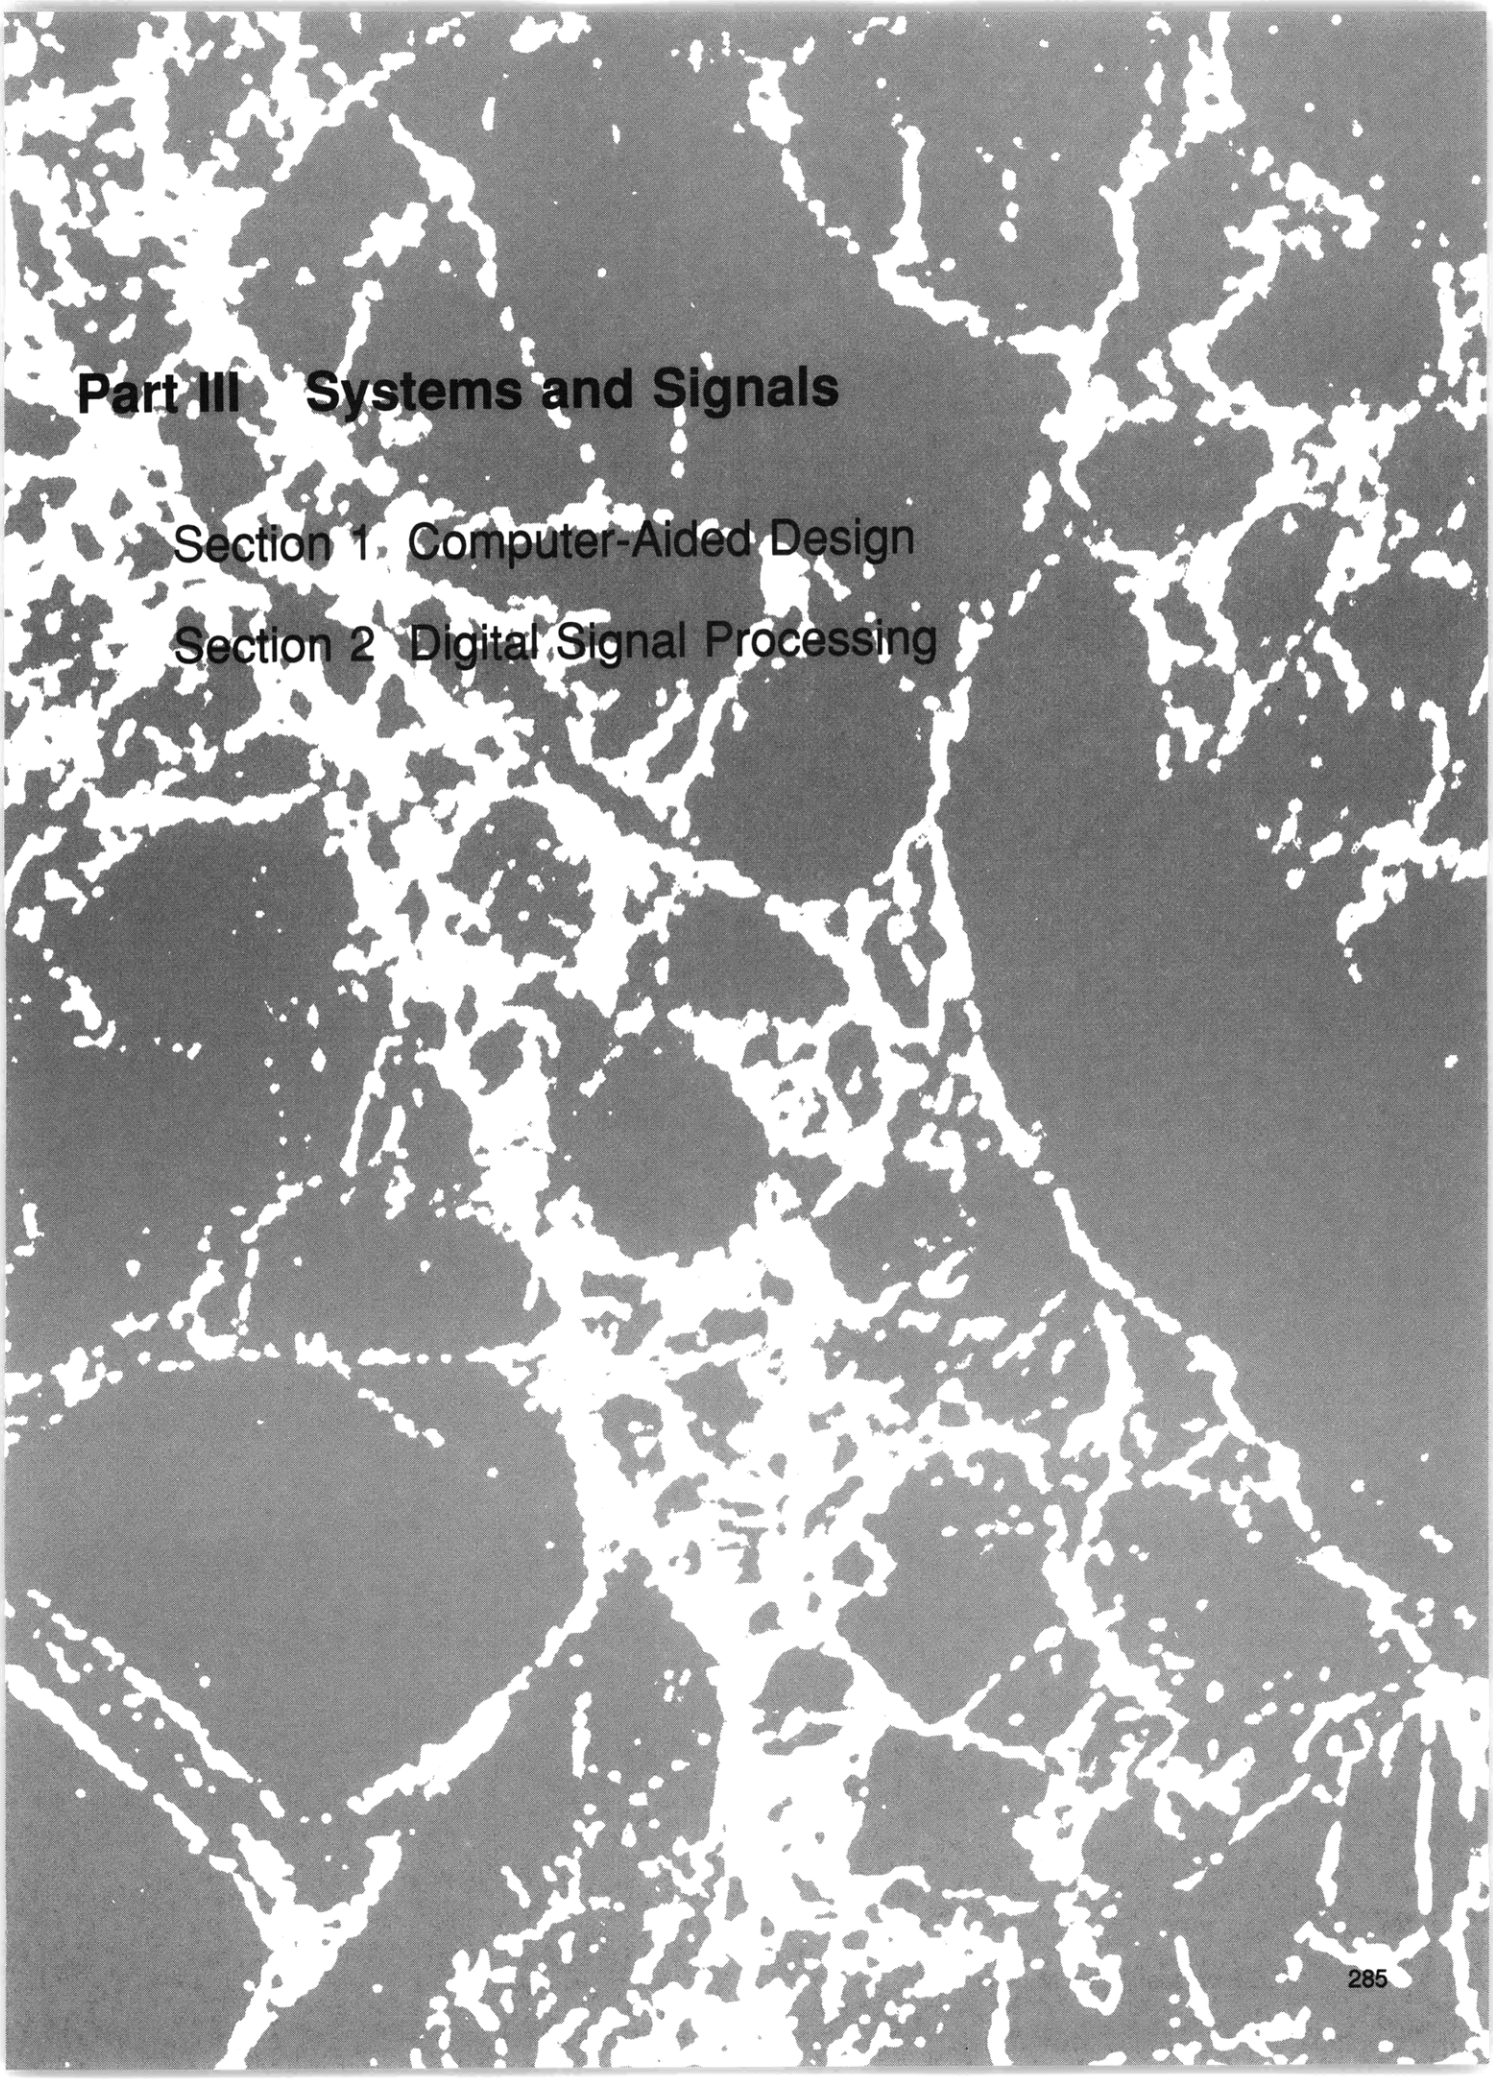

Part III Systems and Signals

Section 1 Computer-Aided Design

Section 2 Digital Signal Processing

Section 1 Computer-Aided Design

Chapter 1. Custom Integrated Circuits

Academic and Research Staff

Professor Jonathan Allen, Professor John L. Wyatt, Jr., Professor Jacob K. White, Professor Srinivas Devadas, Professor Dimitri A. Antoniadis, Professor Berthold K.P. Horn, Professor Hae-Seung Lee, Professor Steven B. Leeb, Professor Terry P. Orlando, Professor Martin A. Schmidt, Professor Stephen D. Senturia, Professor Charles G. Sodini, Dr. Robert C. Armstrong, Dr. Craig L. Keast

Visiting Scientists and Research Affiliates

Dr. Andrew Lumsdaine,[1] Dr. Ichiro Masaki,[2] Dr. Herre S.J. van der Zant, Dr. He Yie[3]

Graduate Students

Mazhar M. Alidina, Donald G. Baltus, Vishal L. Bhagwati, Xuejun Cai, Michael Chou, Steven J. Decker, Lisa G. Dron, Ibrahim M. Elfadel, J. Mikko Hakkarainen, Frederick P. Herrmann, Mattan Kamon, Songmin Kim, Stan Y. Liao, David R. Martin, Robert G. McDonald, Ignacio S. McQuirk, José C. Monteiro, Keith S. Nabors, Joel R. Phillips, Khalid Rahmat, Mark W. Reichelt,[4] W. John Rinderknecht, Mark N. Seidel, Amelia H. Shen, Luis M. Silveira, Chin Hwee Tan, Christopher B. Umminger, Paul C. Yu

Technical and Support Staff

Susan E. Chafe, Dorothy A. Fleischer

1.1 Custom Integrated Circuits

Sponsors

IBM Corporation
National Science Foundation

Project Staff

Professor Jonathan Allen, Dr. Robert C. Armstrong, Donald G. Baltus, Robert G. McDonald, Chin Hwee Tan

The overall goal of VLSI computer-aided design (CAD) research is to provide the means to produce custom integrated circuits quickly, correctly, and economically. Traditionally, correctness has been verified at several representational levels of abstraction, such as required topologies for given circuit families (e.g., static CMOS) and stipulated logic formalisms. These techniques for checking correctness are usually local to the particular representational level involved and, while they are important components of the overall design testing procedure, they do not attempt to provide for the alignment and consistency checks between the different abstract representational levels 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 in a way that ranges over the wide variety of possible performance levels. For this reason, there is a strong need to provide CAD tools that serve as a framework for design exploration when performance is a parameter, thus providing the desired performance together with consistently aligned representations at all levels.

This research group studies a variety of topics with an emphasis on performance-directed synthesis of custom VLSI designs. An overview of the viewpoint that motivates these projects is provided in a major survey paper[5] in which the need for coordinating the design optimization process over the several levels of representation is emphasized. Since design exploration is so central to the production of

---

1 Notre Dame University, Notre Dame, Indiana.

2 Computer Science Department, General Motors Research Laboratory, Warren, Michigan.

3 Microelectronics Center, Southeast University, Nanjing, China.

4 Mathworks, Natick, Massachusetts.

5 J. Allen, "Performance-Directed Synthesis of VLSI Systems," *Proc. IEEE* 78(2): 336-355 (1990).

high-performance designs, emphasis is placed on how performance can be characterized at the several levels of representation and how overall optimal performance can be achieved in an integrated way.

The conversion of a high-level functional specification to a detailed mask layout for a substantial processor is a major task that confronts modern CAD synthesis systems. When the target architecture for such a system is restricted, however, it is often possible to introduce metrics that facilitate optimized performance. With this in mind, graduate student Donald G. Baltus has focused on the optimal design of array logic using systolic representations that allow only nearest neighbor communication between processors in the array. Given this architectural restriction, it is possible to introduce realistic communication constraints into the synthesis process so that meaningful optimization of total delay for a calculation can be achieved. An input language called FLASH has been introduced to represent algorithms functionally. These algorithms are then converted to data-dependence graphs that are augmented to provide indexed representations, thus allowing compact representation of calculations that do not grow with the size of the array.

A major achievement of this work is the introduction of several constraints that allow for substantial reduction of the design space, yet without the possibility of rejecting candidates for the optimal design. One such constraint is due to locality, which requires communication only between nearest neighbor processors. While this constraint may seem restrictive, it introduces the benefit of design compaction through a set of regular transformations that increase the performance of the resultant configuration. There is a strong mathematical basis for this work, and it has been shown that the entire space of all candidate designs is comprehensively searched in an efficient way. The thesis that represents this work is now complete, and its results were recently presented at a meeting on array architectures.[6] It is believed that this is the first time guaranteed delay optimal designs for a broad class of systolic architectures is available. Experience with the system has indicated that many nonintuitive designs are revealed. Current extensions include the coupling of this program (called Descartes) to an augmented layout synthesis system[7] so that it will be possible to generate high-quality layouts directly from a high-level functional specification for this important class of architectures.

These results are an important contribution to systems synthesis and architectural exploration. They are fundamentally and mathematically based to scientifically facilitate high-quality synthesis, rather than being based on heuristic strategies, which are often motivated from a limited understanding of the basic design problem.

In an attempt to minimize power while preserving speed performance, graduate student Chin Hwee Tan has developed a design system based on a standard cell library that seeks to provide the minimum power implementation subject to a given delay constraint. The motivation for this research is the observation that it may be possible to increase the size of devices driving the node in order to gain speed without an excessive power penalty, if indeed the node can be demonstrated to switch at a sufficiently low frequency.

These studies take advantage of new capabilities to estimate the switching frequency of nodes in combinational circuits, and hence, they represent a new class of studies based on realistic power estimates. Ms. Tan has created a standard cell library, including detailed layouts where sizing has been introduced, to achieve a range of trade-offs between delay and power. Specifically, three designs for each gate have been introduced in the library and are available to the design system. Once the design has been mapped to a set of gates, and a delay specification has been introduced, power minimization is achieved by initially finding the critical path through the network. This is the longest delay path in the network, and hence, the first candidate for the introduction of faster cells in order to meet the delay specification while potentially minimizing added power.

A variety of strategies were introduced to increase performance, including the identification of nodes that present a heavy load to the previous stage, while also having a low switching frequency. These nodes identify the previous stage as a candidate for the use of cell versions that have wider transistors and are capable of driving the large load at a higher speed. Reordering of inputs to gates (such as for a three-input NAND gate) can be shown to have a

---

6 D.G. Baltus and J. Allen, "Efficient Exploration of Space-Time Transformations for Systolic Array Synthesis," paper presented at the International Conference on Application-Specific Array Processors, Venice, Italy, October 25-27, 1993.

7 D. Baltus and J. Allen, "SOLO: A Generator of Efficient Layouts from Optimized MOS Circuit Schematics," *Proceedings of the International Symposium on Circuits and Systems,* Espoo, Finland, June 1988, pp. 1647-1650.

substantial effect on performance, while not disturbing the power dissipated by the gate. The algorithms in this system provide a rigorous capability to exploit this potential. Techniques were also introduced to relax and decrease the size of transistors that are not on and do not influence the critical path, and hence, these transistors may be made smaller. This is a desirable direction for several reasons, including the fact that the algorithm tends to equalize the delay of paths to the network, reducing the possibility of glitches and providing a more optimal design from the viewpoint of wave pipelining. Experience with a wide class of combinational circuits has shown that it is possible to reduce delay up to 43 percent with only an 8 percent increase in power.

Possible extensions to the algorithm include the development of a similar yet separate algorithm to optimize delay under a power constraint. In addition, consideration of false paths will be introduced, and the study of basic techniques to optimize convex systems that use discrete variables will be made. Although this study has focused only on static CMOS circuit topologies,[8] other circuit styles may also be investigated.

As reported in previous *RLE Progress Reports*, Armstrong has developed a comprehensive system for the incremental consistency maintenance in a multirepresentational structural VLSI database. Since the original presentation of this research,[9] several extensions have been made. These were motivated by the overall complexity of the code needed for such a system and requirements for flexibility and interfaces to a variety of other CAD systems.

The most significant accomplishment was the generation of the entire VLSI CAD database system from high-level schema descriptions. This capability includes the potential for custom graphic user interfaces when specified as schemas. Using such schemas, a few hundred lines of high-level code can be used to generate many thousands of lines of complex low-level code that is extremely difficult to develop and is highly prone to errors. For example, new techniques for high-speed constraint propagation were introduced through the compilation process from high-level constraint descriptions to on-line propagation code. This provides the designer with greater flexibility, and also allows the system to be custom-made and generated quickly. In addition, the levels of representation were extended to include physical mask layout, whereas previously, only symbolic layout was represented.

A client interface to the original system was developed and used with SPICE for a variety of teaching applications. For example, several different integrated circuit inverter configurations can be introduced and manipulated while almost instantaneously comparing the results available from the system. The SPICE interface is arranged so that whenever any aspect of the circuit representation is changed, a new SPICE simulation is automatically triggered and displayed. This kind of capability is extremely useful for design, and it can exploit the fact that the design can be manipulated at several different levels of representation in such a way that the database system automatically propagates these changes to the circuit level of representation, where they are utilized by SPICE. Thus, for example, the user can manipulate a layout in such a way that changes in the circuit representation will immediately yield new SPICE waveform representations. The user is also able to attach both current and voltage probes to any point in the network, and multiple probes are easily accommodated. Experience has been gained from classroom use, and it has been shown that a greater understanding of circuit behavior can be acquired from the interactive use of this system.

Attention has also been directed to the utilization of this database in relation to existing CAD systems. Feedback from industrial users has indicated the need for a client interface to existing CAD systems, as well as adherence to today's evolving industry standards for database capabilities and interface styles, such as those introduced by the CAD Framework Initiative. Adherence to these specifications increases the portability of the system and its overall usefulness. Consideration is also being given to the possibility of using a commercial object-oriented database in order to provide long-term data storage functionality. Nevertheless, the original database design for this project has demonstrated its ability to provide satisfactory performance in the interactive mode, which has not always been possible with commercially based systems. There are existing languages to express high-level

---

8 C.H. Tan and J. Allen, "Minimization of Power in VLSI Circuits Using Transistor Sizing Input Ordering, and Statistical Power Estimation," paper to be presented at the International Workshop on Low Power Design, Napa, California, April 24-27, 1994.

9 R.C. Armstrong and J. Allen, "FICOM: A Framework for Incremental Consistency Maintenance in Multi-Representation, Structural VLSI Design Databases," *Digest of Technical Papers,* IEEE/ACM International Conference on Computer-Aided Design, Santa Clara, California, November 8-12, 1992, pp. 336-343.

schemas, such as Express, and the utilization of these languages is currently being studied.

In the past, true single-phase clock circuits were studied in Ms. Lisa A. Pickelsimer's thesis,[10] in terms of understanding the high-speed behavior of these circuits and generating a design methodology to systematically provide optimal designs in this circuit style. Now, an additional circuit style is being investigated that shows potential for extremely high performance. Self-resetting CMOS circuits are an interesting class of circuits that can provide for very high speed and utilization within a wave pipelining discipline. These circuits do not utilize level signals (except at clock transitions) and instead generate pulses; the leading edge of which is generated by the normal circuit transition. The falling edge is generated by a rapid feedback loop that automatically limits the length of these pulses. While interesting examples of these self-resetting circuits have been demonstrated in the literature, a detailed investigation of their general properties, configurations, design methodology, and potential applications is currently underway in graduate student Robert G. McDonald's thesis project.

High-performance computing performance for CAD continues to be provided through a variety of configurations. These include the networking of several processors together, using the parallel virtual machine (PVM) methodology, as well as new scalable parallel architectures that provide for a tighter coherence between the existing (workstation-like) processors. Work is underway to acquire such a processor, which will be useful for many CAD algorithms that exhibit substantial parallelism, as well as the support of multiple client utilization of CAD databases with extremely fast communication between processors.

1.1.1 Meeting Paper

Baltus, D.G., and J. Allen. "Efficient Exploration of Space-Time Transformations for Systolic Array Synthesis." Paper presented at the International Conference on Application-Specific Array Processors, Venice, Italy, October 25-27, 1993.

1.2 Analog VLSI Systems for Integrated Image Acquisition and Early Vision Processing

Sponsor

National Science Foundation
Grant MIP 91-17724

Project Staff

Professor Berthold K.P. Horn, Professor Hae-Seung Lee, Professor Charles G. Sodini, Professor Jacob K. White, Professor John L. Wyatt, Jr., Dr. Craig L. Keast, Dr. Ichiro Masaki, Steven J. Decker, Lisa G. Dron, Ibrahim M. Elfadel, J. Mikko Hakkarainen, Frederick P. Herrmann, David R. Martin, Ignacio S. McQuirk, Mark W. Reichelt, Mark N. Seidel, Luis M. Silveira, Christopher B. Umminger, Paul C. Yu

1.2.1 Project Summary

In real-time machine vision the sheer volume of image data to be acquired, managed and processed leads to (1) communications bottlenecks between imagers, memory, and processors, and (2) very high computational demands. We are designing and testing experimental *analog* VLSI systems to overcome these problems. Our goal is to determine how the advantages of analog VLSI—high speed, low power and small area—can be exploited and its disadvantages—limited accuracy, inflexibility and lack of storage capacity—can be minimized. The work is concentrated on *early vision tasks*, i.e., tasks that occur early in the signal flow path of an animal or machine.

The next section details substantial progress on seven different designs, including systems for camera motion estimation, automatic alignment to calibration marks, brightness adaptive imaging, and other tasks. Perhaps the most exciting news concerns the stereo vision system, the initial portion of which was completed by Mikko Hakkarainen[11] in his doctoral work with Professor Hae-Seung Lee last May.

This project has attracted the attention of General Motors, which sent a vision systems engineer, Dr. Ichiro Masaki, to work with us for a year to investigate the potential applications of this type of design

---

[10] L.A. Pickelsimer, *A Structured Design Methodology for Speed-Optimized True Single-Phase-Clock Dynamic CMOS*, S.M. thesis, Dept. of Electr. Eng. and Comput. Sci., MIT, 1992.

[11] J.M. Hakkarainen, *A Real Time Stereo Vision System in CD/CMOS Technology*, Ph.D. diss., Dept. of Electr. Eng. and Comput. Sci., 1992.

to future use in automobiles. In directing our attention to practical automotive applications, Dr. Masaki has made a significant contribution to this project.

1.2.2 Vision Chip Designs

BiCMOS/CCD Focus-of-Expansion Chip for Camera Motion

Graduate student Ignacio McQuirk, working with Professors Berthold Horn and Hae-Seung Lee, has continued development and testing of a chip that determines the direction of camera translation directly from a time-varying image sequence with a real-time analog VLSI system. Mr. McQuirk's approach assumes that a camera moving through a fixed world has translational velocity only. The focus of expansion (FOE) is the projection of the camera translation vector onto the image plane and hence gives the direction of camera motion when the camera parameters are known. The FOE is the image of the point towards which the camera is moving and the point from which other image points appear to be receding. Various possible algorithms for estimating the FOE have been compared in terms of accuracy, robustness, and feasibility for single-chip analog VLSI implementation with on-chip imaging circuitry. An algorithm based on minimizing the sum of squares of the differences at stationary points between the observed time variation of brightness and the predicted variation was chosen for implementation. Mr. McQuirk designed a BiCMOS/CCD chip for estimating the FOE using on-chip CCD imaging circuitry and a row-parallel processing scheme. A complete 64×64 version of this chip has been fabricated through MOSIS.

Mr. McQuirk is currently in the process of building a test setup for the FOE chip. From an electrical standpoint, this involves the completed circuit design and subsequent fabrication of a printed circuit board with the support components required for operation of the FOE system in all its various testing and operational modes. This board consists of pattern generators, CCD clock drivers, FOE chip bias circuitry, I/O drivers, a Feedback Control Processor, and an interface to the Vision Chip Test System. The Vision Chip Test System itself will be upgraded to handle the multiple channels necessary for the FOE chip I/O. This will permit us to test the FOE chip with electrical data, allowing for algorithmic, system, and circuit level characterization. For operation using optical input through the on-chip imager, Mr. McQuirk is also in the process of completing a setup which will optically present the FOE chip with a real-time image sequence due to 3-D motion. This will be accomplished using a flexible fiber optic image guide to transmit the image from a lens system mounted on a precision linear track to the FOE test board. This additionally requires the imager on the FOE chip to be calibrated to ascertain accurately the projection from the 3-D motion to the FOE location.

Brightness Adaptive CCD Imager with Analog and Digital Output

Graduate student Steven Decker, working with Professor Charles Sodini, has been designing an imager to act as the front end of a modular vision system. It will feature a 64×64 array of pixels, large fill factor, operating speed on the order of 1000 frames/second, column parallel ouput, brightness adaptation capability, and user selectable analog or 10-bit digital output.

The cyclic A/D converter is based on an algorithm which requires only addition, division by two, and comparison of two signals. The first two operations are performed in the charge domain because the input analog signal is a charge, and CCDs can add and evenly divide charge packets with good precision using little area. Charge packet splitting accuracy is greatly increased by an algorithm which passes each half-packet through the charge splitter again and appropriately recombines the resultant quarter-packets. Analog output can be obtained by shifting the signal charge onto an integrating capacitor in a feedback loop around an operational amplifier. The linearity of the analog output is much better than for a conventional CCD output because the poly-poly capacitor in the feedback loop is much more linear than the diffusion capacitance of a conventional CCD output.

The brightness adaptive imaging element is based on a design that was previously patented elsewhere. Each pixel is composed of an MOS capacitor biased in deep depletion and an overflow gate. The overflow gate establishes a barrier to the flow of imaged charge from the main potential well to a sink. During the integration period, the height of the barrier $\phi(t)$ is varied as a function of time. A relationship between the light intensity and final imaged charge can be derived for a given $\phi(t)$. In particular, it is possible to obtain a logarithmic compression curve so that, ideally, any light intensity can be sensed without saturating the imager. The advantages of this technique over other schemes which have been considered are simplicity, flexibility, and small size. Professor Wyatt has developed a method to obtain the required $\phi(t)$ to achieve a given compression curve.

Efforts have been directed towards designing a practical realization of the brightness adaptive

imager. In the original patent, $\phi(t)$ is assumed to be a continuous, infinitely precise function of time. The implications of approximating $\phi(t)$ with a digital generated staircase function have been examined. Using the digital generation scheme and digital output, reconstruction of the original light intensity from the imaged charge is easy since this can be done digitally. Digital generation of $\phi(t)$ also makes it possible to alter the compression function in real time based on previously imaged frames.

A test chip has been designed and layed out. This chip has a 4×4 array of wide dynamic range pixels, four adjacent A/D converters, and four adjacent analog output circuits. The chip is currently being fabricated in the MIT CCD/CMOS 1.75 micron process. A test setup is being designed concurrently.

Single-Chip Alignment Sensor

Graduate student Christopher Umminger, working with Professors Charles Sodini and Berthold Horn, has been researching implementations of single-chip sensors to measure alignment error for tasks in automated manufacturing. These sensors will optically acquire an image of the alignment marking and output a measure of the misalignment. Currently, they have settled on two methods for the task of aligning a mark to the sensor. The first is based on the surveyor's mark. The sensor consists of a four quadrant photodiode combined with an edge detection circuit to provide extra alignment accuracy in the face of nonuniform illumination. The second sensor design uses a square grating mark. A position sensitive moiré fringe pattern formed between the sensor and the marking is used to achieve accurate alignment.

The first sensor design using the surveyor's mark was fabricated in 1992. Testing of the imaging photodiodes and associated amplification circuits has shown that they function as expected. In July 1993, a paper was presented on this chip at the SPIE Annual Meeting in the session on optical alignment. A second version of the sensor which significantly improves its speed was submitted for fabrication in August 1993. Design and layout of the moiré fringe sensor was finished in April and the fabricated chip has just been received. Testing on this chip will occur in the fall. A mock alignment system is being constructed to characterize these sensors in an alignment operation.

Content-Addressable Parallel Digital Image Processor

Massively parallel associative processors may be well suited as coprocessors for accelerating machine vision applications. They achieve very fine granularity, as every word of memory functions as a simple processing element. A dense, dynamic, content-addressable memory cell supports fully-parallel operation, and pitch-matched word logic improves arithmetic performance with minimal area cost. An asynchronous reconfigurable mesh network handles interprocessor communication and image input/output, and an area-efficient pass-transistor circuit counts and prioritizes responders.

Graduate student Frederick P. Herrmann has completed the design of a 16K-trit associative processing chip, under the supervision of Professor Charles Sodini. The chip integrates 256 processing elements, each with 64 trits of memory. A 16×16 mesh network connects the processing elements, and larger meshes may be constructed using several chips. Fabrication in MIT's 1.5 μm CCD/CMOS process was completed in February. Some subsystems were tested and their functionality was verified. High leakage current on that run prevented complete testing. The design has been resubmitted to MIT and to MOSIS, with chips expected in September.

Mr. Herrmann has also been working with Jeffrey Gealow and Lawrence Hsu (both with the MIT Associative Processing Project) on a software development system for the associative processor. Simulators for the parallel processor array and its sequential controller were demonstrated in July, and the system is now being used to prepare test vectors for the associative processing chip.

Video Rate Analog-to-Digital Conversion

Graduate student Paul Yu, working with Professor Hae-Seung Lee, has been working on an operational amplifier that uses a replica-amplifier gain-enhancement technique. The op-amp is a key component of a pipelined A/D converter. The replica-amp technique can be generally applied to any conventional op-amp topologies such as folded-cascode and class-AB, single-ended or differential. Unlike conventional techniques such as cascoding which increase gain by increasing the output resistance, the replica-amp technique increases gain by matching main and replica amps. Among the advantages of the replica-amp technique are low supply, high swing, and effectiveness in driving resistive loads. A two-stage topology without any cascode was chosen to maximize the available output swing. The experimental result of this low-

supply, high swing, CMOS op-amp was presented at the 1993 International Solid-State Circuit Conference.

An analytical study has been conducted with regard to the settling time of an op-amp using this replica-amp gain enhancement technique. It was shown that this technique has only a small effect on the settling time. This analytical result is confirmed by a SPICE simulation as well as by the experimental results. The results of this study were recently submitted to IEEE Circuits and Systems II.

In view of the trend toward ever decreasing supply voltages for power minimization and the fact that conventional techniques do not work well at low supply, the replica-amp technique offers a promising alternative.

Our attention has shifted toward the capacitor mismatch problem which contributes to gain-error in a pipelined A/D converter. A simulation program has been written in C programming language to investigate a new technique of minimizing differential nonlinearity (DNL). A Monte Carlo simulation has been carried out to determine the minimum capacitor size necessary under this new technique. It appears that the technique can lead to a low-power pipelined A/D converter operating at video speeds.

Design of Specialized Analog and Digital VLSI Processors for Computing the Relative Orientation Between Two-Camera Systems

Computing the relative orientation between the coordinate systems of two cameras is an important problem for calculating depth from binocular stereo and for determining general camera motion. Graduate student Lisa Dron, working with Professor Berthold Horn, has continued work on the design of a complete, autonomous system for computing relative orientation in real-time with specialized analog and digital VLSI hardware. Such a system would be suitable for mounting on mobile or remote platforms that cannot be tethered to a computer and for which the size, weight and power consumption of the components are critical factors. The problem of computing relative orientation is difficult for two reasons: one is the nonlinearity of the motion equations which must be solved; and the other is the difficulty of determining point correspondences between two images. It is the latter problem which entails the largest I/O and computational requirements and where specialized processors will offer the greatest benefit. Methods already exist for finding point correspondences in software which can be implemented on available general purpose digital image processors. However, the resulting systems are still too large, slow and power-hungry for many applications.

This research consists of two parts. The first, which is now essentially completed, was to develop the theoretical foundations of the relative orientation system. This involved developing and adapting algorithms for determining point correspondences and solving the motion equations which are both robust as well as simple enough to be easily implemented in hardware. It also involved determining the limits of system operation as a function of error in the correspondences, field of view, and image resolution. The second part of the research, which is currently underway, is to design, fabricate and test prototype chips for the specialized processors which will be used to find the point correspondences. Two separate processors are needed: one which computes a binary edge map from the input image data, and the other which determines translational offsets between patches of the edge maps from two different images. Fabrication of the prototype circuits will be done through MOSIS.

The design for the edge detection chip is based on a variation of the CCD circuit developed by Dr. Craig Keast for the focal-plane imaging and segmentation chip, which has been modified to implement the multiscale veto algorithm described in the previous progress report. An important problem to be solved, which was one of the major difficulties encountered by Dr. Keast, is to find an effective method for computing on chip small differences between neighboring pixel values with at least 6 bits of dynamic range. Three test chips containing one- and two-dimensional arrays, as well as their isolated subcircuits, have been fabricated and tested based on an initial design. These circuits did not reach the desired performance level of 6 bits of dynamic range, however, largely due to offsets in the absolute-value-of-difference circuits. An improved design which should reduce the effect of offsets is currently being studied, and a new set of test chips will soon be sent to MOSIS for fabrication. If the tests of these new chips are satisfactory, a full 32×32 array will be sent out for fabrication by the end of the year.

The second processor receives the binary edge maps produced by the edge detection chips and determines point matches using a block correlation method. Because the input is binary, the matching circuits require relatively little hardware and can be implemented with purely digital logic. On the other hand, because the computations are simple and do not require high precision, there may be a net benefit in terms of silicon area in using analog processing for parts of the circuit. In order to test this idea, a design incorporating analog elements has been formulated for the core of the matching circuit.

A test chip based on this design will be fabricated and tested and the result will be compared with an entirely digital implementation.

Switched-Capacitor System for Merging Depth and Slope Estimates

Graduate student Mark Seidel, working with Professors John Wyatt and Thomas Knight, has been designing, fabricating, and testing a switched-capacitor chip to robustly compute the depth of a scene by merging depth estimates (possibly from a stereo algorithm) with surface slope estimates (possibly from a shape-from-shading algorithm). This multisensor image fusion design uses the least-squares properties of certain switched-capacitor systems to compute a dense depth map given noisy or sparse depth and slope input. A MOSIS Tinychip containing some subcircuits of a 1-D version was fabricated and tested; all of the subcircuits were operational. A final 1-D version of the network was designed, fabricated, and partially tested. The protection circuits test positively, as does the connectivity on-chip. The remainder of the testing will be completed before the end of the year.

1.2.3 Mixed Circuit/Device Simulation for CMOS/CCD Systems

Graduate student Mark Reichelt, under the supervision of Professor Jacob K. White, and in collaboration with Professor Andrew Lumsdaine of the University of Notre Dame, is working on algorithms for accelerating mixed circuit/device simulation, as this kind of analysis is required to properly simulate the performance of CCD-based analog vision circuits. Initially, they have focused on developing parallelizable algorithms for semiconductor device transient simulation, and are investigating waveform relaxation (WR), a natural candidate (This work is also partly funded by the ARPA VLSI contract N00014-91-J-1698).

Part of this work is focused on more theoretical aspects of applying parallel WR to device transient simulation. As the WR algorithm decomposes systems of differential equations into subsystems which can be solved independently, the algorithm has easily exploited parallelism. And when using WR on a parallel computer, it is likely that the different subsystems will be assigned to different processors. In that case, communication can be avoided, and overall computational efficiency improved, if each of the separated subsystems can be solved with independently determined timesteps. Using different timesteps for different subsystems implies that a *multirate integration method* has been used to solve the system, and there is no guarantee that this multirate method is stable. In the past year, we have been able to prove that certain multirate methods are A-stable for several model problems associated with semiconductor device simulation.

The more practical part of this work was directed toward accelerated waveform relaxation techniques for parallel semiconductor device simulation. We adapted both the generalized conjugate-direction method and a convolution successive over-relaxation (CSOR) algorithm to performing parallel semiconductor device simulation, and compared the two approaches using both a cluster of workstations and a 32-node i860-based hypercube. Both algorithms are competitive with standard algorithms on a single processor, but achieve nearly linear speed-up on the hypercube. Mark Reichelt finished his doctoral dissertation on the subject of convolution SOR and also received a best student paper award at the SIAM sponsored Copper Mountain Conference on Multigrid methods.

1.2.4 Hybrid Analog/Digital Vision Chips for Intelligent Vehicle Applications

Dr. Ichiro Masaki, a visiting scientist from General Motors, is studying the feasibility of applying hybrid analog/digital vision chip technologies to intelligent vehicle systems in collaboration with Professors Hae-Seung Lee, Charles Sodini, John Wyatt, Berthold Horn, and their students, David Martin and Steve Decker. The demand is increasing rapidly for developing intelligent vehicles that are integral to safe, efficient, and environmentally friendly traffic systems. The technical characteristics of the intelligent vehicle applications match well with the natures of hybrid analog/digital vision chip technologies. Intelligent vehicles need high-speed processing, low mass-production system costs, low power consumption, and small physical sizes. Large development costs will be compensated by the large market size and limited accuracy can be accepted. Examples of applications include on-vehicle systems such as a collision avoidance system and on-road systems including a traffic monitoring system.

A stereo vision system is being developed as the first case study. Installed in a vehicle, the stereo vision system will measure the distance from the vehicle to a vehicle in front of it. The goal is to keep a reasonable distance between the vehicles by controlling a vehicle's speed automatically. The system includes two custom chips: a wide dynamic range imager chip and a hybrid analog/digital array processor chip. The intensity dynamic range of

typical conventional television cameras is large enough for in-door applications but does not cover various weather conditions on highways. Our imager chip will have enough dynamic range for vehicle applications. The array processor chip will carry out edge detection and stereo matching operations using both analog and digital circuits. The chip is expected to require less silicon area, lower power consumption, and lower mass-production cost compared to conventional digital microprocessors. Small-scale test chips are being developed for feasibility studies.

Analog versus Digital Approaches to Machine Vision Hardware

Recent technical progress in analog vision chips has stimulated comparisons between analog and digital processing schemes through case studies. The analog vision chips have demonstrated their small cost in silicon area and high speed characteristics for some applications. The goal of this project is to identify the capabilities and limitations of the analog chips and to find what application areas would be suitable for analog, digital, and hybrid chips, respectively.

In collaboration with Professors Charles Sodini, Berthold Horn, and John Wyatt, Dr. Ichiro Masaki and graduate student David R. Martin, under the supervision of Professor Hae-Seung Lee, are studying stereo vision algorithms. These algorithms are frequently used in various applications including intelligent vehicles and flexible manufacturing systems.

The chip implementation level comparison is underway. Currently a cell in a Multi-Instruction Multi-Data Array Processor architecture, implemented in both analog and digital, is being designed. Each processing element includes an arithmetic unit, a data flow control unit, and a system control unit which includes program memory. Several analog circuit issues for the accuracy of the cell's arithmetic unit have been resolved.

Our next goals include finalization of the chip design, layout using the Cadence design tools, fabrication of the chips, and performance comparisons based on both simulated and actual chips. The criteria for the comparisons will include cost in silicon area, processing speed, manufacturing and development costs, testability, and other related issues.

1.2.5 Solutions of Theoretical Problems in Analog Hardware and Machine Vision

Neural Network Theory

Working under the guidance of Professor John Wyatt, graduate student Ibrahim Elfadel has investigated a generalization of the effective energy function, first introduced by Peterson and Soderberg, and used it to solve constrained optimization problems within the ansatz of saddle-point methods in statistical physics. Mr. Elfadel has shown that within this paradigm, a smooth *convex potential function* is associated with each closed bounded constraint set. He has also shown that using the conjugate of the convex potentials, it is possible to derive, from the effective energy, a cost function that is a natural generalization of the analog Hopfield energy. Descent dynamics and deterministic annealing can then be used to find the global minimum of the original minimization problem. When the conjugate is hard to compute explicitly, we show that minimax dynamics can be used to find the extrema of the effective energy. This general effective-energy framework has allowed Mr. Elfadel to tie up different cost functions that were used previously to solve optimization problems involving local competition of the winner-take-all type. Moreover, he has used the same general framework to propose continuous dynamical systems for finding the fixed points of "rotor" neural networks.

Reduced-Order Models for Interconnect Simulation

To analyze the effects of packaging parasitics on circuit performance, these parasitics are frequently modeled with equivalent circuits which are then included in a circuit simulation. Complicated packages or interconnects can generate extremely large equivalent circuits which must then be simplified, or reduced, for efficient use in a circuit simulator. The most commonly used Padé, or moment-matching, approaches to performing this model order reduction can generate unstable equivalent circuits. To overcome the instability problems, Graduate Students Luis Miguel Silveira and Ibrahim Elfadel, under the guidance of Professor Jacob White, have been investigating a state-space algorithm for model-order reduction based on truncating a balanced system realization. This algorithm has the important characteristic that the reduced order model resulting from the truncation is *guaranteed* to be stable. The simulation of simple RLC circuits modeling transmission lines has demonstrated that the balanced-realization approach generates accu-

rate, stable reduced-order models, whereas the Padé method does not.

Switched-Capacitor Network Theory

Graduate student Mark Seidel, working with Professor John Wyatt, has been studying general networks of switches, capacitors, and sources. These networks are important in the implementation of computational and image processing problems formulated as minimization algorithms. Discrete-time system stability was proved for a large subclass of these networks using contraction mappings; this subclass includes many useful image processing networks, including the SC implementation of the coupled depth and slope network. The explanation of derived and inverted bounding expressions and estimates, along with the important limiting case of large N parameters, was accepted and presented at the European Conference on Circuit Theory and Design.

Under constant input conditions, the switched-capacitor networks can be shown to be generally equivalent to networks comprised of nonreciprocal linear passive resistive multiports and multiterminal elements. Each such element is associated with exactly one of the capacitors in the network. The transformation to these resistive networks results in a resistive network equivalent to the original SC network in the sense that all node voltages are the same, and all currents are equal to the corresponding average transferred charges per period. This equivalence indicates certain rules for the reverse problem of synthesizing a switched capacitor network from a resistive network. This synthesis can always be accomplished for networks of two-terminal resistors and lossy two-port transformers. It has also been shown that even if the resistive network to be implemented includes only the nonreciprocal passive resistive elements discussed earlier, clock phase assignment is sometimes impossible. A conference paper explaining this linear analysis and synthesis of SC networks was accepted and presented at the European Conference on Circuit Theory and Design.

Constant-input switched-capacitor network steady-state analysis and synthesis has been extended to two classes of networks containing certain nonlinear SC elements. This theoretical extension now allows for the analysis of realistic SC implementations, including the reverse-biased-diodes nonlinear parasitic capacitances associated with the MOS switches. More importantly, this extension also allows for implementing arbitrary networks of two-terminal voltage-controlled conductances, two-terminal positive linear resistors, and lossy two-port transformers. An example of this kind of network is the one- and two-dimensional resistive fuse grid. The new networks explain certain existing nonlinear SC networks, as well as offering a synthesis method leading to a resistive fuse network with significant performance improvement. Furthermore, the nonlinear elements are dynamically controlled by two global wires, thereby making continuation method implementation feasible. A conference paper presenting the nonlinear synthesis results was presented at the Midwest Symposium on Circuits and Systems.

First and second drafts of the majority of Mr. Seidel's thesis document were completed. This portion of the thesis covers bounds and estimates theory and develops linear and nonlinear analysis and synthesis theory. The thesis will be completed before December 1994.

1.2.6 Publications

Journal Articles

Dron, L. "The Multi-Scale Veto Model: A Two-Stage Analog Network for Edge Detection and Image Reconstruction." *Int. J. Comput. Vision* 11(1): 45-61 (1993).

Elfadel, I. "Convex Potentials and their Conjugates in Analog Mean-Field Optimization." Submitted to *Neural Computation.*

Elfadel, I.M., and R.W. Picard. "Gibbs Random Fields, Co-Occurrences, and Texture Modeling." *IEEE Trans. Pattern Anal. Machine Intell.* Forthcoming.

Elfadel, I.M. and A.L. Yuille. "Mean-Field Phase Transitions and Correlation Functions for Gibbs Random Fields." *J. Math. Imaging Vis.* 3: 167-186 (1993).

Hakkarainen, J.M., and H.-S. Lee. "A 40×40 CCD/CMOS Absolute-Value-of-Difference Processor for Use in a Stereo Vision System." *IEEE J. Solid-State Circuits* 26(7): 799-807 (1993).

Keast, C.L., and C.G. Sodini. "A CCD/CMOS-Based Imager with Integrated Focal Plane Signal Processing." *IEEE J. Solid-State Circuits* 28(4): 431-437 (1993).

Masaki, I. "ASIC Approaches for Vision-based Vehicle Guidance." submitted to a special issue on Automotive Electronics, *Trans. Inst. Electron., Info. and Commun. Eng.* December 1993.

Reichelt, M., J. White, and J. Allen. "Optimal Convolution SOR Acceleration of Waveform Relaxation with Application to Parallel Simulation of Semiconductor Devices." Submitted to *SIAM J. Sci. Stat. Comp.*

Yu, P.C., and H.-S. Lee. "Settling-Time Analysis of a Replica-Amp Gain Enhanced Operational Amplifier." Submitted to *IEEE Trans. Circuits Syst. II: Analog Digit. Signal Process.*

Book Chapters

Elfadel, I.M., and J.L. Wyatt, Jr. "The "Softmax" Nonlinearity: Derivation Using Statistical Mechanics and Useful Properties as a Multiterminal Analog Circuit Element." *Advances in Neural Information Processing Systems,* Morgan Kaufman, 1994, vol. 6. Forthcoming.

Masaki, I. "Fuzzy Logic Technology and the Intelligent Highway System (IHS)." *Fuzzy Logic Technology and Applications I.* IEEE Press. Forthcoming.

Conference Papers

Dron, L. "Dynamic Camera Self-Calibration from Controlled Motion Sequences." *Proceedings of the DARPA Image Understanding Workshop,* April 1993, pp. 387-396; also *1993 Conference on Computer Vision and Pattern Recognition,* New York, New York, 1993, pp. 501-506.

Elfadel, I.M. "Global Dynamics of Winner-Take-All Networks." *Proceedings of the SPIE Conference on Neural and Stochastic Methods in Image and Signal Processing,* San Diego, California, July 1993, vol. 2032, pp. 127-137.

Elfadel, I.M. "Integration of Shape Deformation and Motion." Submitted to the IEEE Conference on Computer Vision and Pattern Recognition (CVPR '94).

Elfadel, I.M. "Lyaponov Functions for Rotor Neural Networks." The American Control Conference (ACC'94). Forthcoming; also submitted to the Conference on Neural Information Processing Systems, Denver, Colorado, December 1993.

Elfadel, I.M. "On the Convergence of Relaxation Labeling Processes." Submitted to the *Third European Conference on Computer Vision (ECCV'94).*

Elfadel, I.M., and J.L. Wyatt, Jr. "The 'Softmax' Nonlinearity: Derivation Using Statistical Mechanics and Useful Properties as a Multiterminal Analog Circuit Element." Paper presented at Neural Information Processing Systems, December 1993.

Harris, J.G. "Analog Chips and Intelligent Vehicles." Paper presented at the *Intelligent Vehicles '93,* Tokyo, Japan, July 1993, pp. 13-18.

Lumsdaine A., and M. Reichelt. "Waveform Iterative Techniques for Device Transient Simulation on Parallel Machines." *Proceedings of the SIAM Meeting on Parallel Processing for Scientific Computing,* Norfolk, Virginia, March 1993, pp. 237-245.

Masaki, I. "Function-Oriented Chip Approach for Real-Time Vision." *Proceedings of the IEEE IECON Conference,* Hawaii, November 1993, pp. 1610-1615.

Masaki, I. "Three Dimensional Vision Systems for Intelligent Vehicles." *Proceedings of the IEEE IECON Conference,* Hawaii, November 1993, pp. 1712-1717.

Reichelt, M. "Optimal Convolution SOR Acceleration of Waveform Relaxation with Application to Semiconductor Device Simulation." *Proceedings of the Copper Mountain Conference on Multigrid Methods,* Copper Mountain, Colorado, April, 1993, pp. 513-526.

Reichelt, M., A. Lumsdaine, J. White, and J. Allen. "Accelerated Waveform Methods for Parallel Transient Simulation of Semiconductor Devices." *Proceedings of the International Conference on Computer-Aided Design,* Santa Clara, California, November 1993, pp. 270-274.

Reichelt, M., F. Odeh, and J. White. "A-Stability of Multirate Integration Methods, with Application to Parallel Semiconductor Device Simulation." *Proceedings of the SIAM Meeting on Parallel Processing for Scientific Computing,* Norfolk, Virginia, March 1993, pp. 246-253.

Seidel, M.N. "Analysis and Synthesis of Passive Linear Steady-State SC Networks for Image Processing." *Proceedings of the 11th European Conference on Circuit Theory and Design,* Davos, Switzerland, August 30-September 3, 1993, pp. 929-934.

Seidel, M.N. "Nonlinear SC Network Analysis and Synthesis for Image Processing." *Proceedings of the 36th Midwest Symposium on Circuits and Systems,* Detroit, Michigan, August 1993. Forthcoming.

Seidel, M.N., and J.L. Wyatt, Jr. "Step-Response Bounds for a Certain Class of Discrete-Time Systems." *Proceedings of the 11th European Conference on Circuit Theory and Design,* Davos, Switzerland, August 30 - September 3, 1993, pp. 1521-1526.

Silveira, L.M., I.M. Elfadel, and J.K. White. "A Guaranteed Stable Order Reduction Algorithm for Packaging and Interconnect Simulation." Paper presented at the *Second Topical Meeting on Electrical Performance and Electronic Packaging,* Monterey, California, October 1993, pp. 165-168.

Silveira, L.M., I.M. Elfadel, M. Chilukuri, K. Kundert, and J.K. White. "Efficient Frequency-Domain Modeling and Circuit Simulation of Transmission Lines." Submitted to the 31st Design Automation Conference.

Umminger, C.B., and C.G. Sodini. "Alignment Sensor with Integrated Analog Signal Processing." Optical Alignment. *SPIE Proc.* 1996: 2-9 (1993).

Wyatt, J.L., and I.M. Elfadel. "On the Solution of Oja's Equations." Submitted to the 1994 Snowbird Conference on Neural Networks, Snowbird, Utah.

Yu, P.C., and H.-S. Lee. "A High-Swing 2V CMOS Operational Amplifier with Gain Enhancement using a Replica Amplifier." Paper presented at the *1993 IEEE International Solid-State Circuit Conference,* San Francisco, California, February 1993, pp. 116-117.

Patents

Decker, S., H.-S. Lee, and J.L. Wyatt, Jr. *Resistive Fuse Circuits for Image Segmentation and Smoothing.* Patent No. 5223754, June 29, 1993.

Theses

Elfadel, I.M. *From Random Fields to Networks.* Ph.D. diss., Dept. of Mech. Eng., MIT, February 1993.

Reichelt, M.W. *Accelerated Waveform Relaxation Techniques for the Parallel Transient Simulation of Semiconductor Devices.* Ph.D. diss., Dept. of Electr. Eng. and Comput. Sci., MIT, June 1993.

Stein, G.P. *Internal Camera Calibration using Rotation and Geometric Shapes.* S.M. thesis, Dept. of Electr. Eng. and Comput. Sci., MIT, February 1993.

Publicity

Hapgood, F. "The Vision Chip." *Up the Infinite Corridor.* Reading, Mass.: Addison-Wesley, 1993, pp. 173-192.

Internal Publications

Elfadel, I.M. *From Random Fields to Networks.* RLE TR No. 579. Cambridge: MIT Research Laboratory of Electronics, 1993.

Elfadel, I.M., and R.W. Picard. *Gibbs Random Fields, Co-Occurrences, and Texture Modeling.* Perceptual Computing TR No. 204. Cambridge: MIT Media Laboratory, 1993.

Elfadel, I.M., M. Kaneko, A. Koike, and Y. Hatori. *Facial Feature Detection Using Multiresolution Template Matching.* Telematics Group Memo No. 3/4. Cambridge: KDD Laboratories, 1993.

Abstracts for Talks

Reichelt, M., and J. White. "Waveform Frequency-dependent SOR Methods for Parallel Semiconductor Device Simulation." *IBM Technical Seminar,* Yorktown, New York, March 1993.

1.3 Parallel Algorithms for Device Simulation

Sponsors

Advanced Research Projects Agency/
National Science Foundation
Grant MIP 91-17724
IBM Corporation
U.S. Navy - Office of Naval Research
Grant N00014-91-J-1698

Project Staff

Dr. Andrew Lumsdaine, Khalid Rahmat, Mark W. Reichelt, Professor Jacob K. White, Professor Jonathan Allen

The growing importance of mixed circuit/device simulation, its enormous computational expense, as well as the increasing availability of parallel computers, have made the development of specialized, easily parallelized, algorithms for transient simulation of MOS devices necessary. In earlier work on the Waveform Overrelaxation Device Simulator

(WORDS) program, the easily parallelized waveform relaxation (WR) algorithm was shown to be a computationally efficient approach to device transient simulation even on a serial machine. The WR algorithm, though, typically requires hundreds of iterations to achieve an accurate solution.

To use WORDS in a mixed circuit/device simulator, we have been investigating ways of making the program more robust and efficient. We determined how to compute the terminal currents accurately using different timepoints at different mesh nodes. We also improved the timestep selection procedure by determining how to refine the timesteps as WR iterations proceed and reducing the total computation by as much as a factor of two by using only a few coarse timesteps in early iterations. The more accurate, electric field dependent mobility model was also implemented. Recent work on theoretical aspects of these methods have answered several long-standing questions about multirate stability.[12]

We found experimental evidence that WR using standard overrelaxation acceleration can produce oscillatory results, and are investigating methods for eliminating this phenomenon. A frequency-dependent overrelaxation algorithm using lowpass filtering was developed,[13] as well as a waveform conjugate-direction approach.[14] Experimental results indicate that both approaches reduces the number of waveform iterations required by more than a factor of seven. Finally, experimental results using a parallel computer show that although the accelerated WR methods are only competitive with the best of the standard algorithms for device transient simulation on a serial machine, the WR algorithms are substantially faster on a parallel machine.[15]

1.4 Numerical Simulation of Short Channel MOS Devices

Sponsors

IBM Corporation
National Science Foundation
Fellowship MIP 88-58764
U.S. Navy
Contract N00174-92-Q-0133

Project Staff

Khalid Rahmat, Professor Dimitri A. Antoniadas, Professor Jacob K. White

The model used in conventional device simulation programs is based on the drift-diffusion model of electron transport, and this model does not accurately predict the field distribution near the drain in small geometry devices. This is of particular importance for predicting oxide breakdown due to penetration by "hot" electrons. There are two approaches for more accurately computing the electric fields in MOS devices, one is based on adding an energy equation to the drift-diffusion model and the second is based on direct solution of Boltzman's equation.

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 is numerically less tame than the standard approach, and must be solved carefully. We have developed a two-dimensional device simulator in which an energy balance equation is solved for electron temperature along with the usual drift-diffusion equations. The program avoids temperature instabilities produced by previous discretization approaches through the use of a careful application of exponential-fitting to the energy equation. Drain currents for silicon MOSFETs predicted by the simulator, using one set

---

[12] M. Reichelt, F. Odeh, and J. White, "A-Stability of Multirate Integration Methods, with Application to Parallel Semiconductor Device Simulation," *Proceedings of the SIAM Meeting on Parallel Processing for Scientific Computing,* Norfolk, Virginia, March 1993.

[13] M. Reichelt, J. White, and J. Allen, "Frequency-Dependent Waveform Overrelaxation for Transient Two-Dimensional Simulation of MOS Devices," *Proceedings of NUPAD IV,* Seattle, Washington, May 1992, pp. 161-166; A. Lumsdaine and J. White, "Accelerating Waveform Relaxation Methods with Application to Parallel Semiconductor Device Simulation," *Numerical Functional Analysis and Optimization,* forthcoming.

[14] A. Lumsdaine and J. White, "Accelerating Waveform Relaxation Methods with Application to Parallel Semiconductor Device Simulation," *Numerical Functional Analysis and Optimization,* forthcoming.

[15] A. Lumsdaine and M. Reichelt, "Waveform Iterative Techniques for Device Transient Simulation on Parallel Machines," *Proceedings of the SIAM Meeting on Parallel Processing for Scientific Computing,* Norfolk, Virginia, March 1993; A. Lumsdaine and J. White, "Accelerating Waveform Relaxation Methods with Application to Parallel Semiconductor Device Simulation," *Numerical Functional Analysis and Optimization,* forthcoming; M. Reichelt, "Optimal Convolution SOR Acceleration of Waveform Relaxation with Application to Semiconductor Device Simulation," Third place, best student paper competition, *Proceedings of the Copper Mountain Conference on Multigrid Methods,* Copper Mountain, Colorado, April 1993.

of model parameters, match well with experimental data for devices over a range of channel lengths from 90 μm to 0.16μm. Also, a method to compute substrate current has been derived which uses the electron temperature provided by the simulator. The computed substrate currents match well with measured data, for the regime above subthreshold, for MOSFET's with channel lengths as short as 0.16μm.[16]

The limited ability to accurately predict hot-electron effects from an energy-balance based simulator has led to our current investigation. We are now trying to solve the full Boltzman equation, using a spherical harmonics based approach, but only in portions of the MOS device which are likely to create hot carriers.[17]

1.5 Coupled Simulation Algorithms for Microelectromechanical CAD (MEMCAD)

Sponsors

Federal Bureau of Investigation
Contract J-FBI-92-196
National Science Foundation
Fellowship MIP 88-58764

Project Staff

Xuejun Cai, Mattan Kamon, Dr. He Yie, Professor Martin A. Schmidt, Professor Stephen D. Senturia, Professor Jacob K. White

High fabrication costs and increasing microsensor complexity is making computer simulation of the realistic geometries necessary, both to investigate design alternatives and to perform verification before fabrication. At MIT, we are developing a microelectromechanical computer-aided design (MEMCAD) system to make it possible for microsensor designers to easily perform realistic simulations. Carefully selected commercial software packages have been linked with specialized database and numerical programs to allow a designer to easily enter a three-dimensional microsensor geometry and quickly perform both mechanical and electrical analysis. The system currently performs electromechanical analyses, such as calculating the capacitance versus pressure (or force) curve for both a square diaphragm deformed by a differential pressure and can be used to calculate levitation forces in structures as complicated as a comb drive.[18]

To support design of electromechanical structures, we are currently investigating two approaches to combining finite-element mechanical analysis with multipole-accelerated electrostatic analysis. The first method is the obvious relaxation algorithm[19] and the second method is a more sophisticated surface/Newton generalized conjugate-residual scheme.[20] By comparing the two methods, we have demonstrated both theoretically and by example that our surface/Newton-GCR algorithm is faster and more robust than the simplier relaxation scheme.[21]

---

16 K. Rahmat, J. White, and D. Antoniadis, "Simulation of Very Short Channel MOSFET's Including Energy Balance," *IEEE Trans. Comput.-Aided Des.* 12(6): 817-825 (1993).

17 K. Rahmat, J. White, and D. Antoniadis, "A Galerkin Method for the Arbitrary Order Expansion in Momentum Space of the Boltzmann Equation using Spherical Harmonics," *Proceedings of NUPAD V,* Honolulu, Hawaii, June 1994, forthcoming.

18 S.D. Senturia, R.M. Harris, B.P. Johnson, S. Kim, K. Nabors, M.A. Shulman, and J.K. White, "A Computer-Aided Design System for Microelectromechanical Systems (MEMCAD)," *IEEE J. Microelectromech. Syst.* 1(1): 3-13 (1992); J.R. Gilbert, P.M. Osterberg, R.M. Harris, D.O. Ouma, X. Cai, A. Pfajfer, J. White, and S.D. Senturia, "Implementation of a MEMCAD System for Electrostatic and Mechanical Analysis of Complex Structures From Mask Descriptions," *Proceedings of the IEEE Conference on Microelectromechanical Systems,* Fort Lauderdale, February 1993, forthcoming.

19 X. Cai, H. Yie, P. Osterberg, J. Gilbert, S. Senturia, and J. White, "A Relaxation/Multipole-Accelerated Scheme for Self-Consistent Electromechanical Analysis of Complex 3-D Microelectromechanical Structures," *Proceedings of the International Conference on Computer-Aided Design,* Santa Clara, California, November 1993, pp. 270-274; X. Cai, P. Osterberg, H. Yie, J. Gilbert, S. Senturia, and J. White, "Self-Consistent Electromechanical Analysis of Complex 3-D Microelectromechanical Structures using a Relaxation/Multipole Method," *Int. J. Sensors Mater.,* forthcoming.

20 X. Cai, H. Yie, P. Osterberg, J. Gilbert, S. Senturia, and J. White, "A Relaxation/Multipole-Accelerated Scheme for Self-Consistent Electromechanical Analysis of Complex 3-D Microelectromechanical Structures," *Proceedings of the International Conference on Computer-Aided Design,* Santa Clara, California, November 1993, pp. 270-274.

21 H. Yie, X. Cai, P. Osterberg, S. Senturia, and J. White, "Convergence Properties of Relaxation versus the Surface-Newton Generalized-Conjugate Residual Algorithm for Self-consistent Electromechanical Analysis of 3-D Micro-Electro-Mechanical Structures," *Proceedings of NUPAD V,* Honolulu, Hawaii, June 1994, forthcoming.

1.6 Numerical Techniques for Simulating Josephson Junction Arrays

Sponsors

Advanced Research Projects Agency/
Consortium for Superconducting Electronics
Contract MDA972-90-C-0021
National Defense Science and Engineering
Graduate Fellowship
National Science Foundation
Fellowship MIP 88-58764

Project Staff

Joel R. Phillips, Dr. Herre S.J. van der Zant, Professor Terry P. Orlando, Professor Jacob K. White

Vortices play a central role in determining the static and dynamic properties of two-dimensional (2D) superconductors. Artificially fabricated networks of superconducting islands weakly coupled by Josephson junctions are model systems for studying the behavior of vortices. Through simulation, we have discovered that the static properties of vortices in an array of Josephson junctions can be significantly influenced by magnetic fields induced by the vortex currents. The energy barrier for vortex motion is enhanced, nearly doubling for penetration depths on the order of a cell size. Moreover, we have found that correct calculation of the vortex current distribution, magnetic moment, and lower critical field require modeling mutual inductance interactions between all cell pairs in the array. To make numerical simulation of the system with all inductive effects computationally feasible, a novel FFT-accelerated integral equation solver was derived. This algorithm is sufficiently efficient to allow study of large (500 x 500 cells) arrays.[22]

The development of the FFT-accelerated algorithm for inductance effects in Josephson Junction arrays led to a more general result useful for packaging analysis. For the important special case of conductors over non-ideal ground planes, extremely fast FFT-accelerated inductance extraction methods have been developed. The FFT technique is so fast compared to its direct rival method that it reduces analysis time from hours to seconds even for relatively simple problems.[23]

1.7 Efficient 3-D Interconnect Analysis

Sponsors

Digital Equipment Corporation
IBM Corporation
National Science Foundation
Fellowship MIP 88-58764
U.S. Navy - Office of Naval Research
Contract N00014-91-J-1698

Project Staff

Mattan Kamon, Songmin Kim, Keith S. Nabors, Joel R. Phillips, Professor Jacob K. White

We have developed multipole-accelerated algorithms for computing capacitances and inductances of complicated 3-D geometries and implemented these algorithms in the programs FASTCAP and FASTHENRY. The methods are accelerations of the boundary-element or method-of-moments techniques for solving the integral equations associated with the multiconductor capacitance or inductance extraction problem. Boundary-element methods become slow when a large number of elements are used because they lead to dense matrix problems which are typically solved with some form of Gaussian elimination. This implies that the computation grows as n^3, where n is the number of panels or tiles needed to accurately discretize the conductor surface charges. Our new algorithms, which use generalized conjugate residual iterative algorithms with a multipole approximation to compute the iterates, reduces the complexity so that accurate multiconductor capacitance and inductance calculations grow nearly as nm where m is the number of conductors. For practical problems which require as many as 10,000 panels or filaments, FASTCAP and FASTHENRY are more than two orders of magnitude faster than standard boundary-element based programs.[24] Manuals and source code for

---

[22] J. Phillips, H. Van der Zant, J. White, and T. Orlando, "Influence of Induced Magnetic Fields on the Static Properties of Josephson-Junction Arrays," *Phys. Rev. B*, forthcoming.

[23] J.R. Phillips, M. Kamon, and J. White, "An FFT-Based Approach to Including Non-ideal Ground Planes in a Fast 3-D Inductance Extraction Program," *Proceedings of the Custom Integrated Circuits Conference,* San Diego, California, May 1993.

[24] J.R. Phillips, M. Kamon, and J. White, "An FFT-Based Approach to Including Non-ideal Ground Planes in a Fast 3-D Inductance Extraction Program," *Proceedings of the Custom Integrated Circuits Conference,* San Diego, California, May 1993; K. Nabors, S. Kim, and J. White, "Fast Capacitance Extraction of General Three-Dimensional Structures," *IEEE Trans. Microwave Theory Tech.* 40(7): 1496-1507 (1993); K. Nabors and J. White, "Multipole-Accelerated Capacitance Extraction Algorithms for 3-D Structures with Multiple Dielectrics," *IEEE Trans. Circuits Syst.* 39(11): 946-954 (1992); M. Kamon, M. Tsuk, C. Smithhisler, and J. White, "Efficient Techniques for Inductance Extraction of Complex 3-D Geometries," *Proceedings of the International Conference on Computer-Aided*

FASTCAP and FASTHENRY are available directly from MIT.

1.8 Adaptive Gridding Techniques for Multipole-Accelerated Solution of Integral Equations

Sponsors

National Science Foundation
Semiconductor Research Corporation
Contract SRC 93-SJ-360

Project Staff

Michael Chou, F. Thomas Korsmeyer,[25] Professor Jacob White

Finding computationally efficient numerical techniques for simulation of three dimensional structures has been an important research topic in almost every engineering domain. Surprisingly, the most numerically intractable problem across these various discplines can be reduced to the problem of solving a three-dimensional Laplace problem. Such problems are often referred to as potential problems, and application examples include electrostatic analysis of sensors and actuators; electro- and magneto-quasistatic analysis of integrated circuit interconnect and packaging; and potential flow based analysis of wave-ocean structure interaction.

The recent development of extremely fast multipole-accelerated iterative algorithms for solving potential problems has renewed interest in integral equation formulations. In such methods, the fast multipole algorithm is used to implicitly construct a sparse representation of the dense matrix associated with a standard integral equation formulation, and this implicit representation is used to quickly compute the matrix-vector products required in a Krylov subspace based iterative method like GMRES. That the multipole-accelerated approach leads to practical engineeering analysis tools has been demonstrated for electrostatic analysis of integrated circuit interconnect and packaging problems,[26] as well as general Laplace problems.[27]

We are extending our own work on multipole-accelerated iterative methods for potential problems to address automatic error control, which is the key difficultly which prevents such techniques from finding broader use in engineering applications. In particular, we are interested in first investigating the straight-forward problem of how best to include higher-order boundary or volume elements. We are then planning to use such a result, along with other error estimation procedures, to perform automatic element refinement. The spatial hierarchy already introduced by the multipole algorithm and our use of a locally preconditioned iterative scheme make it possible to perform very efficient local element refinement and may allow for an efficient variable element order scheme. Because we have been using realistic examples from a wide variety of engineering disciplines, we can determine the robustness of our procedures.

1.9 Circuit Simulation Algorithms

Sponsors

IBM Corporation
National Science Foundation
Fellowship MIP 88-58764
Grant MIP 91-17724

Project Staff

Luis M. Silveira, Steven B. Leeb, Professor Jacob K. White

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

Design, Santa Clara, California, November November 1992, pp. 438-442; M. Kamon, K. Nabors, J. White, "Multipole-Accelerated 3-D Interconnect Analysis," *Proceedings of the International Workshop on VLSI Process and Device Modeling (VPAD),* Nara, Japan, May 1993, forthcoming.

[25] MIT Department of Ocean Engineering.

[26] K. Nabors, T. Korsmeyer, F.T. Leighton, and J. White, "Multipole Accelerated Preconditioned Iterative Methods for Three-Dimensional Potential Integral Equations of the First Kind," *SIAM J. Sci. Stat. Comput.*, forthcoming.

[27] T. Korsmeyer, D.K.P. Yue, K. Nabors, and J. White, "Multipole-Accelerated Preconditioned Iterative Methods For Three-Dimensional Potential Problems," Boundary Element Methods 15, Worcester, Massachusetts, August 1993.

to exploit only the property that the behavior of such a circuit in a given high frequency clock cycle is similar, 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. Such approaches are known as "envelope-following" algorithms.

In our recent work, we have been trying to make the envelope-following algorithm more robust and efficient by exploiting the fact that the envelope of "quasi-algebraic" components in the solution vector do not need to be computed. An automatic method for determining the quasi-algebraic solution components has been derived; experimental results demonstrate that this modified method reduces the number of computed clock cycles when applied to simulating closed-loop switching power converters.[28]

To help support our work in interconnect analysis, we are developing algorithms for efficient SPICE-level simulation of elements with arbitrary frequency-domain descriptions, such as scattering parameters. That is, an element can be represented in the form of a frequency-domain model or a table of measured frequency-domain data. Our approach initially uses a forced stable decade-by-decade l_2 minimization approach to construct a sum of rational functions approximation, but the approximation has dozens of poles and zeros. This unnecessarily high-order model is then reduced using a guaranteed stable model order reduction scheme based on balanced realizations.[29] Once the reduced-order model is derived, it can be combined with any inherent delay (for transmission line models) to generate an impulse response. Finally, following what is now a standard approach, the impulse response is efficiently incorporated in the circuit simulator SPICE3 using recursive convolution.

1.10 Techniques for Embedded System Design, Formal Verification, and Synthesis for Low Power Dissipation

Sponsors

Advanced Research Projects Agency/
U.S. Navy - Office of Naval Research
Grant N00014-91-J-1698
National Science Foundation
Fellowship MIP 92-58376

Project Staff

Professor Srinivas Devadas, Mazhar M. Alidina, Vishal L. Bhagwati, Stan Y. Liao, José C. Monteiro, W. John Rinderknecht, Amelia H. Shen

1.10.1 Introduction

Complex digital systems have always combined application-specific integrated circuits (ASICs) and reprogrammable components, such as off-the-shelf microprocessors, where the application-specific component is often referred to as the *hardware component,* and the program running on the reprogrammable component is referred to as the *software component.* More recently, however, to achieve higher performance and density, as well as lower power dissipation and cost, designers have turned to more complicated packaging and are much more careful about distributing functionality between ASICs and reprogrammable components. Examples of such systems include embedded control systems for printing and plotting, telecommunications systems, and other signal processing systems.

While most digital functions can be implemented by software programs, a major reason for building dedicated ASICs is to satisfy performance constraints. Specialized integrated circuits are needed to assist the reprogrammable component for certain performance-critical tasks. In some cases, reprogrammable processors can meet performance constraints. However, a second important reason for using ASICs is that ASICs have significantly lower power dissipation than reprogrammable processors. Thus, system-wide power dissipation constraints may be met by using dedicated ASICs. However, when using standard printed circuit boards to

---

[28] M. Silveira, J. White, and S. Leeb, "A Modified Envelope-Following Approach to Clocked Analog Circuit Simulation," *Proceedings of the International Conference on Computer-Aided Design,* Santa Clara, California, November 1991, pp. 20-23.

[29] A. Lumsdaine, L.M. Silveira, and J. White, "Massively Parallel Simulation Algorithms for Grid-Based Analog Signal Processors," *IEEE Trans. Comput.-Aided Des.,* forthcoming.

provide the interconnect between ASICs and programmable components, some of the performance and power dissipation advantages of using ASICs may be eliminated. For this reason, the use of novel interconnect structures like multichip modules (MSMs) and three-dimensional packaging is increasing.

Our focus is on solving several critical problems in the design and verification of these large-scale digital systems and targeting maximal performance and minimal power dissipation. Currently, large-scale digital systems exploit sophisticated packaging methods. However, in the near future most of the hardware for such systems will be integrated on a single chip. In this context, we are addressing the problems described in the following sections.

1.10.2 Embedded System Design Aids

Currently, we lack the CAD tools that would allow for the quick and efficient design of a high level mixed hardware/software system or an embedded system. Most ASICs that correspond to application-specific processors in mixed hardware/software systems contain a ROM which occupies a large fraction of the area of the chip and contributes significantly to the power dissipation. The size of the ROM is determined by the amount of code that it must store, which, in turn, is generated by a compilation process from a higher level software specification. Separate ROM chips may also be associated with the reprogrammable processors in the system.

Compilation methods play a very important part in embedded system design. While compilation techniques are well-developed for purely software systems, a direct usage of these techniques to generate code in mixed hardware/software systems is not recommended. Typically, these techniques attempt to maximize system performance, i.e., the speed at which the threads of code are executed without regard to code density. In a mixed hardware/software system during compilation, one should maximize code density and minimize ROM size while meeting a system-level performance constraint. Maximizing code density will result in improved area and power dissipation.

At present, global, efficient methods to maximize code density during compilation have not been developed. It is possible to formulate this problem as one of data compression allowing the use of modified text compression algorithms to minimize code size. Graduate student Stan Liao is developing global and efficient optimization methods that will result in minimally-sized ROMs after compilation, improving the area and power dissipation characteristics of the system.

1.10.3 Formal Verification

It is critical to verify that the circuits obtained after application of sophisticated synthesis and optimization methods to the initial specification are error-free. Formal verification methods can guarantee correctness of an implementation with respect to a specification, even if the specification and implementation are described at vastly different levels of abstraction.

Most verification problems in VLSI design require the use of Boolean manipulation. Graduate student Amelia Shen has developed a representation of Boolean functions termed the Free Boolean Diagram (FBD) that has proven to be superior to other graph representations of Boolean functions. A FBD package capable of manipulating complex Boolean functions has been developed. This package serves as an equivalence checking engine for combinational circuits and is used in our behavioral verification methodology.

Our work on the behavioral verification of ASICs has been successful in checking the equivalence of descriptions at vastly different abstraction levels, such as a dependency graph of computations and a logic-level circuit. In the past year, graduate student Vishal Bhagwati has succeeded in developing a formal methodology for checking the equivalence of pipelined microprocessors against unpipelined specifications. The specification is the instruction set of the microprocessor with respect to which the correctness property is to be verified. A relationship, namely the β-relation, is established between the input/output behavior of the implementation and specification. The relationship corresponds to changes in the input/output behavior that result from pipelining, and takes into account data hazards and control transfer instructions that modify pipelined execution. The correctness requirement is that the β-relation hold between the implementation and specification. Symbolic simulation of the specification and implementation is used to verify their functional equivalence. The pipelined and unpipelined microprocessor can be characterized as definite machines (i.e., a machine in which for some constant k, the output of the machine depends only on the last k inputs) for verification purposes. It can be shown that only a small number of cycles, rather than exhaustive state transition graph traversal and state enumeration, must be simulated for each machine to verify whether the implementation is in β-relation with the specification.

The verification problems associated with mixed hardware/software systems are much more difficult to solve than those encountered in ASICs. While simulation is an integral part of verification methodology, it can be very expensive and time-consuming. Formal verification methods for such systems are thus very attractive. The next challenge is to develop a verification methodology for checking the equivalence of a software description of the functionality of the entire system against a mixed hardware/software implementation.

1.10.4 Logic Synthesis for Low Power

For many consumer electronic applications, low average power dissipation over a system-wide basis is desirable, and, for certain special applications, low power dissipation is of critical importance. For applications such as personal communication systems like hand-held mobile telephones, low-power dissipation may be the tightest constraint in the design. More generally, with the increasing scale of integration, we believe that power dissipation will assume greater importance, especially in multi-chip modules (MCMs) where heat dissipation is one of the biggest problems.

The average power dissipation of a circuit, like its area or speed, may be significantly improved by changing its architecture or technology. But once these architectural or technological improvements have been made, the switching of the logic and leakage currents will ultimately determine its power dissipation. Traditionally, logic synthesis has been applied to improve the area or speed of a circuit. We are developing methods for fast, accurate logic-level power estimation methods that will help development of architectural and logic-level synthesis strategies that target low power. We believe that the development of logic optimization algorithms that directly target low power dissipation poses the next major challenge for logic synthesis.

Switching activity is the primary cause of power dissipation in CMOS circuits. Accurate, average switching activity estimation for sequential circuits is considerably more difficult than for combinational circuits, because the probability of the circuit being in each of its possible states has to be calculated. The Chapman-Kolmogorov equations can be used to accurately estimate the power dissipation of sequential circuits by computing the exact state probabilities in steady state. However, the Chapman-Kolmogorov method requires the solution of a linear system of equations of size 2N, where N is the number of flip-flops in the machine.

Graduate student José Monteiro has developed a computationally efficient scheme to approximate average switching activity which requires the solution of a non-linear system of equations of size N, where the variables correspond to state line probabilities. The approximation method is within three percent of the exact method, but is up to 100 times faster for large circuits.

Graduate student Mazhar M. Alidina has developed a powerful sequential logic optimization method that is based on (1) selectively precomputing the output logic values of the circuit one clock cycle before the values are required and (2) using the precomputed values to reduce internal switching activity in the succeeding clock cycle.

The primary optimization step is the synthesis of the precomputation logic, which computes the output values for a subset of input conditions. If the output values can be precomputed, the original logic circuit can be "turned off" in the next clock cycle and will not have any switching activity. The size of the precomputation logic determines the power dissipation reduction, area increase, and delay increase relative to the original circuit.

Given a logic-level sequential circuit, we are developing automatic methods of synthesizing precomputation logic to achieve maximal reductions in power dissipation.

1.10.5 Publications

Journal Articles

Ashar, P., S. Devadas, and K. Keutzer. "Path-Delay-Fault Testability Properties of Multiplexor-Based Networks." *Integrat. VLSI J.* 15(1): 1-23 (1993).

Bryan, M.J., S. Devadas, and K. Keutzer. "Analysis and Design of Regular Structures for Robust Dynamic Fault Testability." *VLSI Des.: Int. J. Custom-Chip Des. Simulat. Test.* 1(1): 45-60 (1993).

Cheng, K-T., S. Devadas, and K. Keutzer. "Robust Delay-Fault Test Generation and Synthesis for Testability Under a Standard Scan Design Methodology." *IEEE Trans. Comput.-Aided Des.* 12(8): 1217-1232 (1993).

Devadas, S., K. Keutzer, and S. Malik. "Computation of Floating Mode Delay in Combinational Logic Circuits: Theory and Algorithms." *IEEE Trans. Comput.-Aided Des.* 12(12): 1913-1923 (1993).

Devadas, S., K. Keutzer, S. Malik, and A. Wang. "Computation of Floating Mode Delay in Combinational Logic Circuits: Practice and Implementation." *IEEE Trans. Comput.-Aided Des.* 12(12): 1924-1936 (1993).

Devadas, S., K. Keutzer, S. Malik, and A. Wang. "Event Suppression: Improving the Efficiency of Timing Simulation for Synchronous Digital Circuits." *IEEE Trans. Comput.-Aided Des.* Forthcoming.

Devadas, S., K. Keutzer, S. Malik, and A. Wang. "Verification of Asynchronous Interface Circuits with Bounded Wire Delays." *J. VLSI Sig. Process.* Forthcoming.

Jyu, H.-F., S. Malik, S. Devadas, and K. Keutzer. "Statistical Timing Analysis of Combinational Logic Circuits." *IEEE Trans. VLSI Syst.* 1(2): 126-137 (1993).

Liao, S., and S. Devadas. "Automatic Generation and Verification of Sufficient Correctness Properties for Synchronous Array Processors." *IEICE Trans. Info. Syst.* E76-D(9): 1030-1038 (1993).

Van Aelten, F., J. Allen, and S. Devadas. "Event-Based Verification of Synchronous Globally Controlled Logic Designs against Signal Flow Graphs.", *IEEE Trans. Comput.-Aided Des.* 1(1): 122-134 (1994).

Van Aelten, F., J. Allen, and S. Devadas. "Verification of Relations Between Synchronous Machines." *IEEE Trans. Comput.-Aided Des.* 12(12): 1947-1959 (1993).

Conference Papers

Bhagwati, V., and S. Devadas. "Automatic Verification of Pipelined Microprocessors." *Proceedings of the 31st Design Automation Conference,* San Diego, California, June 1994. Forthcoming.

Camposano, R., S. Devadas, K. Keutzer, S. Malik, and A. Wang. "Implicit Enumeration Techniques Applied to Asynchronous Circuit Verification." *Proceedings of the 26th Hawaii International Conference on System Sciences,* Koloa, Hawaii, January 1993, pp. 300-309.

Liao, S., S. Devadas, and A. Ghosh. "Boolean Factorization using Multiple-Valued Logic Minimization." *Proceedings of the International Conference on Computer-Aided Design,* Santa Clara, California, November 1993, pp. 606-611.

Monteiro, J., S. Devadas, and A. Ghosh. "Retiming Sequential Circuits for Low Power." *Proceedings of the International Conference on Computer-Aided Design,* Santa Clara, California, November 1993, pp. 398-402.

Monteiro, J., J. Kukula, S. Devadas, and H. Neto. "Bitwise Encoding of Finite State Machines." *Proceedings of the Seventh International Conference on VLSI Design,* Calcutta, India, January 1994.

Monteiro, J., S. Devadas, and B. Lin. "A Methodology for Efficient Estimation of Switching Activity in Sequential Circuits." *Proceedings of the 31st Design Automation Conference,* San Diego, California, June 1994. Forthcoming.

Shen, A., S. Devadas, and A. Ghosh. "Probabilistic Construction and Manipulation of Free Boolean Diagrams." Proceedings of the International Conference on Computer-Aided Design, Santa Clara, California, November 1993, pp. 544-549.