only one level of metal, it is now being generalized to multiple levels of metal. Recent comparisons of this program with other techniques at a benchmark competition have demonstrated the effectiveness of these procedures.

Now that there is assurance of the generation of a high quality layout from any circuit specification, we are turning our attention to design performance exploration at the architectural level. For this research, our goal is to study the performance of quasi-regular architectures through manipulations of an initial behavioral specification. A new language has been developed for algorithms which can be compiled into either regular systolic architectures, or structures with a limited degree of irregularity. From the input language specification, a prototype descriptor is generated which compactly and efficiently represents both the structural hierarchy and the degree of replication of architectural elements. This descriptor can be readily expanded into a fully parallel data flow graph if necessary, but has the advantage that its complexity does not appreciably grow with circuit element replication.

After obtaining the prototype descriptor, we apply a series of three types of transformations in order to explore the architectural design space. First, various amounts of serialization can be instantiated, cutting down the amount of inherent parallelism in the input language specification, but also reducing the area of the resultant implementation. Next, we explore a variety of architectural transformations that map the index spaces associated with each array variable into a common global index space. Finally, we explore pipelining to localize communication between architectural elements. This approach is expected to lead to a flexible but powerful framework for architectural design exploration in the context of known capability for the conversion of the circuit form of the architectural elements into efficient layout. One important research issue is the degree of irregularity that is desirable in the resultant arrays due to non-local communication. It may be possible to make the distance of communication a designer-controlled parameter, thus providing a means to depart from regularity when the communication needs of the algorithm favor it, thereby expanding the class of designs that can be efficiently converted to layout.

Another example of the mapping of a high level specification to the circuit level representation has focused on the technology mapping process. Much research has focused on the conversion of a high-level functional specification to a highly optimized logic specification. This logic representation is usually then converted to the circuit form by means of a technology mapping in terms of a cell library. Heretofore the cell library was made up of typical restoring gates, such as

---

[2] D.G. Baltus and J. Allen, "SOLO: A Generator of Efficient Layouts from Optimized MOS Circuit Schematics," *Proceedings of the 25th Design Automation Conference,* June 1988.

NAND and NOR cells, but Kaplan[3] has recently expanded this process to include pass transistor gates as well as restoring forms. By broadening the class of circuit forms, and hence basic library gate cells, new flexibility is introduced into the technology mapping process. New logic-level representations have been introduced to facilitate the optimization of these circuits. By coupling to the existing SOLO layout generator, savings in the area of approximately 20 percent have been demonstrated for a number of designs that take advantage of the use of pass transistor gates as well as restoring gates. This research exemplifies the need to coordinate design at several levels. Logic synthesis must contemplate the class of library cells that will be used in implementations, while increased circuit flexibility leads to new layout savings by means of a more powerful technology mapping.

Driven by the requirements of performance optimization, we have devoted a considerable amount of research to delay characterization associated with circuit waveform transitions. McCormick[4] has developed an efficient scheme for computing these transitions accurately using a small number (four or less) of circuit transform moments. Interactions between interconnection lines, even as general as RLC lines, can be readily computed with these moments. In addition to passive interconnect, we have expanded these techniques to also include the efficient characterization of active circuits. Using these techniques, efficiency of delay calculation, together with accuracy, is obtained, thus avoiding the substantial numerical load of comprehensive circuit simulation. Moment calculations are thus part of an overall strategy of performance estimation that must be done quickly and accurately in order to facilitate rapid exploration of design alternatives. It is also important to note that delay due to interconnect is now dominating the delay caused by active circuits in large systems. The ability to accurately characterize circuit delay in complex branching interconnect structures, including the effect of interconnect material together with driving and terminating impedances, is essential to the accurate exploration of circuit performance.

With the advent of very small devices operating at very high speeds, we need to revise device models for circuit simulation. Rather than use average value parameters, we must now solve the device equations directly to provide the accuracy required. Unfortunately, direct use of these equations implies very long simulation run times, making it important to find ways of obtaining required accuracy with reduced computational expenditure. Reichelt[5] has applied waveform relaxation techniques to this problem in a way that readily allows for the use of parallel waveform relaxation simulators. Initial results have proved the effectiveness of this approach, which guarantees convergence of the simulation. Investigations of possible parallel architectures for this purpose are now under way.

Recently, our emphasis on verification techniques has grown as the complexity of designs, and the need for overall correctness, has increased. There are a variety of approaches to verification, including logical proofs that implementation structures correspond to behavior specifications, and symbolic simulation. Because the need for formal characterizations of design has grown, we are searching for representations that can compactly characterize the class of correct

---

[3] J.T. Kaplan, *Pass and Restoring Logic Technology Mapping for the Efficient Automated Generation of Combinational Leaf Cells,* S.M. thesis, Dept. of Electr. Eng. and Comput. Sci., MIT, 1989.

[4] S.P. McCormick and J. Allen, "Waveform Moment Methods for Improved Interconnection Analysis." *Proceedings of the 27th Design Automation Conference,* Orlando, Florida, June 1990.

[5] M. Reichelt, J. White, and J. Allen, "Waveform Relaxation for Transient Simulation of Two-Dimensional MOS Devices." *Proceedings of the International Conference on Computer Aided Design,* November 1989.

designs. Bamji[6] has applied context free formal grammars to the specifications for the class of MOS digital integrated circuits. Compact grammars can represent classes of infinite cardinality, so that, for example, the class of all CMOS static combinatorial circuits can be generated by only five grammatical rules. Many circuit styles can be characterized in this way, and the corresponding parsing procedures can deal with mixed styles in an efficient way. Bamji has represented the class of all layouts generated by a regular grammar by use of the Regular Structure Generator, and he has shown how to verify the correctness of the resulting layout by parsing techniques. Finally, both of these grammatical representations, one at the circuit level and the other at the layout level, have been combined into a coupled grammar that can be used to verify the proper correspondence between layout and circuit. Thus, not only can both layout and circuit structures be verified independently by similar grammatical techniques, but we can show that their alignment is correct so that they are both coordinated representations of the same underlying abstract circuit in a consistent fashion. Successful verification demonstrates that all circuit and layout structures have correspondents at related layout and circuit representations, respectively, so that no loose fragments at either level are isolated from those at other design levels. Since these grammatical techniques have been successful at the circuit and layout levels, we are now focusing on their application at higher levels of design.

Grammatical techniques are appropriate for characterizing correct structure of circuits, but they are not easily adapted to the verification of nonstructural circuit properties such as device sizing and allowable capacitance sizes that could lead to undesirable charge sharing. To deal with these attributes, Van Aelten[7] has developed an efficient technique for verification of circuit properties. This approach builds on the grammatical techniques of Bamji, parsing the circuit first to reveal its constituent structure. Semantic checks are then applied to the circuit to verify its specified performance. The circuit is effectively mapped into logic level representations by a process of recursive interpretation. Checks for electrical constraints are made at the level of the individual subcircuits, and the behavioral description of these elements is composed in order to provide a match with the user supplied specification. Extensions to this approach require more comprehensive circuit grammars that may have to be interleaved with the semantic interpretation process, thus breaking down the traditional separation of syntax and semantics.

To continually maintain a consistent set of design representations that are properly aligned, it is important to build a unified framework for all levels of representation. Armstrong[8] has developed a technique for specifying design at all levels of abstraction with a uniform formalism. The way in which each of these abstractions is a consistent projection of one underlying abstract design entity is explicitly described, thus facilitating the automatic alignment of these representations. In this way, when one design projection at a specific level of representation (e.g., circuit level) is manipulated, all other levels are changed automatically to provide the correct alignment. In general, this is a one-many process, since there is usually an infinite number of less abstract design projections corresponding to a more abstract projection of the same overall design. All that is required, however, is that all levels are consistent, so that performance optimization algorithms can select the best design projection at a given level, subject to some

---

[6] C.S. Bamji and J. Allen, "GRASP: A Grammar-based Schematic Parser." *Proceedings of the 26th Design Automation Conference,* June 1989.

[7] F. Van Aelten and J. Allen, "Efficient Verification of VLSI Circuits Based on Syntax and Denotational Semantics," *Proceedings, IMEC/IFIP Workshop on Applied Formal Methods for Correct VLSI Design,* Houthalen, Belgium, 1989.

[8] R.C. Armstrong, *A Formal Approach to Incremental Consistency Maintainence in Multirepresentation VLSI Databases,* Ph.D. diss. proposal, Dept. of Electr. Eng. and Comput. Sci., MIT, 1987.

overall optimization criterion. Upon completion, this new formal framework will provide a basis for interconnecting an entire suite of design tools and will easily contemplate the addition of new levels of representation without the need to provide an entirely new framework.

The above descriptions show how several aspects of performance-driven synthesis are being studied and adapted to new CAD tools. There is a clear trend toward adopting more formal notions of behavior, structure, and attributes, together with techniques for coordinating these several levels of description in one overall consistently maintained structure. In the future, we will emphasize completion of the theoretical framework and unification of optimization procedures across all levels of design.

6.2 The Vision Chip Project

Sponsors

DuPont Corporation
National Science Foundation
Grant MIP 88-14612

Project Staff

Professor John L. Wyatt, Jr., Professor Berthold K.P. Horn, Professor Hae-Seung Lee, Professor Tomaso Poggio, Professor Charles G. Sodini, Steven J. Decker, Ibrahim M. Elfadel, Mikko Hakkarainen, Craig Keast, Ignacio McQuirk, Mark N. Seidel, David L. Standley, Christopher B. Umminger, Woodward Yang

6.2.1 Overall Project Description

Problem Definition and Methods

Computational Demands of Vision and the Smart Sensor Paradigm: A major problem in machine vision is the extraordinary volume of input data that must be acquired, managed and processed. This large quantity of data in real-time grey-level images leads to communication bottlenecks between imager, memory and processors, while the computational demands remain high once data have reached the processors. The result is that conventional machines are incapable of any but the most rudimentary forms of real-time vision processing.

The *smart sensor paradigm* shows great promise as a new approach to this problem. Although potentially applicable to a variety of sensor types, e.g., tactile and acoustic arrays, in this project we are restricting its application to vision. The key idea is to incorporate signal processing as early as possible into a system's signal flow path to reduce demands for transmission bandwidth and subsequent computation.

In this paradigm, the major themes are:

- sensors and processing circuitry integrated on a single chip,
- parallel computation,
- processing circuits distributed throughout the array for spatially local operations in situ with minimal wiring overhead,
- use of analog circuits for early processing tasks to avoid the speed and area costs of analog-to-digital conversion on high-bandwidth data streams,
- selection of tasks and algorithms requiring low to moderate precision, and
- special emphasis on computations that map naturally to physical processes in silicon, e.g., to relaxation processes or to resistive grids.

Analog Integrated Circuits: Both continuous-time and sampled-data analog integrated circuits play a key role in our research. Analog circuits offer substantial advantages in speed, functionality per unit silicon area, and ease of interfacing to inherently analog input data, while digital circuits are easier to design and can obtain unlimited precision through use of progressively longer wordlengths.

The Gilbert 2-quadrant multiplier is a dramatic example of analog capability. Requiring only four transistors, its similar, although slightly more complex circuits, achieve 0.02 percent precision multiplication in under 200 nanoseconds with a dynamic

range of about five orders of magnitude![9] More generally, analog circuits that perform basic operations such as addition, subtraction, multiplication, thresholding, absolute value, logarithm, exponentiation, time integration and simple linear filtering are much more compact and, in most instances, faster than corresponding digital implementations. These circuits can be fabricated in array form on the same chip as transducers and operate directly on transducer outputs.

The trade-off is in the inevitable drift and imprecision of analog integrated circuits, which require careful design. Although these problems make analog methods inappropriate for certain tasks, they are not major drawbacks in early vision applications, in which the input data (pixel intensities) are rarely accurate to much better than 1 percent (about 7 bits equivalent). In vision applications, the demand is for a *massive volume of computation*, rather than high precision.

In our research, we use the speed and area advantages of analog *systems,* as distinct from single components like multipliers. For example, the CCD 2-D *binomial convolver* we are building is conceptually an analog systolic array, capable of massively parallel computation at an estimated 10 MHz rate using minimal area. The *parallel spatial correlator* will achieve a similar level of parallelism in a small space using the special properties of resistive grids.

Advances in parallel computation and VLSI technology will also enable the continued growth of general purpose digital computer capabilities. We expect that advanced digital machines will remain the tool of choice for high-level, i.e., *late,* vision tasks such as object recognition, while the analog implementation of the smart sensor paradigm will enable enormous further reductions in machine size, power consumption, speed and cost required for low-level or *early* vision processing.

Project Direction and Goals

The goal of this project is to develop novel, high-performance silicon systems for several real-time processing tasks in early vision. We are pursuing two parallel routes. The first leads to *single-chip sensor and processor systems,* while the second leads to a *modular early vision system.*

As used here, the term "single-chip sensor and processor system" refers to any chip that acquires image data through a photoreceptor array, has processors spatially distributed throughout the array, and performs a specialized form of parallel computation resulting in low-bandwidth output. Most systems designed in Carver Mead's Laboratory[10] are of this form.[11] These systems are small in physical size, inexpensive to manufacture, and specialized in application. The term "modular early vision system" refers to a larger, much less specialized system that will require board-level implementation or, perhaps eventually, wafer-scale integration. Its primary output will consist of arrays of intensity, depth and velocity information, and it will be useful as a front-end processor for a variety of vision applications.

---

[9] B. Gilbert, "Four Quadrant Analog Divider/Multiplier with 0.01 Percent Distortion," *ISSCC Digest of Technical Papers,* pp. 248-249, February 1983.

[10] The silicon retina only partially fits this category as defined above. That system passes entire images off chips at a high data rate and may not be considered very specialized, within the vision domain.

[11] J. Tanner, and C. Mead, "An Integrated Optical Motion Sensor," *VLSI Signal Processing II,* paper presented at the ASSP Conference on VLSI Signal Processing, University of California, Los Angeles, November 5-17, 1986; S.Y. Kung, R.E. Owen, and J.G. Nash, eds., IEEE Press, 59-76, 1987; M.A. Sivilotti, M.A. Mahowald, and C.A. Mead, "Real Time Visual Computations Using Analog CMOS Processing Arrays," in *Advanced Research in VLSI: Proceedings of the 1987 Stanford Conference,* ed. P. Losleben (Cambridge: MIT Press, 1987), 295-312; C.A. Mead and M.A. Mahowald, "A Silicon Model of Early Visual Processing," *Neural Networks* 1 (1):91-97 (1988); S. DeWeerth and C.A. Mead, "A Two-Dimensional Visual Tracking Array," in *Advanced Research in VLSI: Proceedings of the 1988 MIT Conference* (Cambridge: MIT Press, 1988); C.A. Mead, *Analog VLSI and Neural Systems* (Reading, Massachusetts: Addison-Wesley, 1989).

These two routes occupy complementary positions in our program. The former will yield a working system quickly, enabling us to evaluate our approach at an early stage. Although the latter allows for continued leverage from our past work for future system designs, it requires a major, time-consuming initial investment in developing the early vision modules.

This project has distinct goals that lie both within and beyond the three-year funding period of this project. Within three years, we will make components that are useful for a variety of early vision tasks. In five years, we will make a more general early vision system using these components and a more advanced specialized component useful for robot navigation. Our long-term goal is an advanced early vision system with highly processed, low-bandwidth output representing robustly segmented images.

Three-Year Goals: For the modular early vision system, we are building three components that will be useful for a wide range of image processing tasks: (1) an analog charged-coupled-device (CCD) restoring *image memory*, (2) a parallel pipelined CCD 2-D *binomial convolver*, and (3) a parallel *spatial correlator* using resistive grid averagers. These three systems are the foundation for our long-term plans outlined on the following pages.

MOSIS presently has no CCD capability. Therefore, we will develop a CCD enhancement to MIT's baseline CMOS process to fabricate the CCD binomial convolver and image memory. We are exploring both MOS and bipolar designs for the resistor circuits and arithmetic circuits used in the correlator. The fairly rapid turnaround CMOS process at MOSIS and the BiCMOS process on campus give us valuable flexibility in developing this system.

During the third year, we will complete design and simulation of the modular early vision system which will employ these components. Basing these simulations in part on experimental data from test circuits, we will modify circuit specifications as a function of simulated overall system performance.

For the single-chip sensor and processor part of the project, we will fabricate a working single-chip *image sensor and moment extractor* through MOSIS in the second year. Experimental data from testing and evaluating it will play a role in our development of a more general performance evaluation methodology for early vision systems.

Five-Year Goals: The work in this project is organized around two five-year goals. The first is development of a *single-chip motion sensor and analyzer* useful for robot navigation. Based on a new algorithm, it is a complex system using many of the circuits developed for the simpler image sensor and moment extractor.

The second goal is a *modular early vision system:* a relatively inexpensive, low-power, single-board system that will operate in real time. Its sensor arrays can be placed at the focal plane of a pair of lenses to acquire binocular image input directly from the environment. Its outputs will consist of the intensity field, velocity field, and depth map of the scene. Edge detector circuits will detect and enhance discontinuities in the individual fields.

The CCD imager, image memory, spatial correlator, and binomial convolver with edge detector circuits are key components in this system. Our current plans call for discrete- and continuous-time circuits with continuous voltage levels used to represent pixel values — only the edge detector outputs will be binary. This system will be of industrial value as a front-end processor for vision applications such as robot navigation and object recognition.

6.2.2 Progress in 1989

The National Science Foundation provided the initial funding for this project in September 1988; DuPont Corporation provided additional funding in December 1988.

The following describes our accomplishments in 1989:

- Mr. Craig Keast, working with Professor Charles Sodini, has designed, implemented and tested the additional process

steps that enable the MIT CMOS fabrication line to produce CMOS/CCD chips. We can now make 4-phase buried-channel CCDs.

- Mr. David Standley, working with Professors Berthold Horn and John Wyatt, has designed and tested a novel analog CMOS chip to determine the position and orientation of objects using linear and quadratic moments. The chip operates in continuous time on an optical image focused directly on its surface. A phototransistor array senses the image, processor cells transduce image intensity to current, and a uniform grid of resistors (connecting each processor to its four nearest neighbors) carries the currents to the edge of the array. The exciting feature of this design is that there is no need to address each processor separately — the resistive grid automatically conveys the required information to its periphery for further processing.

- Mr. Woodward Yang, working with Dr. Alice Chiang at Lincoln Laboratories and Professor Tomaso Poggio at MIT, has designed and begun testing a parallel, pipelined 2D CCD binomial convolver. This is a very high speed image bandpass filtering device useful in edge detection and a variety of vision applications.

- Resistive grid networks are frequently useful in parallel analog image processing operations on chips. Mr. Christopher Umminger, working with Professor Charles Sodini, has begun experimental work on switched capacitor networks as an alternative to resistor networks. The major advantages of switched capacitors are their smaller area and controllable effective resistance. Mr. Umminger has designed a switched-capacitor test chip to investigate this alternative method of fast image processing.

- Nonlinear resistors with incrementally active regions of operation are useful for image smoothing and segmentation tasks, but the design of small, accurate controllable circuits for this purpose remains a challenge. Mr. Steve Decker, working with Professors John Wyatt and Charles Sodini, has created some promising designs for these circuits and will fabricate a test chip that requires custom processing available at MIT.

- As a person or robot moves through a fixed environment, most of the scene projected on the retina or focal plane appears to move as well. But the image of the point toward which the motion is directed, i.e., the eventual point of impact if the motion were to continue in a straight line, expands without moving. The image of this point in the focal plane is called the *focus of expansion.* The magnitude of the divergence of the velocity field at this point is known to be inversely related to the *time of impact.* Ignacio McQuirk, working with Professors Berthold Horn and John Wyatt, is simulating algorithms that will automatically locate the focus of expansion of a moving scene projected onto it. His simulation studies are aimed toward the design of an analog chip that will rapidly perform this computation.

Publications

Harris, J., C. Koch, J. Luo, and J. Wyatt. "Resistive Fuses: Analog Hardware for Detecting Discontuities in Early Vision." In *Analog VLSI Implementations of Neural Systems,* eds. C. Mead and M. Ismail, 27-56. New York: Kluwer, 1989.

Hengeveld, J.A. *Smooth Surface Reconstruction in VLSI.* S.B. thesis. Dept. of Electr. Eng. and Comput. Sci., MIT, 1989.

Horn, B.K.P., H.-S. Lee, T. Poggio, C. Sodini, and J. Wyatt. *The First Year of the MIT Vision Chip Project.* VLSI Memo No. 89-555. MIT, 1989.

Keast C.L., P. Tedrow and C.G. Sodini. "A CCD/CMOS Process for Integrated Image Acquisition and Early Vision Signal Processing." Abstract presented at the VLSI Research Review, MIT, May 22, 1989.

Keast, C.L. *A CCD/CMOS Process for Integrated Image Acquisition and Early Vision Signal Processing.* S.M. and E.E. thesis,

Dept. of Electr. Eng. and Comput. Sci., MIT, 1989.

Lumsdaine, A., J. Wyatt, Jr., and I. Elfadel. "Parallel Distributed Networks for Image Smoothing and Segmentation in Analog VLSI." Paper presented at the *Proceedings of the 28th IEEE Conference on Decision and Control,* Tampa, Florida, December 1989, pp. 272-279.

Lumsdaine, A., J. Wyatt, and I. Elfadel. "Parallel Distributed Networks for Image Smoothing and Segmentation in Analog VLSI." Abstract presented at the VLSI Research Review, MIT, December 18, 1989.

Standley, D.L., and J.L. Wyatt, Jr. "Stability Criterion for Lateral Inhibition and Related Networks that is Robust in the Presence of Integrated Circuit Parasitics." *IEEE Trans. Circuits Syst.* 36 (5):675-681 (1989).

Standley, D.L. "Design Criteria Extensions for Stable Lateral Inhibition Networks in the Presence of Circuit Parasitics." Paper presented at the *IEEE Symposium on Circuits and Systems.* Portland, Oregon, May 1989, pp. 837-840.

Standley, D.L. "Criteria for Robust Stability in a Class of Lateral Inhibiton and Related Networks Coupled Through a Resistive Grid." *Proceedings of the 1989 Conference on Information Sciences and Systems,* Johns Hopkins University, Baltimore, Maryland, March 1989, pp.441-446.

Standley, D., and J.L. Wyatt, Jr. "Stability Criterion for Lateral Inhibition and Related Networks that is Robust in the Presence of Integrated Circuit Parasitics." Abstract presented at the VLSI Research Review, MIT, December 18, 1989.

Wyatt, J.L., Jr., and D.L. Standley. "Criteria for Robust Stability in a Class of Lateral Inhibition Networks Coupled Through Resistive Grids." *Neural Computation* 1 (1):58-67 (1989).

Yang, W. "A Charged-Coupled Device Architecture for On Focal Plane Image Signal Processing." *Proceedings of the International Symposium on VLSI Technology, Systems and Applications.* Taipei, Taiwan, R.O.C., May 1989, pp. 266-270.

Yang, W. "Method and Charge Coupled Apparatus for Algorithmic Computations." Patent application submitted to U.S. Patent Office on May 1, 1989, receipt number 07/352,765.

Yang, W., and A.M. Chiang. "VLSI Processor Architectures for Computer Vision," *Proceedings of the Image Understanding Workshop.* Palo Alto, CA, May 1989, pp. 193-199.

6.3 Techniques for Logic Synthesis, Testing and Design-for-Testability

Sponsors

U.S. Navy - Office of Naval Research
Contract N00014-87-K-0825

Project Staff

Professor Srinivas Devadas, Michael J. Bryan, Curtis S. Chen, Kevin Lam

This research focuses on the development of Computer-Aided Design (CAD) tools for integrated circuits with area, performance and reliability as design parameters.

To produce a highly reliable large-scale computer, it is vital for its hardware components to be individually reliable. It is essential to detect failures in VLSI components caused by defective manufacturing processes. Each component of a computer must be *tested* for correct functionality before assembly.

Test strategies must keep pace with increasing circuit complexity. Increased emphasis must be placed on finding a defect as early as possible in the manufacturing cycle, new algorithms must be devised to create tests for logic circuits, and more attention must be given to design-for-test techniques that require active participation by the logic designers.

Today's logic designer must work toward a multivariate objective — the designer must meet area, performance and technological constraints while also trying to ensure that the design is testable with a small set of patterns. Needless to say, it is almost impossible for the designer to find an optimal tradeoff of all these competing objectives for anything except the most simple circuits. For this reason, designers will often attempt to find a feasible design that meets performance and area goals and later modify the design to meet test requirements. Testing, even for combinational circuits, has been a post-design activity.

Current research focuses on automated logic synthesis and optimization strategies, targeting area and performance objectives. The logic restructuring steps used in the optimization of combinational and sequential circuits have a *very* profound affect on the testability of the eventual circuit. The purpose of our research is to *completely integrate reliability and testability considerations* into automated synthesis and optimization procedures instead of the "after-thought" post-design-for-test methodologies in vogue today. We believe that this integration can result in design of logic circuits which are not only optimal for area and performance, but also have *unprecedented levels of reliability, testability and diagnosability.*

Our work focuses on several critical areas which are elaborated on in the following sections.

6.3.1 Synthesis of Fault-Tolerant Designs

Use of static redundancy for system level fault-tolerance is an established practice. However, massive hardware redundancy introduced to mask faults can also prevent the detection of potential faults in the circuit during manufacturing testing and system diagnostics. Therefore, a fault-tolerant circuit must be completely testable in terms of both original circuit function and fault masking capability.

For example, to ensure desired reliability, digital systems used in a deep-space probe must be tested for both correct functional operation as well as for the capability of the circuit to mask faults. If the fault masking capability allows the original circuit function to perform properly when faults are present, the fault-tolerant circuit may not be reliable during system operation. In other words, the next fault sustained by the system could exceed the threshold of the fault masking capability so that the circuit no longer performs the original logic function correctly. Therefore, complete testability of the original circuit function as well as the fault masking capability is necessary during manufacturing testing. We are currently developing synthesis techniques targeting circuit testability and diagnosability.[12]

6.3.2 Combinational Logic Synthesis for Testability

Combinational logic synthesis and optimization is a well-understood problem. The traditional cost functions used in synthesis are layout area and circuit performance. Researchers have only recently begun investigating the effect of logic transformations on single and multiple stuck-at fault testability of a circuit. This work concentrates on deriving algebraic and Boolean transformations for logic circuits to improve testability (and remove redundancy) with minimal or no impact on area and performance. Methods for *implicit* test generation, i.e., obtaining test vectors to detect stuck-at faults as a by-product of the optimization steps, are being investigated. These methods are potentially much more efficient than explicit test gener-

[12] S. Devadas, H-K.T. Ma, A.R. Newton, and A. Sangiovanni-Vincentelli, "Irredundant Sequential Machines Via Optimal Logic Synthesis," *IEEE Trans. Comput.-Aided Des.* 9 (1):8-18 (1990); S. Devadas and K. Keutzer, "Necessary and Sufficient Conditions for Robust Delay-Fault Testability of Logic Circuits," in the *Proceedings of the Sixth MIT Conference on Advanced Research on VLSI* (Cambridge: MIT Press, 1990).

ation algorithms, especially for multiple stuck-at faults and delay faults.[13]

6.3.3 Sequential Logic Synthesis for Testability

A digital computer can be viewed as a set of interacting finite state machine (FSM) controllers and datapath components. Test generation for large sequential circuits is recognized as a very difficult problem. We are studying this problem by using both synthesis-for-testability and explicit test generation techniques. We began preliminary investigations into the nature of the relationships between testability of interacting FSMs and the optimization steps used during FSM synthesis such as state minimization, decomposition, state assignment and combinational logic optimization.[14]

A "perfect" set of optimization steps can ensure that all stuck-at faults in the synthesized FSMs are testable. In reality, "perfect" optimization is difficult to attain. However, it appears feasible that highly testable control portions and test sequences that detect faults within these controllers can be obtained implicitly during synthesis with reasonable CPU time expenditure. We will use Register-Transfer (RT) level as well as State Transition Graph descriptions of sequential circuits as a starting point for these testability-driven synthesis approaches. We will also investigate alternate representations that are more compact and/or easy to manipulate. This work can potentially reduce the area and performance overheads imposed by conventional design-for-testability schemes (e.g., LSSD) and dramatically reduce the time required to test the chip.

6.3.4 Test Generation for Sequential Circuits

Datapath components in a VLSI circuit have a large number of memory elements with a very regular structure. We are developing strategies based on our work[15] that exploit the *specific topology* of a datapath component to efficiently generate tests for faults within the component. Exploiting design knowledge available in the register-transfer level specification of the component can dramatically speed up test generation.[16]

For example, one can use to advantage the knowledge that an arithmetic unit performing addition and subtraction is embedded in a datapath during the backward justification step in test generation. If one requires a set of binary values at the outputs of the adder, rather than justifying back across many logic gates in a gate-level description, one can simply find the different pairs of numbers that sum up to the given output of the adder (and justify the output at a higher level of abstraction).

---

[13] S. Devadas and K. Keutzer, "Necessary and Sufficient Conditions for Robust Delay-Fault Testability of Logic Circuits," in the *Proceedings of the Sixth MIT Conference on Advanced Research on VLSI* (Cambridge: MIT Press, 1990).

[14] S.Devadas, H-K.T. Ma, A.R. Newton, and A. Sangiovanni-Vincentelli, "Irredundant Sequential Machines Via Optimal Logic Synthesis," *IEEE Trans. Comput.-Aided Des.* 9 (1):8-18 (1990); S. Devadas and K. Keutzer, "A Unified Approach to the Synthesis of Fully Testable Sequential Machines," in the *Proceedings of the 23rd Hawaii Conference on System Sciences* (IEEE Computer Society Press, 1990); S. Devadas, "Delay Test Generation for Synchronous Sequential Circuits," in the *Proceedings of the International Test Conference* (Washington, DC: August 1989), 144-152.

[15] A. Ghosh, S. Devadas, and A.R. Newton, "Test Generation for Highly Sequential Circuits," paper presented at the International Conference on Computer-Aided Design, Santa Clara, California, November 1989.

[16] A. Ghosh, S. Devadas, and A.R. Newton, "Sequential Test Generation at the Register Transfer and Logic Levels," paper presented at the 27th Design Automation Conference, Orlando, Florida, June 1990.

6.3.5 Synthesis for Delay-Fault Testability

Physical defects in integrated circuits (ICs) can degrade performance without affecting logic functionality. It is important to ensure that a given IC satisfies its timing specifications — this requires that a *delay fault test* be performed. This is especially important for aggressive design methodologies that push processing technology to the utmost.

Unfortunately, most circuits being designed/synthesized today are typically difficult to test for delay faults. This is because the conditions for *robust* delay-fault testability are quite stringent, requiring the single sensitization of paths to ensure hazard-free error propagation. In this part of our research project, we concentrated on synthesis methodologies for completely and robustly delay-fault-testable combinational and sequential logic circuits, continuing our preliminary work.[17] We found necessary and sufficient conditions for robust delay-fault testability of combinational circuits,[13] and used these conditions as desirable requirements of synthesis procedures.[18]

6.3.6 Performance-Directed Logic Synthesis

To design high-performance computers, each of the processing elements constituting the computer is required to have minimal delay. Processing elements are comprised of combinational and sequential logic circuits. We are focusing on the automated, performance-directed optimization of logic functions with a view of synthesizing logic that is as fast as, if not faster than, implementations designed manually.

For purely combinational logic functions the Boolean and algebraic transformations in use today are mature and well-developed as far as minimizing for layout area is concerned. However, circuit delay, which depends on complex technological and topological considerations, proved to be a more difficult objective function. We recently proposed a form of logic optimization that we term four-level Boolean minimization.[19] We are investigating a formal and elegant way of incorporating delay considerations into four-level Boolean minimization. This incorporation can result in a logic optimization strategy that exactly minimizes the delay of a logic function.

The optimization of interacting sequential circuits for area and performance has received very little attention. Optimizing for performance entails the redistribution or migration of logic across latch boundaries to minimize the required clock period. Classical retiming techniques cannot be used if the basic blocks forming the interconnection are finite state machines (FSMs), rather than being simply combinational. To address this problem, we developed automata-theoretic FSM decomposition strategies that *redecompose* a set of interacting FSMs to minimize critical path delay.[20]

Publications

Ashar, P., S. Devadas, and A.R. Newton. "Optimum and Heuristic Algorithms for Finite State Machine Decomposition and Partitioning." Paper presented at the International Conference on Computer-Aided Design, Santa Clara, California, November 1989.

---

[17] S. Devadas, "Delay Test Generation for Synchronous Sequential Circuits," in *Proceedings of the International Test Conference* (Washington, DC: August 1989), 144-152.

[18] S. Devadas and K. Keutzer, "Synthesis and Optimization Procedures for Delay-Fault Testable Combinational Logic," paper to be presented at the 27th Design Automation Conference, Orlando, Florida, June 1990.

[19] S. Devadas and A.R. Newton, "Exact Algorithms for Output Encoding, State Assignment and Four-Level Boolean Minimization," *Proceedings of the 23rd Hawaii Conference on System Sciences.*

[20] K. Lam and S. Devadas, "Performance-Oriented Synthesis of Finite State Machines," paper presented at the International Conference on Circuits and Systems, New Orleans, May 1990.

Ashar, P., S. Devadas, and A.R. Newton. "A Unified Approach to Decomposition and Re-Decomposition of Sequential Machines." Paper presented at the 27th Design Automation Conference, Orlando, Florida, June 1990.

Ashar, P., S. Devadas, and A.R. Newton. "Multiple-Fault Testable Sequential Machines." Paper presented at the International Conference on Circuits and Systems, New Orleans, May 1990.

Devadas, S. "Delay Test Generation for Synchronous Sequential Circuits." Paper presented at the International Test Conference, Washington, DC, August 1989.

Devadas, S. "Minimization of Functions with Multiple-Valued Outputs: Theory and Applications." Paper presented at the 20th International Symposium on Multiple-Valued Logic, Charlotte, North Carolina, May 1990.

Devadas, S. "Optimal Layout Via Boolean Satisfiability." Paper presented at the International Conference on Computer-Aided Design, Santa Clara, California, November 1989.

Devadas, S. "Optimal Layout Via Boolean Satisfiability." *Int. J. Comput.-Aided VLSI Des.* Forthcoming.

Devadas, S., and K. Keutzer. "A Unified Approach to the Synthesis of Fully Testable Sequential Machines." Paper presented at the 23rd International Conference on System Sciences, Hawaii, January 1990.

Devadas, S., and K. Keutzer. "Synthesis and Optimization Procedures for Robust Delay-Fault Testability of Logic Circuits." Paper presented at the 27th Design Automation Conference, Orlando, Florida, June 1990.

Devadas, S., K. Keutzer, and J. White. "Estimation of Power Dissipation in CMOS Combinational Circuits." Paper presented at the Custom Integrated Circuits Conference, Boston, May 1990.

Devadas, S., H-K.T. Ma, A.R. Newton, and A. Sangiovanni-Vincentelli. "Irredundant Sequential Machines Via Optimal Logic Synthesis." Paper presented at the 23rd International Conference on System Sciences, Hawaii, January 1990.

Devadas, S., H-K.T. Ma, A.R. Newton, and A. Sangiovanni-Vincentelli. "A Synthesis and Optimization Procedure for Fully and Easily Testable Sequential Machines." *IEEE Trans. Comput.-Aided Des.* 8 (10):1100-1107 (1989).

Devadas, S., and A.R. Newton. "Exact Algorithms for Output Encoding, State Assignment and Four-Level Boolean Minimization." Paper presented at the 23rd International Conference on System Sciences, Hawaii, January 1990.

Devadas, S., and K. Keutzer. "Logic Synthesis and Testability." *Proceedings of the International Conference on Circuits and Systems,* New Orleans, May 1990.

Devadas, S., and K. Keutzer. "Validatable Nonrobust Delay-fault Testable Circuits Via Logic Synthesis." *Proceedings of the International Conference on Circuits and Systems,* New Orleans, May 1990.

Devadas, S., H-K. T. Ma, and A.R. Newton. "Redundancies and Don't Cares in Sequential Logic Synthesis." Paper presented at the International Test Conference, Washington, DC, August 1989.

Devadas, S., H-K. T. Ma, A.R. Newton, and A. Sangiovanni-Vincentelli. "The Relationship Between Logic Synthesis and Test." Paper presented at VLSI '89, Munich, August 1989.

Devadas, S., and K. Keutzer. "Boolean Minimization and Algebraic Factorization Procedures for Fully Testable Sequential Machines." Paper presented at the International Conference on Computer-Aided Design, Santa Clara, California, November 1989.

Devadas, S., and A.R. Newton. "Algorithms for Hardware Allocation in Datapath Syn-

thesis." *IEEE Trans. on Comput.-Aided Des.* 8(7):768-791 (1989).

Devadas, S., and A.R. Newton. "Decomposition and Factorization of Sequential Finite State Machines." *IEEE Trans. Comput.-Aided Des.* 8 (11):1206-1217 (1989).

Devadas, S., H-K.T. Ma, A.R. Newton, and A. Sangiovanni-Vincentelli. "Irredundant Sequential Machines Via Optimal Logic Synthesis." *IEEE Trans. Comput.-Aided Des.* 9 (1):8-18 (1990).

Devadas, S., and K. Keutzer. "Necessary and Sufficient Conditions for Robust Fault-Delay Testability of Combinational Logic Circuits." Paper presented at the *Sixth MIT Conference on Advanced Research in VLSI,* Cambridge, April 1990.

Devadas, S., and K. Keutzer. "A Unified Approach to the Synthesis of Fully Testable Sequential Machines." *IEEE Trans. Comput.-Aided Des.* Forthcoming.

Devadas, S., and A.R. Newton. "Exact Algorithms for Output Encoding, State Assignment and Four-Level Boolean Minimization." *IEEE Trans. Comput.-Aided Des.* Forthcoming.

Devadas, S., H-K.T. Ma and A.R. Newton. "Redundancies and Don't Cares in Sequential Logic Synthesis." *J. Electronic Testing: Theory and Appl.* 1:15-30 (1990).

Devadas, S., and H-K.T. Ma. "Easily Testable PLA-based Finite State Machines." *IEEE Trans. Comput.-Aided Des.* Forthcoming.

Ghosh, A., S. Devadas, and A.R. Newton. "Test Generation for Highly Sequential Circuits." Paper presented at the International Conference on Computer-Aided Design, Santa Clara, California, November 1989.

Ghosh, A., S. Devadas, and A.R. Newton. "Verification of Interacting Sequential Circuits." Paper presented at the 27th Design Automation Conference, Orlando, Florida, June 1990.

Ghosh, A., S. Devadas, and A.R. Newton. "Sequential Test Generation at the Register-Transfer and Logic Levels." Paper presented at the 27th Design Automation Conference, Orlando, Florida, June 1990.

Lam, K., and S. Devadas. "Performance-Oriented Synthesis of Finite State Machines." Paper presented at the International Conference on Circuits and Systems, New Orleans, May 1990.

6.4 Mixed Circuit/Device Simulation

Sponsors

International Business Machines Corporation
U.S. Navy - Office of Naval Research
Contract N00014-87-K-0825

Project Staff

Mark W. Reichelt, Professor Jacob White, Professor Jonathan Allen

For critical applications, the four-terminal lumped models for MOS devices used in programs such as SPICE are not sufficiently accurate. Also, using lumped models to relate circuit performance to process changes is difficult. We can perform sufficiently accurate transient simulations if instead of using a lumped model for each transistor, we compute some of the transistor terminal currents and charges by numerically solving the drift-diffusion based partial differential equation approximation for electron transport in the device. However, because accurate solution of the transport equations of an MOS device requires a two-dimensional mesh with more than a thousand points, simulating a circuit with even a few of the transistors treated by solving the drift-diffusion equations is very computationally expensive.

One way of accelerating mixed device and circuit simulation is with waveform relaxation, which performs the transient simulation, not only at the circuit level, but also inside the devices being simulated with a drift-diffusion description. In this investigation, we apply the waveform relaxation (WR) algorithm to the sparsely-connected system of algebraic and ordinary differential equations in time generated by standard spatial

discretization of the drift-diffusion equations that describe MOS devices. Two theoretical results indicate that the method will be effective. One result suggests that the WR algorithm will converge in a uniform manner independent of the time interval,[21] and the other implies that the multirate integration will be stable.[22]

In addition, in preliminary experiments, we have used waveform relaxation to perform transient two-dimensional simulation of MOS devices. Experiments demonstrate that WR converges in a uniform manner and that there is typically some multirate behavior in a device that the WR algorithm can exploit. Speed and accuracy comparisons between standard direct methods, red/black Gauss-Seidel WR and red/black overrelaxed WR, indicate that, for the experiments examined, calculated terminal currents match well between the methods and that overrelaxed WR was between two and five times faster than direct methods.[23] A recently implemented modification based on a waveform-Newton algorithm increased this from a factor of 5 to 11.[24]

6.5 Circuit Simulation Algorithms for Specialized Applications

Sponsors

American Telephone and Telegraph
Analog Devices
Digital Equipment Corporation
International Business Machines Corporation
National Science Foundation
Grant MIP 88-58764
U.S. Navy - Office of Naval Research
Contract N00014-87-K-0825

Project Staff

Andrew Lumsdaine, Luis M. Silveira, Professor Jacob White, Professor John L. Wyatt

For all practical purposes, the general circuit simulation problem has been solved. Given enough time, programs like SPICE or ASTAP are capable of simulating virtually any circuit. Unfortunately, for some types of circuits, "enough time" is too much time for simulation to be a practical part of a circuit design cycle. To obtain large performance increases over a program such as SPICE, one must exploit special properties of the specific problem to be solved. In particular, we are developing fast and accurate simulation algorithms for two application areas: (1) the simulation of clocked analog circuits like switching filters, switching power supplies, or phase-locked loops and (2) the simulation of analog signal processing circuits used for early vision.

The simulation of clocked analog circuits such as switching filters, switching power supplies, and phase-locked loops poses a challenging problem. These circuits are computationally expensive to simulate using conventional techniques because they are all clocked at a frequency with a period orders of magnitude smaller than the time interval of interest to the designer. To construct such a long time solution, a program such as SPICE or ASTAP must calculate the behavior of the circuit for many high frequency clock cycles. The basic approach to making simulation of

---

[21] M. Reichelt, J. White, J. Allen and F. Odeh, "Waveform Relaxation Applied to Transient Device Simulation," paper presented at the International Symposium on Circuits and Systems, Espoo, Finland, June 1988.

[22] M. Crow, J. White and M. Ilic, "Stability and Convergence Aspects of Waveform Relaxation Applied to Power System Simulation," in *Proceedings of the International Symposium on Circuits and Systems,* Portland, Oregon, May 1989.

[23] M. Reichelt and J. White, "Techniques for Switching Power Converter Simulation," invited paper presented at the NASECODE Conference, Dublin, Ireland, June 1989.

[24] M. Reichelt, J. White and J. Allen, "Waveform Relaxation for Transient Simulation of Two-Dimensional MOS Devices," paper published the *Proceedings of the International Conference Computer-Aided Design,* Santa Clara, California, October 1989; R. Saleh, J. White, A. S. Vincentelli, A. R. Newton, "Accelerating Relaxation Algorithms for Circuit Simulation Using Waveform-Newton and Step-Size Refinement," *IEEE Trans. Comput.-Aided Des.* (October 1990), forthcoming.

these circuits more efficient is to exploit only the property that the behavior of such a circuit in a given high frequency clock cycle is similar to, but not identical to, the behavior in the preceding and following cycles. Therefore, by accurately computing the solution over a few selected cycles, an accurate long time solution can be constructed.

Simulating clocked analog systems is an old problem, but this novel approach has led to a very efficent algorithm for the distortion analysis of switched-capacitor filters. The idea is based on simulating selected cycles of the high-frequency clock accurately with a standard discretization method and pasting together the selected cycles by computing the low frequency behavior with a truncated Fourier series. If carefully constructed, the nonlinear system that must be solved for the Fourier coefficients is almost linear, and can be solved rapidly with Newton's method.[25] Transient behavior, important for switching power supply designers, was also accelerated using similar techniques.[26] In particular, the "envelope" of the high-frequency clock can be followed by accurately computing the circuit behavior over occasional cycles.

The "vision circuits" form a class of circuits which, for the most part, cannot be simulated with traditional programs. These circuits are necessarily very large and must be simulated exactly at an analog level (i.e., one cannot perform simulations at a switch or gate level as is commonly done with very large digital circuits). Typical analog circuit simulators are not able to handle vision circuits simply because of their immense size — the computation time for these simulators grows superlinearly with the size of the circuit. However, because these circuits are somewhat similar to certain discretized partial differential equations, one can exploit their special characteristics to obtain efficient simulation techniques.

During the past year, we completed the development of Simlab, an extensible circuit simulation program constructed to facilitate algorithm development for circuit simulation.[27] After adding vision simulation extensions to Simlab, we used the program to investigate the behavior of a class of nonlinear smoothing and segmentation networks. In addition, we added several conjugate-gradient based matrix solution algorithms to Simlab, testing their performance on resistive grid and vision circuits. Finally, we are investigating what we hope will be a novel and very general extension to the algebraic multigrid algorithm suitable for vision circuit simulation.

6.6 Numerical Simulation of Short Channel MOS Devices

Sponsors

Analog Devices
U.S. Navy - Office of Naval Research
Contract N00014-87-K-0825

Project Staff

Jennifer A. Lloyd, Joel R. Phillips, Khalid Rahmat, Professor Dimitri A. Antoniadas, Professor Jacob White

The model used in conventional device simulation programs is based on the drift-diffusion model of electron transport, but this model does not accurately predict the field distribution near the drain in small geometry

---

[25] K. Kundert, J. White, and A. Sangiovanni-Vincentelli, "A Mixed Frequency-Time Approach for Finding the Steady State Solution of Clocked Analog Circuits," paper presented at the Custom Integrated Circuits Conference, Rochester, New York, May 1988; K. Kundert and J. White, A. Sangiovanni-Vincentelli, "A Mixed Frequency-Time Approach for Distortion Analysis of Switched Capacitor Filters," *IEEE J. Solid State Circuits* 24 (2):? (1989).

[26] K. Kundert and J. White, A. Sangiovanni-Vincentelli, "An Envelope-Following Method for the Efficient Transient Simulation of Switching Power Converters," paper published in the *Proceedings of the International Conference on Computer-Aided Design,* Santa Clara, California, October 1988; M. Reichelt and J. White, "Techniques for Switching Power Converter Simulation," invited paper presented at the NASECODE Conference, Dublin, Ireland, June 1989.

[27] A. Lumsdaine, L. Silveira and J. White, "Simlab Users Guide," *MIT VLSI Memo*, September 1989; A. Lumsdaine, L. Silveira and J. White, "Simlab Programmers Guide," *MIT VLSI Memo,* September 1989.

devices. This is particularly important for predicting oxide breakdown due to penetration by "hot" electrons. There are two approaches for more accurately computing the electric fields in MOS devices: the first is based on adding an energy equation to the drift-diffusion model and the second on particle or Monte-Carlo simulations.

In the first approach, an energy balance equation is solved along with the drift-diffusion equations so that the electron temperatures are computed accurately. This combined system, which is numerically less tame than the standard approach, must be solved carefully. We have developed a 2-D simulator for MOS devices based on the drift-diffusion plus energy equations, uncovering a source of instability in the computation. The instability occurs when the semiconductor mobility is made a function of the inverse of carrier temperature. Using a very fine mesh removes this instability, but makes this method computationally expensive. Currently, we are trying to understand the cause of the instability so that we can use a coarse mesh and still produce accurate results.

In the area of Monte Carlo device simulation, we are focusing on transient calculations with self-consistent electric fields. Specifically, we are trying to apply the recently developed implicit particle methods to semiconductors by decomposing the field calculation into a part due to charged particles, a part due to dopant ions, and a part due to boundaries. This allows us to calculate the electric field acting on every charged particle in the system rapidly and accurately. Specifically, we are able to use the fast multipole algorithm for the particle-particle interactions. After rewriting (with J. Phillips) a Silicon Monte Carlo code from the National Center for Computational Electronics for use with ensemble Monte Carlo methods, we are now including the electric field calculations.

6.7 Efficient Capacitance Extraction Algorithms

Sponsors

International Business Machines Corporation
U.S. Navy - Office of Naval Research
Contract N00014-87-K-0825

Project Staff

Keith S. Nabors, Professor Jacob White

To achieve the dense packing of processors and memory required for high performance parallel computers, we must design interconnection aggressively. It is not possible to achieve high interconnect wire density by following commonly used conservative design rules for avoiding unintended coupling between wires. However, to assure some degree of reliable data transfer, we must examine signal corruption due to interconnect coupling. Extracting an equivalent resistor-capacitor-inductor (RLC) circuit for the interconnect is an effective approach that lends insight into the effects of coupling. However, this parasitic extraction calculation is computationally expensive because it implies solving electrostatic and magnetostatic equations in three space dimensions.

We developed and tested a fast algorithm for computing the capacitance of a complicated 3-D geometry of ideal conductors in a uniform dielectric.[28] The method is an acceleration of the standard integral equation approach for multiconductor capacitance extraction. These integral equation methods are slow because they lead to dense matrix problems typically solved with some form of Gaussian elimination. This implies the computation grows like n^3, where n is the number of tiles needed to accurately discretize the conductor surface charges. We developed a preconditioned conjugate-gradient iterative algorithm with a multipole approximation to compute the iterates. This reduces the complexity of the multiconductor capacitance calculations to grow as nm where m is the number of conductors. Recently, we began to improve the imple-

[28] K. Nabors and J. White, "A Fast Multipole Algorithm for Capacitance Extraction of Complex 3-D Geometries," in *Proceedings of the Custom Integrated Circuits Conference*, San Diego, California, May 1989.

mentation of the capacitance calculation algorithm. Specifically, we developed a faster code which includes a novel adaptive multipole algorithm. In the future, our work in this area will be to include dielectric interfaces in the development of a multipole algorithm to accelerate the calculation of inductances.

6.8 Parallel Numerical Algorithms

Sponsors

U.S. Navy - Office of Naval Research
Contract N00014-87-K-0825

Project Staff

Andrew Lumsdaine, Mark W. Reichelt, Luis M. Silveira, Ricardo Telichevsky, Professor Jacob White

We are trying to develop parallel algorithms for circuit and device simulation that are effective on either massively parallel machines similar to the Connection Machine, or on hypercube machines like the Intel hypercube. In addition, we are also trying to understand the fundamental problem involved in developing these algorithms by investigating the interaction between architecture and certain numerical algorithms (with Professor William J. Dally).

The key problem in parallelizing many of the numerical algorithms used in circuit and device simulators is finding efficient techniques for solving large sparse linear systems in parallel. Most parallel matrix solution algorithms fall into one of two general categories: the direct (Gaussian-elimination based) and the iterative; each algorithm presents different problems. The computation in the direct approach is not very structured, and is, therefore, difficult to parallelize. Iterative methods are easily parallelized, but are not as numerically robust.

The direct solution of circuit simulation matrices is particularly difficult to parallelize. Methods such as parallel nested dissection are ineffective because finding good separators is difficult. For that reason, we are studying general sparse matrix techniques and, specifically, the interaction between sparse matrix data structures, computer memory structure, and multiprocessor communication (with Professor Dally). The most interesting results from simulations so far is that communication throughput, and not latency, is more crucial to final performance.

As mentioned above, relaxation algorithms for solving matrices are easily parallelized. It is also possible to apply relaxation directly to the differential equations, a method referred to as waveform relaxation (WR), to develop a parallel algorithm in which different differential equations are solved on different processors. A recently developed variant of the WR algorithm, referred to as waveform-relaxation Newton (WRN), allows for additional parallelism. Much of the computation for each of the discretization timepoints for a single differential equation can be computed in parallel. This method may be appropriate for massively parallel computers. In recent theoretical work, we proved that WRN converges globally even when applied to circuits with nonlinear capacitors.[29] We also successfully applied WRN to device simulation. Currently, we are implementing a WR-based device simulation program on an Intel hypercube.

6.9 Integrated Circuit Reliability

Sponsors

Digital Equipment Corporation
U.S. Navy - Office of Naval Research
Contract N00014-87-K-0825

Project Staff

Professor Srinivas Devadas, Professor Jacob White

---

[29] R. Saleh, J. White, A. S. Vincentelli, A. R. Newton, "Accelerating Relaxation Algorithms for Circuit Simulation Using Waveform-Newton and Step-Size Refinement," *IEEE Trans. Comput.-Aided Des. (October 1990), forthcoming.*

The high transistor density now possible with CMOS integrated circuits makes peak power dissipation and peak current density important design considerations. However, because peak quantities in a logic circuit are usually a function of the input vector or vector sequence applied, accurate estimation of peak quantities is extremely difficult. The number of input sequences that must be simulated to find the sequence that produces the peak is *exponential* in the number of inputs to the circuit. By using simplified models of power and current dissipation, we can relate peak quantities such as power or current density to maximizing gate output activity, weighted to account for differing load capacitances or transistor sizes. Then, we can derive transformations that convert a logic description of a circuit into a multiple-output Boolean function of the input vector or vector sequence in which each output of the Boolean function is associated with a logic gate output transition. Next, we must solve a *weighted max-satisfiability* problem to find the input or input sequence that maximizes the quantity of interest. For solving the problem of estimating peak power dissipation, we derived algorithms for constructing the Boolean function for dynamic CMOS circuits, as well as for static CMOS, which take into account dissipation due to glitching. Lastly, we developed exact and approximate algorithms for solving the associated weighted max-satisfiability problem.[30]

---

[30] S. Devadas, K. Keutzer, J. White, "Estimation of Power Dissipation in CMOS Combinational Circuits," in *Proceedings of the Custom Integrated Circuits Conference,* Boston, 1990.