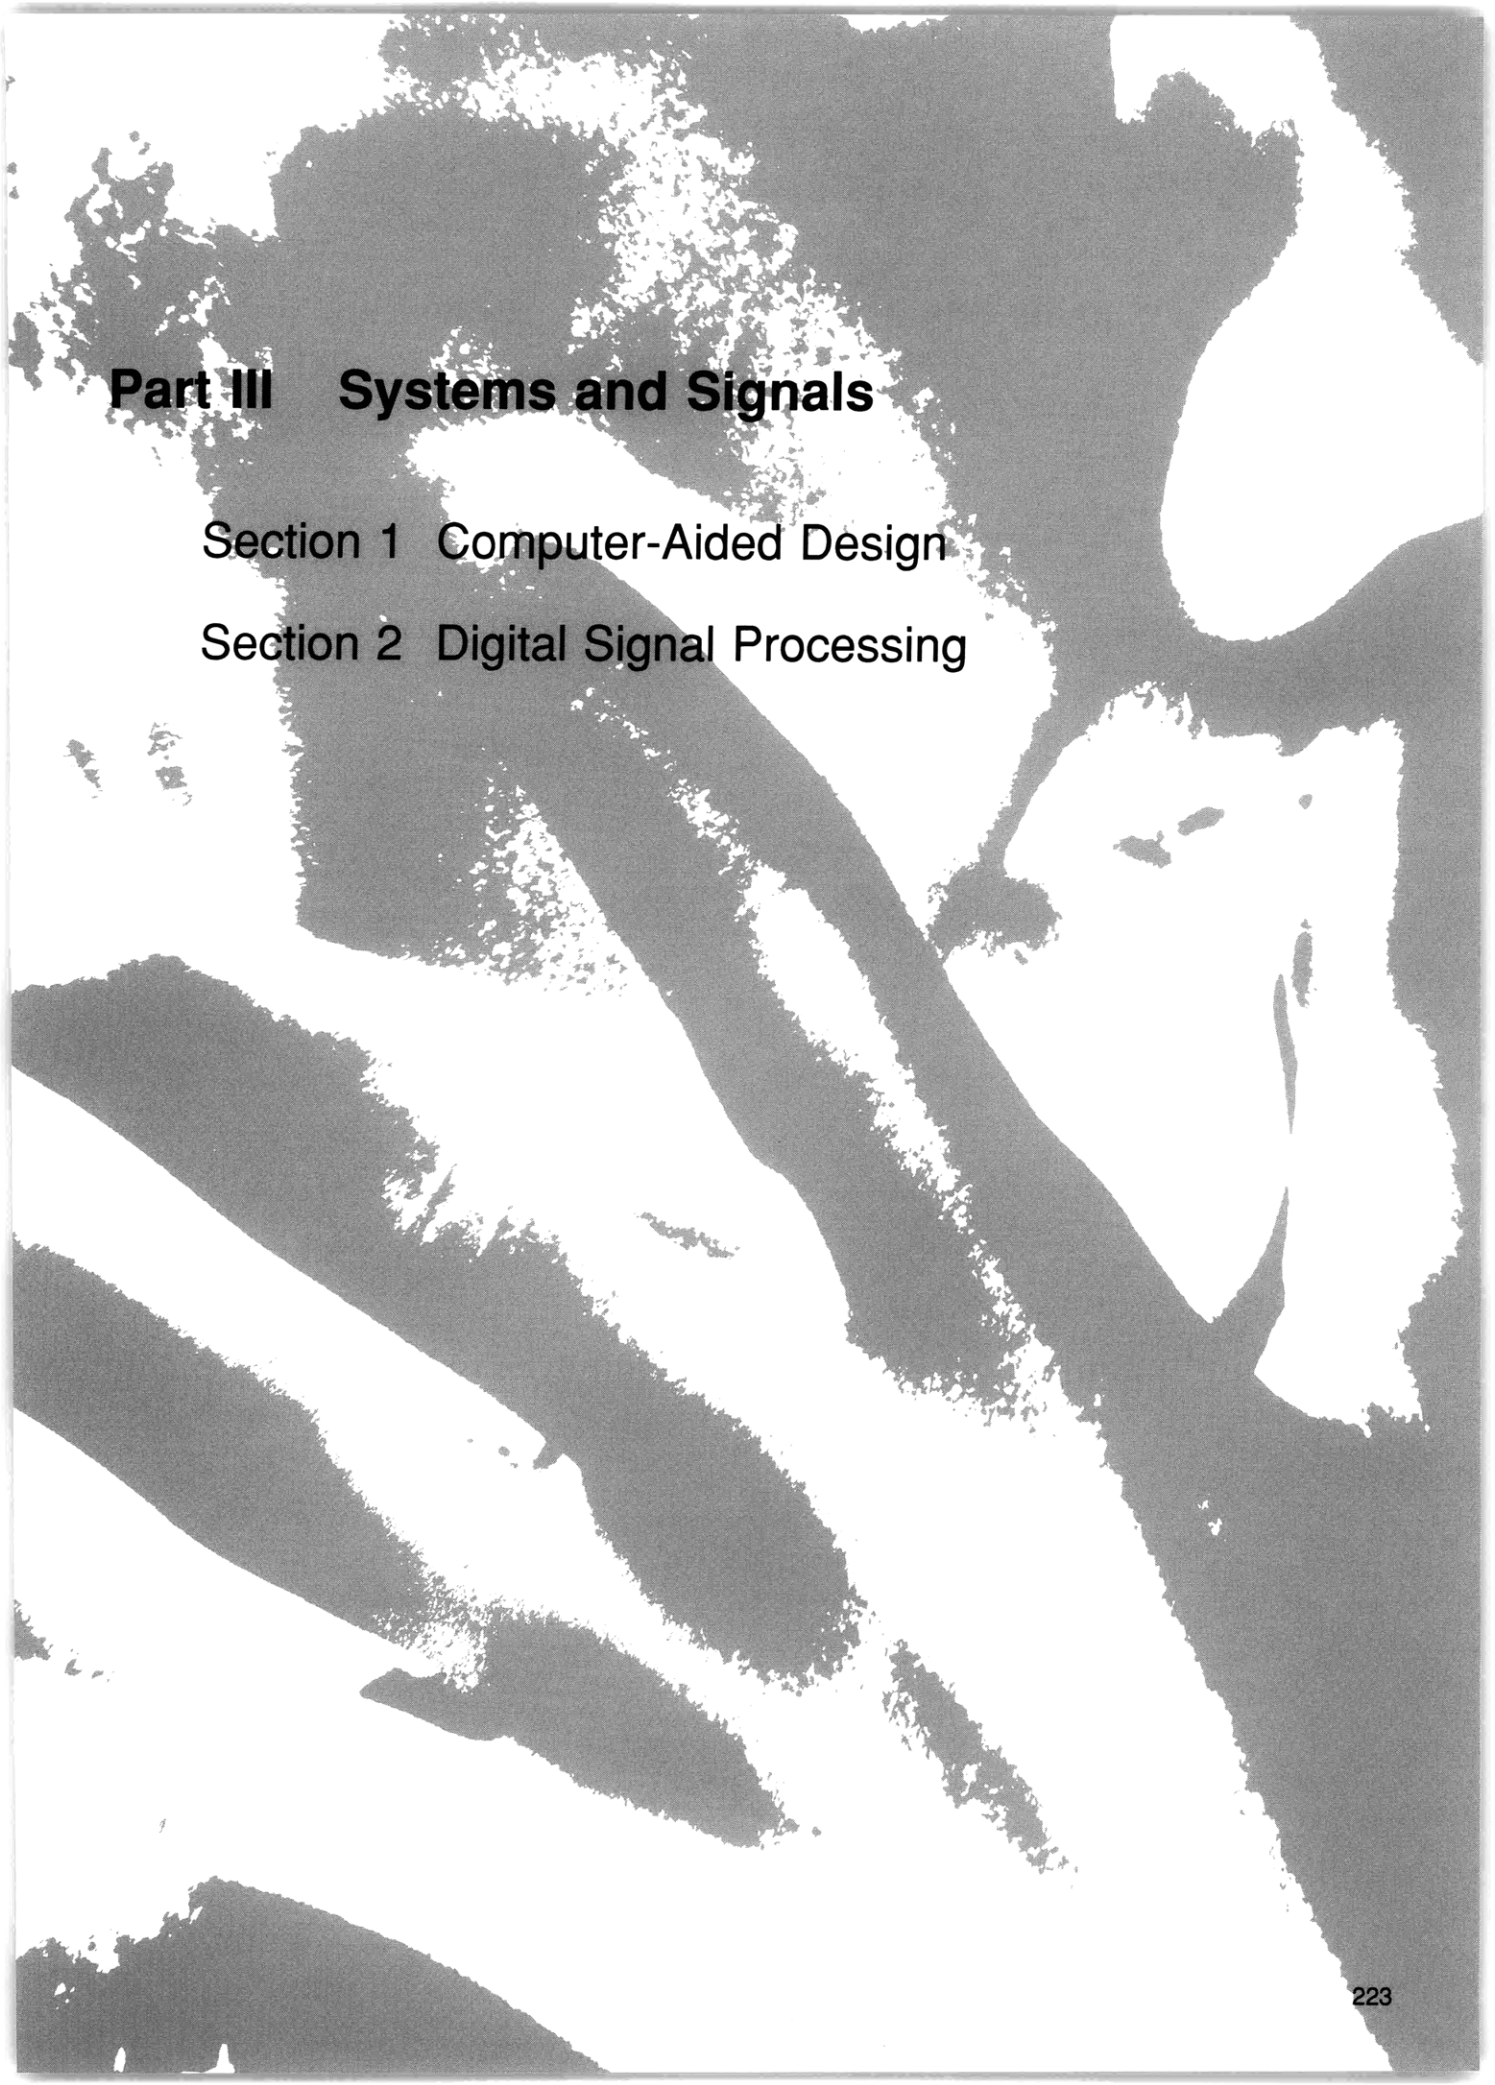

Part III Systems and Signals

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 William J. Dally, Professor Berthold K.P. Horn, Professor Hae-Seung Lee, Professor Terry P. Orlando, Professor Martin A. Schmidt, Professor Stephen D. Senturia, Professor Charles G. Sodini, Dr. Sami M. Ali, Dr. Robert C. Armstrong, Dr. John G. Harris

Visiting Scientists and Research Affiliates

Dr. Ichiro Masaki,[1] Jose C. Monteiro, Dr. Filip J. Van Aelten, Dr. Herre S.J. van der Zant

Graduate Students

Mazhar M. Alidina, Kelly S. Bai, Donald G. Baltus, Vishal L. Bhagwati, Xuejun Cai, Steven J. Decker, Lisa G. Dron, Ibrahim M. Elfadel, Kamyar Eshghi, J. Mikko Hakkarainen, Frederick P. Hermann, Mattan Kamon, Craig L. Keast, Songmin Kim, Steven B. Leeb, Stan Y. Liao, Andrew Lumsdaine, David R. Martin, Ignacio S. McQuirk, Keith S. Nabors, Gene T. Osgood, Joel R. Phillips, Lisa A. Pickelsimer, Khalid Rahmat, Mark W. Reichelt, Mark N. Seidel, Gideon P. Stein, Chin Hwee Tan, Amelia H. Shen, Luis M. Silveira, Ricardo Telichevsky, Christopher B. Umminger, Paul C. Yu

Technical and Support Staff

Susan E. Chafe, Dorothy A. Fleischer

1.1 Custom Integrated Circuits

Sponsors

Analog Devices, Inc.
IBM Corporation

Project Staff

Professor Jonathan Allen, Robert C. Armstrong, Donald G. Baltus, Lisa A. Pickelsimer, Mark W. Reichelt, Chin Hwee Tan, Filip J. Van Aelten

The overall goal of VLSI 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 layout (via design rule checking), and circuit and logic representations (both via simulation). These techniques for checking correctness are usually local to the particular representational level involved and, while these techniques 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 an increasing need to provide CAD tools that serve as a framework for design exploration, 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 has recently been provided in a major survey paper[2] 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 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.

In addition to the basic issues of circuit optimization, architectures for digital signal processing have been studied because of the highly parallel nature

---

[1] Computer Science Department, General Motors Research Laboratory, Warren, Michigan.

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

of the algorithms involved and the need for a very high level of real-time performance in these systems. Emphasis on developing formally specified systems has increased so that the design space can be comprehensively searched, and verification of the resultant design can be confirmed in terms of the initial behavioral specification. Current projects focus on design methodologies for very high-speed clocked circuits, the trade-offs between power dissipation and circuit delay with transistor sizing as a parameter, design of highly optimized array architectures for digital signal processing, and provision of a VLSI design database that provides for incremental consistency maintenance.

There has been a continuing evolution of circuit styles to provide high performance in a minimal space. For example, in CMOS technology, static designs are utilized, but precharge-evaluate methodologies have also been employed to provide speed, minimize the number of transistors, and increase utilization of n-channel devices (rather than slower p-channel devices). In recent years, examples of a new circuit style called *True Single-Phase Clocking* (TSPC) have been published.[3] This style provides for clock speeds of several hundred megahertz in modest CMOS technologies. While these results have been confirmed, an analysis that explains the speed has not been provided, and no overall design methodology for these circuits has been demonstrated. In her master's thesis research project, Pickelsimer[4] has analyzed the performance of TSPC circuits as well as other precharge-evaluate forms, in the context of a ripple-carry binary adder. A library of basic TSPC circuit elements has been provided, as well as a comprehensive design methodology for the optimal structuring of TSPC and TSPC-related circuits for minimal delay. In this way, it is possible to systematically derive an optimal circuit form from an input logical specification through the appropriate combination of matched circuit forms. All of these circuits require only a single-phase clock, without the need for clock waveform inversion, and hence avoid the problems of clock skew and waveform overlap.

It has frequently been observed that sizing of transistors (i.e., variation of MOSFET channel widths) can significantly improve the speed of many circuit forms. In particular, optimal procedures have been provided to examine the trade-off of circuit delay with aggregate transistor channel width. Even more recently, techniques have been devised to probabilistically characterize switching activity on circuit nodes in terms of the logic function being performed and hence deriving power estimates for the circuit, which can be modified through logical transformations.[5] Tan, in her master's thesis work, is studying the trade-off between circuit delay and power dissipation (as characterized by probabilistic power estimates) with transistor channel width size as a parameter. The results of this study will be useful in maximizing the performance of low-power circuitry or minimizing the power of circuits designed to meet a given delay specification.

In a doctoral dissertation project by Baltus, comprehensive optimization of array logic designs is being achieved over all design representations from the input function specification to the final layout. The goal is to generate a custom layout specification from an input functional description using a language called FLASH, which is well suited to functionally characterizing signal processing algorithms, and can be viewed as an extension of the SILAGE language.[6] FLASH designs are converted to data dependence graphs, where some elements of these graphs are reused and indexed to provide a compact representation that does not grow with the size of the array. Since algorithms are mapped onto a class of architectures with well-defined structural and interconnection characteristics, important implementation-level costs can be accurately modeled at the architectural level. More specifically, the structured nature of the target architecture allows the temporal and spatial relationships between computations to be seen at the architectural level.

The availability of an abstract model which accurately predicts implementation-level characteristics allows the process of performance-directed synthesis to be formalized. Powerful techniques from

---

3 J. Yuan and E. Svensson, "High-Speed Circuit Technique," *IEEE J. Solid-State Circ.* 24(1): 62-70 (1989).

4 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.

5 A. Shen, A. Ghosh, S. Devadas, and K. Keutzer, "On Average Power Dissipation and Random Pattern Testability of CMOS Combinational Logic Networks," *Digest of Technical Papers,* IEEE/ACM International Conference on Computer-Aided Design, Santa Clara, California, November 1992, pp. 402-407.

6 P.N. Hilfinger, "A High-Level Language and Silicon Compiler for Digital Signal Processing," *Proceedings of the IEEE 1985 Custom Integrated Circuits Conference,* Portland, Oregon, May 1985, pp. 213-216.

graph theory, optimization theory, and linear algebra can be used to allow the space of possible implementation architectures to be efficiently, comprehensively, and systematically explored. A systolic array synthesis method has been developed which exploits these techniques. The space of possible implementation architectures is efficiently explored, and a set of provably optimal register-transfer-level implementations is generated. The exhaustive nature of the design exploration process ensures that all optimal designs are generated. Compared to other work in this area, the developed technique is applicable to a wider class of algorithm and is the only system of its kind which incorporates exhaustive architectural exploration.

Armstrong has built a multirepresentation VLSI database that provides for incremental consistency maintenance.[7] The goal of this system is to automatically, incrementally, and continuously provide for consistent alignment of all levels of design abstraction. The system provides a design environment which maintains these consistency linkages while allowing the designer to initiate changes at any level of representation. Alignment is achieved through the introduction of explicit correspondence zones between the adjacent representation levels. Some of the necessary translation rules to maintain correspondence are built into the system, but there is also provision for the user to introduce new relationships between levels of design representation. The initial implementation of this system has required over 60,000 lines of code, written in C, but database schemas that are written in C++ will now be introduced to vastly compress the size of the code, increase the speed of the database, and minimize errors and the need for debugging. In addition, the ability to explore designs at a given level of representation is being coupled to the basic database structure. A circuit simulator is being connected at the circuit-level of representation, and a transistor-sizing algorithm is also being used at that level. Since there will always be design requirements that have not been contemplated in the existing design, a default capability is being provided that will indicate structural regions that must be changed in a design in response to a user-initiated change at another level without automatically providing the translation rule that maintains correspondence. In this way, users will not be restricted to a system-provided set of correspondences, but may continuously introduce new translations that may eventually be assimilated into the overall system.

There is increasing interest in overall CAD systems where the database is central and surrounded by specific applications for design and verification at a variety of different levels of representation, such as in the database just described. This research group is currently using a network of workstations that provides an implicit client-server architecture together with coupling to a mainframe and its attached large tape and disk capabilities interconnected with high-speed fiberoptic switching. Scalable, parallel architectures, particularly useful for highly parallel simulation codes, are now being studied for introduction into this framework. They will also support multiple client utilization of the database with extremely fast communication between the contributing processors. This system will serve as a test bed for the combined utilization of high-performance workstations and mainframe-based storage for high-speed, reliable database access. This environment also provides a structure in which distributed parallel algorithms can run coherently within an overall environment that supports other specialized, nonparallel CAD algorithms.

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

Sponsor

National Science Foundation
Grant MIP 91-17724

Project Staff

Professor John L. Wyatt, Jr., Professor Berthold K.P. Horn, Professor Hae-Seung Lee, Dr. Ichiro Masaki, Professor Charles G. Sodini, Professor Jacob K. White

1.2.1 Project Summary

In real-time machine vision the sheer volume of image data to be acquired, managed and processed leads to communications bottlenecks between imagers, memory, and processors, and also to very high computational demands. We are designing and testing experimental *analog* VLSI

[7] 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 1992, pp. 336-343.

systems to overcome these problems. The 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 animal or machine.

The next section details substantial progress on eleven different designs, including systems for camera motion estimation, edge detection, 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 J. Mikko Hakkarainen 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 Boston for a year to work with us and investigate the possibilities of this type of design for future use in automobiles. He has made a real contribution to this project in directing our attention to practical automotive applications.

1.2.2 Vision Chip Designs

Real-Time Stereo Vision System in CCD/CMOS Technology

J. Mikko Hakkarainen completed, for his doctoral thesis under Professor Hae-Seung Lee, the design of an analog VLSI design for a high-speed binocular stereo vision system used for the recovery of scene depth. In this context he has attempted to exploit the principal advantages of analog VLSI—small area, high speed, and low power—while minimizing the effects of its traditional disadvantages—limited accuracy, inflexibility, and lack of storage capacity. A CCD/CMOS stereo system implementation was proposed, capable of processing several thousand image frame pairs per second for 40×40 pixel binocular images, which should yield approximately frame rate operation on 256×256 pixel image pairs. A shift-and-compare module of the stereo system has been fabricated and characterized.

The algorithmic focus of this research was on binocular stereopsis and the stereo correspondence problem in particular. Three stereo algorithms (Marr-Poggio-Drumheller, Pollard-Mayhew-Frisby, and Prazdny) were considered in detail as solutions to this problem and as candidates for VLSI implementation. The most crucial step in the recovery of stereo correspondence was found to be the correlation of left and right images through candidate matches and local support scores. The Marr-Poggio-Drumheller (MPD) algorithm was shown to be best suited for an efficient VLSI implementation. A detailed simulation study of this algorithm was presented. The simulations were specifically designed to address issues related to analog VLSI implementation.

A full MPD stereo system consists of seven modules: two imagers (left and right views), two image filters, a shift-and-compare module ("correlator"), a local support filter, and a vote taker. The shift-and-compare module together with the local support filter perform the largest number of operations per pixel in the stereo system, and thus they present the speed bottleneck. Of the two modules, the shift-and-compare function is more demanding for hardware implementation, particularly because the local support filter has a simple low-pass characteristic that has already been demonstrated previously using the CCD/CMOS approach. For these reasons the shift-and-compare circuit was selected for design and fabrication. A 40×40 pixel absolute-value-of-difference (AVD) array, the core processor of the shift-and-compare function, was fabricated through MOSIS in a 2 μm CCD/CMOS process. The array can perform the AVD operation between two arrays a programmable number of times. At a 10 MHz clock rate for the CCD processor, approximately 1000 array (image frame) pairs can be processed per second for a total disparity range of 11 pixels. Special purpose test circuits were designed that are compatible with such high speed processing requirements. Individual AVD processors as well as the full array were characterized. Array functionality was demonstrated in the stereo algorithm context by running the MPD stereo system with actual data from the AVD chip. Upon comparison, the chip output was seen to correspond closely with fully computerized stereo simulations for both artificial and real input image pairs.

A Dynamic-Wires Chip for Pixel Clustering

Dr. John G. Harris, working in collaboration with Shih-Chii Liu from the Rockwell Science Center, has continued to develop the *dynamic wires* paradigm for solving a new class of computer vision problems with analog hardware. This methodology provides dedicated lines of communication among groups of pixels of an image which share common properties. In simple applications, object regions can be grouped together to compute the area or the center of mass of each object. Alternatively, object boundaries may be used to compute curvature or contour length. These measurements are useful for higher level tasks such as object recognition or structural saliency. The dynamic wire meth-

odology is efficiently implemented in fast, low-power analog hardware. Switches create a true electrical connection among selected pixels—dynamically configuring wires or resistive networks on the fly. Dynamic wires provide a different processing approach from that of present vision chips, which are limited to pixel-based or image-based operations. Using this methodology, they have successfully designed and demonstrated a custom analog VLSI chip which computes contour length.

Single-Chip CMOS Imager and Smoothing and Segmentation System Using Resistive Fuses

Paul Yu, working with Professor Hae-Seung Lee, has designed, built, and tested an integrated CMOS resistive-fuse processor capable of smoothing out noise while preserving the edges of a 32x32 image. The on-chip photo-transistor imager converts the optical image into electrical currents. Through a system-level optimization, a single-pixel processing element using only 15 transistors for the horizontal resistive fuses and only one transistor for the vertical resistor was realized. The processed output is read out using a row decoder and a column MUX. Since the on-resistance and off-voltage of the resistive fuses can be electronically controlled, one can vary the degree to which the image is smoothed and segmented electronically. The chip was fabricated through MOSIS and tested; it works well. The chip contains over 19,000 transistors and occupies a 9.2 mm x 7.9 mm area. Operating from a single 5-V supply, it typically dissipates only 10 mW.

Novel Operational Amplifier for Video Rate Analog-to-Digital Conversion

Paul Yu is working with Professor Hae-Seung Lee on novel operational amplifiers to serve as key components in a video-rate pipelined analog-to-digital converter. A new gain-enhancement technique, using a replica amplifier, is being studied. Although this technique is general and can be applied to any conventional op-amp topologies such as folded-cascode, class-AB, differential, or single-ended, a two-stage topology without any cascode was chosen to maximize the available output swing and input common-mode range. Experimental results show that the gain enhancement technique improves the open-loop gain by a factor of 13 without adversely affecting the output swing or the settling behavior. The amplifier achieves an effective open-loop gain of 10,800 using only minimum-length transistors in the signal path and no cascode. Even with a 1 K ohm load resistor, an effective open-loop gain of over 10,000 is maintained. The output swings within 100 mV from either rail. The gain-bandwidth product is 63 MHz with a 27 pF load capacitance and 9 mW power dissipation. These results are obtained using only a 2-V supply. The chip was fabricated in a 1.2um CMOS technology.

CMOS/CCD Imager and Focus-of-Expansion Chip for Camera Motion

Ignacio McQuirk, working with Professors Berthold Horn and Hae-Seung Lee, has continued development of a chip that determines the direction of camera translation directly from a time-varying image sequence with a real-time analog VLSI system. His approach assumes a camera moving through a fixed world with only translational velocity. Also, the surfaces in the environment must have some spatial variations in reflectance. 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. 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. Knowledge of the location of the FOE in the case of pure translation also allows coarse calculation of the depth map of the imaged world up to a scale factor ambiguity, as well as the associated time to impact.

Previously, various possible algorithms for estimating the FOE were compared in terms of accuracy, robustness, and feasibility for single-chip analog VLSI implementation with on-chip imaging circuitry. As a result of these studies, 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. Additionally, several system architectures and circuit designs were explored, and one was chosen for realization in analog VLSI. Mr. McQuirk designed a preliminary CMOS/CCD version of the FOE chip using on-chip CCD imaging circuitry and a row-parallel processing scheme. Several components of this design were fabricated through MOSIS and tested with excellent results.

During the period covered by this report, the design of the FOE chip was revised, and a preliminary version of the complete chip was also fabricated through MOSIS. He is currently in the process of designing and fabricating a test printed circuit board for the FOE system consisting of pattern generators, CCD clock drivers, FOE chip bias circuitry, I/O drivers, and an interface to the Vision Chip Test System. This will allow for stand-alone operation of the FOE system as well as operation with the

Vision Chip Test System for algorithmic, system, and circuit level testing.

Robust Analog Image Segmentation Chip for Edge Detection and Camera Motion

Lisa Dron, working with Professor Horn, has invented a robust method of image segmentation that has been shown to work well in computer simulations. The basic idea is to segment an image between two nodes of a grid wherever the difference in intensities at the two nodes exceeds a threshold at *all of several different scales* of spatial low-pass filtering. (The threshold can vary with the scale.) The method is *robust* in the sense that local noise spikes in the image are averaged out with filtering and do not contribute false edges, and it is *accurate* in the sense that the unfiltered image plays a role, since threshold must be exceeded there as well. Bandwidth compression is achieved by retaining intensity values only at nodes adjacent to edges. Intensities elsewhere are obtained by spatial averaging using a resistive grid.

A preliminary design has been studied for the segmentation circuit using a two-dimensional CCD array with added circuitry for computing differences and performing threshold tests. The structure of the array uses the same principles developed by Dr. Craig Keast for smoothing and segmentation; however, it has been modified to implement the multi-scale veto (MSV) algorithm. In addition, several improvements suggested by Dr. Keast to improve the performance of the CCD circuits have been incorporated. A test chip containing a one-dimensional array, as well as its isolated subcircuits, has been fabricated through MOSIS and tested. A second chip with a two-dimensional array has now been designed and fabricated and is currently being tested.

The two primary applications of the multi-scale veto (MSV) chip are image reconstruction and edge detection. Ms. Dron has chosen to pursue the second application in order to design a system for computing general camera motion in real-time with analog processing. The binary edge signals produced by the MSV chip can be used to obtain correspondences between successive frames by matching patches of the edge maps in the two images. It is well known that it is possible to compute the relative camera motion between frames from as few as five correspondences, provided the internal camera parameters are known; although in practice, several more than five are used in order to obtain a more reliable estimate of the motion.

Conventional digital methods for estimating camera motion typically utilize special purpose DSP chips to find point correspondences by matching gray levels over small patches in the two images using either cross-correlation or mean absolute value of difference. Due to the complexity of the computations with 8-bit gray-level data, the maximum displacement which can be computed at video-rate with a single chip is currently limited on commercially available chips to around ±8 pixels in any direction. In addition, the displacement estimates obtained from matching gray levels are in general less reliable than those which can be obtained by matching edges. It is possible to define functions on two patches of binary edges which have a large non-zero value only when there is a good match between their edge patterns. Robust displacement estimates are obtained by identifying patches which produce a well-peaked unimodal response when tested against all like-sized patches within a large search area in the other image. An analog processor can be designed to effectively implement such a matching function because it is not necessary to accurately compute the exact value of the function, but only to find the location of its peak value and to determine if it is significantly greater than zero.

Ms. Dron has recently completed a study of different matching functions which can be computed by analog methods using the binary edge signals produced by the MSV chip. She has designed several test structures which will eventually form the basis of the analog matching circuit and has sent these out for fabrication by MOSIS. In addition to designing the circuits, she has also developed the algorithms needed to use the point correspondences in a real-time system. She has derived a robust least-squares method suitable for implementation on a digital processor to compute camera motion, as well as an algorithm to determine the internal camera calibration parameters, required for motion computation, from point matches in a sequence of images for which the translational motion is known.

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.

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 x 16 mesh network connects the processing elements, and larger meshes may be constructed using several chips.

The chip has been submitted for fabrication in MIT's 1.5 μm CCD/CMOS process. Fabrication is complete, and testing will begin in January 1993.

Single-Chip Alignment Sensor

Christopher Umminger, working with Professors Sodini and Horn, has been researching implementations of single-chip sensors to measure alignment error for tasks in automated manufacturing. The sensor 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 Moirè 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 finished in June and submitted for fabrication. The completed chips were received in mid-September and preliminary tests show that the circuits are working as expected. Characterization of the sensor in an alignment system awaits the construction of a prototype aligner. Design of the Moirè fringe sensor will begin in November and will be submitted for fabrication at the end of the year.

Switched-Capacitor System for Merging Depth and Slope Estimates

Mark Seidel, working with Professor John Wyatt and Principal Research Scientist Thomas Knight, has been designing and fabricating 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 multi-sensor 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. Testing of a 1-D version fabricated through the MOSIS system revealed a shorting problem that inhibits further testing. Laser "surgery" done at MIT Lincoln Laboratory failed to resolve the origin of the problem. A MOSIS Tinychip containing some subcircuits has also been fabricated, but not tested. The architecture of the chip was presented at ISCAS 1992 in San Diego. In the final two months of the year, the MOSIS Tinychip will be tested.

Brightness Adaptive CCD Imager with Analog and Digital Output

Steven Decker, working with Professor 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/sec, column parallel ouput, a brightness adaptation capability, and user selectable analog or 10-bit digital output.

A CCD-based algorithmic A/D converter has been developed which requires only CCD gates and one operational amplifier. It avoids the need for subtraction of charge packets required by a similar previous design, a step which severely limited the attainable accuracy. Analog output can also be obtained by shifting the 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.

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 signal 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 all be done digitally. A layout scheme which allows column parallel output while maintaining a fill factor exceeding 50 percent has been achieved.

1.2.3 Mixed Circuit/Device Simulation for CMOS/CCD Systems

Andrew Lumsdaine and Mark Reichelt, under the supervision of Professor Jacob K. White, are 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. They found experimental evidence that accelerating WR convergence using standard overrelaxation can produce oscillatory results and have recently developed methods for eliminating this phenomenon. In particular, they have developed two new techniques for accelerating WR convergence, one based on a generalized conjugate-direction method and the other based on a frequency-dependent successive over-relaxation (SOR) algorithm. Proofs of the optimality of the SOR algorithm and the guaranteed convergence of the C-D method have been completed, and experimental results indicate that both methods can reduce the number of waveform iterations as much as an order of magnitude.

1.2.4 Vision Chip Test System

In order to test VLSI early vision processors that are being developed by this project, it is important to input synthetic data or real image data to the chip and study processed image output from the chip. Kamyar Eshghi and Professor Sodini have constructed a Demonstration System for Early Vision processors (DSEV) that will perform this high-level testing. DSEV can capture images from video cameras and transmit them to a device under test (DUT). Processed images retrieved from the DUT are stored by DSEV and may be displayed for visual comparison to the input image. Also, the same images that were input to the DUT may be processed by software algorithms on a workstation and compared with the DUT output, allowing the sensitivities of the integrated circuit design to be studied. DSEV has been standardized in such a way that a minimum of effort is required to interface various vision chips to it.

The hardware of DSEV consists of a camera, display monitor, frame grabber, display interface, image memory, SUN workstation, S-bus to VME adapter, MaxBus interface, DUT interface, and DUT board. A MaxVideo20 module was purchased from DataCube Incorporated, which includes a frame grabber, display interface and image memory. Software has been written to control MaxVideo20, and the MaxBus interface board and DUT interface boards were designed, built and tested.

Each student will have to design and build his or her own DUT board specifically for the chip to be tested, but the rest of the system is shared, at a substantial savings in overall group effort. Ignacio McQuirk is now modifying and altering the system for use in testing his motion chip. He is currently moving the level-shifting and scaling circuitry from the DUT board to the A/D and D/A converter boards and building more of those boards.

1.2.5 Novel Camera Calibration Method Using Images of Lines and Spheres

A stereo vision system and two systems for camera motion determination were described earlier in this report. Both these tasks require accurate knowledge of camera parameters, specifically, principal point, focal length and distortions such as horizontal scaling and radial distortion, for any analog or digital implementation.

Although camera calibration is important for 3D measurements, it is often neglected because it is a laborious process and requires special apparatus and skill. The aim of this project is to devise a method for calibration which will be simple to perform and accurate to better than one percent error in camera parameters.

Most of the techniques currently used in the field of robotics use images of a 3D object whose world coordinates are accurately known and then proceed to find both the camera parameters and the position and orientation of the camera relative to the 3D object. These methods have a few inherent weaknesses. It is often hard to decouple the camera parameters from the camera orientation parameters (external parameters). To have any chance of success, a very precise 3D object is required with easily identifiable calibration points. These objects are relatively hard to obtain, and the number of calibration points in the object is small (on the order of 100).

Gideon Stein has been working with Professors Wyatt and Horn to develop new techniques for calibration of internal camera parameters which will not require finding all the external calibration parameters.

The first method finds the radial distortion parameters. The image of straight lines under perspective projection are straight lines. In real images these lines are curved due to radial distortion. By correcting for the radial distortion the lines can be made straight. This is an adaptation of the "plumb line" method known in the literature.

The second method finds the focal length and the principal point. It assumes the image has been corrected for radial distortion. The image of a sphere under perspective projection is an ellipse whose major axis is on a line which passes through the principal point. The eccentricity of the ellipse is a function of the focal length. By taking the image of a number of spheres the intersection of the lines passing through the major axis gives the principle point. The focal length can be found from the eccentricity. This method finds the principal point to within one pixel but so far does not provide accurate estimates of the focal length.

The third method involves rotating the camera around the y-axis through a known angle. By tracking points in the scene the focal length and radial distortion can be recovered. Currently the estimates are repeatable to within 0.5 percent. To a lesser degree of accuracy, the principal point can also be recovered. This method is suitable for self calibration of autonomous robots working in unstructured outdoor environments.

Future work involves refining these three methods and evaluating their performance compared to the more traditional methods.

1.2.6 Analog versus Digital Approaches to Machine Vision Hardware

Recent technical progress in analog vision chips has stimulated comparisons of 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, David A. Martin, under the supervision of Professor Hae-Seung Lee, and Ichiro Masaki, visiting scientist from General Motors, chose direct image correlation and three-dimensional stereo vision algorithms as the cases. These algorithms are frequently used in various applications including intelligent vehicles and flexible manufacturing systems.

As the first step, two algorithms were compared at a conceptual level: a binary algorithm for digital implementation and a multi-valued algorithm for analog chips. The computer simulation is being developed for this algorithm level comparison. The chip implementation level comparison is also underway. Currently a Multi-Instruction Multi-Data Array Processor architecture, implemented in both analog and digital, is being studied. Each processing element includes an arithmetic unit, a data flow control unit, and a system control unit which includes program memory.

Our next goals include finalization of the chip architecture, design of the chips, and performance comparisons based on simulations of the designed chips. The criteria for the comparisons will include cost in silicon area, processing speed, manufacturing and development costs, testability, and other related issues.

1.2.7 Solutions of Theoretical Problems in Analog Hardware and Machine Vision

Gibbs Random Field Models

Gibbs random fields (GRFs) are probabilistic models inspired by statistical mechanics and used to model images in computer vision and image processing. In this research, conducted by Ibrahim Elfadel under the supervision of Professors John Wyatt, Berthold Horn and Alan Yuille, the analytical methods of statistical mechanics are brought to bear on these models. Specifically, the following fundamental problems have been addressed and contributions have been made to their solutions:

1. Mean-field estimation of a constrained GRF model: The configuration space of a GRF model is often constrained to produce "interesting" patterns. Mr. Elfadel has developed mean-field equations for estimating the means of these constrained GRFs. The novel feature of these equations is that the finiteness of graylevels is incorporated in a "hard" way in the equations.

2. Correlation-field estimation of a GRF model: GRF correlation functions are generally hard to compute analytically and expensive to compute numerically. Mr. Elfadel has used the mean-field equations developed above to propose a new procedure for *estimating* these correlation functions. His procedure, which is valid for both unconstrained and constrained models, is applied to the quadratic interaction model, and a new closed-form approximation for its corre-

lation function in terms of the model parameters is derived.

3. Network representation: Mr. Elfadel has shown how the mean-field equations of the GRF model can be mapped onto the fixed point equations of an analog winner-take-all (WTA) network. The building block of this network is the *generalized sigmoid mapping*, a natural generalization of the sigmoidal function used in artificial neural networks. This sigmoidal mapping has a very simple VLSI circuit implementation with desirable circuit-theoretic properties such as reciprocity and local passivity.

4. Solution algorithms: Iterated map methods and ordinary differential equations (ODEs) are proposed to solve the network fixed-point equations. In the former, Mr. Elfadel has shown, using Lyapunov stability theory for discrete systems, that the worst that could happen during synchronous iteration is an oscillation of period 2. In the latter, he has shown that the ODEs are the gradient descent equations of energy functions that can be derived from the mean-field approximation. One of his gradient descent algorithms can be interpreted as the generalization to analog WTA networks of Hopfield's well-known algorithm for his analog network.

5. Temperature dependence: The GRF temperature parameter reflects the thermodynamic roots of the model. Using eigenstructure analysis, he has studied the temperature effect on the stability of the WTA network fixed points. In particular, he derived new closed-form formulas for the *critical temperature* of a large class of models used in grayscale texture synthesis. The stability study is used to gain insight into the phase transition behavior of GRFs.

The technical details of these results will appear in Ibrahim Elfadel's Ph.D. thesis titled *From Random Fields to Networks,* which was submitted in February 1993.

In future work, Mr. Elfadel will investigate the implications of these results in the fields of image modeling, optimization, and analog hardware implementation of image processing and computer vision algorithms.

Switched-Capacitor Network Theory

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. For nonpathological networks, conditions for stability during each clock phase have been proven. Stability for the discrete time system has been proven for a large subclass of these networks. We also expect to prove stability for general networks.

Under constant input conditions, the switched-capacitor networks can be shown to be equivalent to networks comprised of, in general, nonreciprocal 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 is equivalent in the sense that all node voltages are the same, and all currents are equal to the corresponding average transferred charge over the 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. We will be generating several examples and writing a journal article.

Finally, additional work has been completed on bounding the discrete time system settling time of large switched capacitor networks. The bounds have been completely inverted and can be shown to be equivalent to the continuous time bounds in the appropriate limit. We will be generating several examples.

Generating a Uniformly Spaced Set of 3-D Rotations for Object Recognition Purposes

In some pattern recognition tasks, e.g., vehicle identification from aerial photography, one wishes to decide which object from a finite, previously determined set is represented in some portion of a single image. One seeks the best match by rotating each of the known objects to various positions, creating 2-D images in various sizes, and finding the best match. One problem is that there is no known way of writing down closed form expressions for a uniformly spaced set of N rotations in three dimensions, except for very special values of N corresponding to the number of faces on a regular

polyhedron. (The solution in two dimensions is trivial—the rotations are through angles of $(K/N) \cdot 360^\circ$, where K = 1,2,..., N – 1.)

Gene Osgood, working with Professors Wyatt and Horn, has explored an iterative method for producing uniformly spaced rotations using the quaternion number system. In this system each rotation in three dimensions is represented abstractly as a point on the three-dimensional sphere in four dimensions. Gene represents each point as a positive point charge and simulates the movement of these charges under mutual repulsion as they spread apart on the surface of the sphere. The method yields a substantial savings in computer time over conventional approaches that use unevenly spaced points.

1.2.8 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.* Forthcoming.

Elfadel, I.M., and A.L. Yuille. "Mean-Field Phase Transitions and Correlation Function for Gibbs Random Fields." *J. Math. Imag. Vision.* Forthcoming.

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

Herrmann, F., and C. Sodini. "A Dynamic Associative Processor for Machine Vision Applications." *IEEE Micromag.* 12(3): 31-41 (1992).

Liu, S.C., and J.G. Harris. "Dynamic Wires: An Analog VLSI Model for Object Processing." *Int. J. Comput. Vision* 8(3): 231-239 (1992).

Picard, R.W., and I.M. Elfadel. "Structure of Aura and Co-occurrence Matrices for the Gibbs Texture Model." *J. Math. Imag. Vision* 2: 5-25 (1992).

Umminger, C.B., and C.G. Sodini. "Switched Capacitor Networks for Focal Plane Image Processing Systems." *IEEE Trans. Circuits Syst. Video Technol.* 2(4): 392-400 (1992).

Wyatt, J.L., Jr., C. Keast, M. Seidel, D. Standley, B. Horn, T. Knight, C. Sodini, H-S. Lee, and T. Poggio. "Analog VLSI Systems for Image Acquisition and Fast Early Vision Processing." *Int. J. Comput. Vision* 8(3): 217-230 (1992).

Yu, P.C., S.J. Decker, H-S. Lee, C.G. Sodini, and J.L. Wyatt, Jr. "CMOS Resistive Fuses for Image Smoothing and Segmentation." *IEEE J. Solid-State Circuits* 27(4): 545-553 (1992).

Meeting Papers

Dron, L. "System-Level Design of Specialized VLSI Hardware for Computation Relative Orientation." *Proceedings of the IEEE Workshop on Applications of Computer Vision*, Palm Spring, California, November 30 - December 2, 1992.

Dron, L. "Dynamic Camera Self-Calibration from Controlled Motion Sequences." Submitted to the *1993 Conference on Computer Vision and Pattern Recognition.*

Elfadel, I.M., and A.L. Yuille. "Mean-Field Phase Transitions for Gibbs Random Fields." *SPIE '92 Proceedings*, San Diego, California, July 1992.

Hakkarainen, J.M., and H.-S. Lee. "A 40×40 CCD/CMOS AVD Processor for Use in a Stereo Vision System." *1992 European Solid-State Circuits Conference*, Copenhagen, Denmark, September 1992, pp. 155-157.

Keast, C.L., and C.G. Sodini. "A CCD/CMOS Based Imager with Integrated Focal Plane Signal Processing." Paper presented at the *1992 Symposium on VLSI Circuits*, Seattle, Washington, June 1992, pp. 38-39.

Lumsdaine, A., and J.K. White. "Accelerating Dynamic Iteration Methods with Application to Semiconductor Device Simulation." *Proceedings of the Copper Mountain Conference on Iterative Methods*, Copper Mountain, Colorado, April 1992.

Masaki, I. "Function-oriented Vision Chips for Factory Automation." *1992 International Conference on Industrial Electronics, Control, and Instrumentation*, San Diego, California, November 1992.

Masaki, I. "Vision-based Vehicle Guidance." *1992 International Conference on Industrial Electronics, Control, and Instrumentation*, San Diego, California, November 1992.

Reichelt, M., 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.

Seidel, M.N. "Analysis and Synthesis of Steady-State SC Networks." Submitted to *IEEE International Symposium on Circuits and Systems*, Chicago, Illinois, 1993.

Wyatt, J.L., Jr., C. Keast, M. Seidel, D. Standley, B. Horn, T. Knight, C. Sodini, H-S. Lee, and T. Poggio. "Analog VLSI Systems for Early Vision Processing." *Proceedings of the 1992 IEEE International Symposium on Circuits and Systems*, San Diego, California, May 1992, pp. 1644-1647.

Wyatt, J.L., Jr., C. Keast, M. Seidel, D. Standley, B. Horn, T. Knight, C. Sodini, and H-S. Lee. "Small, Fast Analog VLSI Systems for Early Vision Processing." Paper presented at the Conference on *Intelligent Vehicles '92*, Detroit, Michigan, June 1992, pp. 69-73.

Yu, P.C., and H-S. Lee. "A CMOS Resistive-fuse Processor for 2-D Image Acquisition, Smoothing and Segmentation." *1992 European Solid-State Circuits Conference*, Copenhagen, Denmark, September 1992, pp. 147-149.

Yu, P.C., and H-S. Lee. "A High-Swing 2-V CMOS Operational Amplifier with Gain Enhancement Using a Replica Amplifier." *1993 International Solid-State Circuit Conference*, San Francisco, California, February 1993.

Theses

Eshghi, K. *A System Demonstration of VLSI Early Vision Processors.* S.M. thesis, Dept. of Electr. Eng. and Comput. Sci., MIT, 1992.

Hakkarainen, J.M. *A Real-Time Stereo Vision System in CCD/CMOS Technology.* Ph.D. diss., Dept. of Electr. Eng. and Comput. Sci., MIT, 1992.

Keast, C.L. *An Integrated Image Acquisition, Smoothing and Segmentation Focal Plane Processor.* Ph.D. diss., Dept. of Electr. Eng. and Comput. Sci., MIT, 1992.

Lumsdaine, A. *Theoretical and Practical Aspects of Parallel Numerical Algorithms for Initial Value Problems, with Applications.* Ph.D. diss., Dept. of Electr. Eng. and Comput. Sci., MIT, 1992.

Osgood, G.T. *Evenly Sampling a Rotation Space.* S.B. thesis, Dept. of Electr. Eng. and Comput. Sci., MIT, 1992.

1.3 Parallel Algorithms for Device Simulation

Sponsors

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

Project Staff

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

Enormous computational expense and the growing importance of mixed circuit/device simulation, as well as the increasing availability of parallel computers, suggest that specialized, easily parallelized algorithms should be developed for transient simulation of MOS devices. In earlier work on the WORDS program (Waveform Overrelaxation Device Simulator), the easily parallelized waveform relaxation (WR) algorithm was shown to be a computationally efficient approach to device transient simulation even on a serial machine. However, the WR algorithm typically requires hundreds of iterations to achieve an accurate solution.

In order to use WORDS in a mixed circuit/device simulator, we have been investigating ways of making WORDS 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 (reducing the total computation by as much as a factor of 2 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.[8]

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

Through experiments, we found 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[9] as well as a waveform conjugate-direction approach.[10] Experimental results indicate that both approaches reduce 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.[11]

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, but 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 of model parameters, match well with experimental data for devices over a range of channel lengths from 0.90 μm to 0.16 μm. 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 0.16 μm.[12]

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 in portions of the MOS device which are likely to create hot carriers.

1.5 Circuit Simulation Algorithms

Sponsors

IBM Corporation
National Science Foundation
Fellowship MIP 88-58764

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

---

9 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.

10 A. Lumsdaine and J. White, "Accelerating Dynamic Iteration Methods with Application to Semiconductor Device Simulation," (Third place, best student paper competition) *Proceedings of the Copper Mountain Conference on Iterative Methods,* Copper Mountain, Colorado, April 1992.

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

12 K. Rahmat, J. White, and D. Antoniadis, "Simulation of Very Short Channel MOSFET's Including Energy Balance," *IEEE Trans. Comput.-Aided Des.*, forthcoming.

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

In our recent work, we are 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 need not be computed. An automatic method for determining the quasi-algebraic solution components has been derived, and experimental results demonstrate that this modified method reduces the number of computed clock cycles when applied to simulating closed-loop switching power converters.[13]

1.6 Parallel Circuit Simulation Algorithms

Sponsor

U.S. Navy - Office of Naval Research
Grant N00014-91-J-1698

Project Staff

Luis M. Silveira, Ricardo Telichevsky, Professor William J. Dally, Professor Jacob K. White

We are taking an application-oriented approach to developing parallel numerical algorithms and focusing on circuit and device simulation. Application programs and techniques are being developed for both massively parallel machines SIMD machines like the Connection Machine or MIMD machines like the Intel hypercube. In addition, we are also trying to understand fundamental aspects of the interaction between architecture and certain numerical algorithms.

For example, the direct solution of circuit simulation matrices is particularly difficult to parallelize, in part because methods like parallel nested dissection are ineffective due to the difficulty of finding good separators. For that reason, general sparse matrix factorization techniques are being studied and, in particular, the interaction between sparse matrix data structures, computer memory structure, and multiprocessor communication is being investigated. To focus this work, a special-purpose processor for circuit simulation, the Numerical Engine (NE), is under development. Preliminary design is complete, and register-transfer level simulation results indicate that the specialized processor can achieve up to 80% of its peak floating-point performance for sparse matrix factorization and nearly 90% of its peak performance on model evaluation.

Explicit integration methods avoid matrix solution and, therefore, are also interesting algorithms to use on parallel processors. For this reason, we investigated some properties of the recently developed explicit exponentially-fit integration algorithms. The results were not very encouraging, although our theoretical investigation yielded several new insights about these methods.[14]

1.7 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, 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

---

[13] 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.

[14] H. Neto, L.M. Silveira, J. White, and L.M. Vidigal, "On Exponential Fitting For Circuit Simulation," *IEEE Trans. Comput.-Aided Des.* 11(5): 566-574 (1992).

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 analyses. The system currently performs electromechanical analyses, such as calculating the capacitance versus pressure (or force) curve for a square diaphragm deformed by a differential pressure and can be used to calculate levitation forces in structures as complicated as a comb drive.[15]

1.8 Numerical Techniques for Simulating Josephson Junction Arrays

Sponsors

National Defense Science and Engineering
Graduate Fellowship
National Science Foundation
Fellowship MIP 88-58764

Project Staff

Joel R. Phillips, 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 to study 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 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, the magnetic moment, and the 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.[16]

1.9 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

---

[15] 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 Micro Electro Mechanical Systems Conference,* Fort Lauderdale, Florida, February 1993.

[16] 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.

based programs.[17] Manuals and source code for FASTCAP and FASTHENRY are available directly from MIT.

1.10 Techniques for Logic Synthesis, Formal Verification and Testing

Sponsors

Defense Advanced Research Projects Agency/
U.S. Navy - Office of Naval Research
Contract N0014-91-J-1698
IBM Corporation Grant/Fellowship

Project Staff

Professor Srinivas Devadas, Mazhar M. Alidina, Vishal L. Bhagwati, Kelly S. Bai, Stan Y. Liao, Jose C. Monteiro, Amelia H. Shen, Dr. Filip J. Van Aelten

1.10.1 Introduction

The design of microelectronic integrated circuits is a complex process because the final product must meet many different requirements and specifications. Specifications concerning the size and performance of a circuit are invariably placed on the design process. For example, reliability restrictions requiring that a circuit be testable for a chosen fault model place limitations on the design process. With an increasing number of transistors being fabricated within a small area on a silicon substrate, the power consumption of a circuit has become a very important design parameter. Indeed, within the next decade, the power consumed by a circuit may soon become the fundamental constraint limiting the density, i.e., transistors per unit area, of integrated circuits. At MIT, we are working on various aspects of the problem of designing integrated circuits for low power consumption. Some solutions include efficient power estimation methods as well as automated methods for logic optimization directly targeting low power consumption in the optimized circuit.

In addition, a circuit being designed must be checked for correct functionality and must meet performance goals. Formal verification methods can guarantee correctness of the design—however, a successful and practical method must be efficient and use a realistic circuit model. We are addressing various verification problems at the architectural, logic, and circuit levels. These problems include verification of pipelines, combinational logic verification, and asynchronous circuit verification. Verifying asynchronous circuits is particularly difficult since a sequential circuit must be directly checked for correct functionality, while in the synchronous case one can separate the combinational logic of the synchronous circuit from the memory elements. For instance, verifying the functionality of an asynchronous circuit requires timing simulation on the combinational logic to identify the presence of hazards that might cause its malfunction.

In the following sections, we will describe our recent work in the areas of design for low power consumption and formal verification techniques.

1.10.2 Design for Low Power Consumption

Probabilistic Power Estimation

We have addressed the problem of estimating the average switching activity of VLSI sequential circuits under random input sequences. We have developed a model that can be used to gauge the power dissipation of the sequential circuit and to make architectural or design-style decisions during the VLSI synthesis process.

Because switching activity is strongly affected by gate delays, we have used a general delay model in estimating switching activity. This model considers glitching at gate outputs in the circuit while modeling the inertial delay of a gate. Our method considers correlation caused at internal gates in a circuit due to reconvergence of input signals. In combinational circuits, uncorrelated input patterns with uniform switching rates are typically assumed, simplifying the problem of switching activity esti-

---

[17] K. Nabors, S. Kim, and J. White, "Fast Capacitance Extraction of General Three-Dimensional Structures," *IEEE Trans. Microwave Theory Tech.* 40(7): 1496-1507 (1992); K. Nabors, T. Korsmeyer, and J. White, "Multipole Accelerated Preconditioned Iterative Methods for Three-Dimensional Potential Integral Equations of the First Kind," *SIAM J. Sci. Stat. Comp.* Forthcoming; K. Nabors, and J. White, "Multipole-Accelerated 3-D Capacitance Extraction Algorithms for Structures with Conformal Dielectrics," *Proceeding of the 29th Design Automation Conference,* Anaheim, California, June 1992, pp. 710-715; 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 Design,* Santa Clara, California, November 1992, pp. 438-442; 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; M. Kamon, K. Nabors, and J. White, "Multipole-Accelerated 3-D Interconnect Analysis," (Invited paper) *Proceedings of the International Workshop on VLSI Process and Device Modeling (VPAD),* Nara, Japan, May 1993.

mation. However, in sequential circuits the input sequence applied to the combinational portion of the circuit is highly correlated because some of the inputs to the combinational logic are flip-flop outputs representing the state of the circuit. We have developed approximate methods to probabilistically estimate switching activity in sequential circuits that automatically compute switching rates and correlations between flip-flop outputs.

Ongoing work includes relaxing our current assumptions in sequential circuit power estimation and improving the efficiency of our algorithms.

Combinational Logic Synthesis

We are also investigating the use of retiming methods to improve power dissipation. In our experiments with probabilistic power analysis, we noted that power consumption depended significantly on the gate-level structure of a combinational or sequential circuit. Then we became interested in gauging exactly what logic structures corresponded to low power realizations. With this information we can tailor a logic optimization strategy to produce circuit topologies that consume the least amount of power.

We have developed several strategies to optimize a combinational circuit to lower its power consumption: (1) timing optimization and V_{DD} scaling, (2) selective collapsing, (3) don't-care-based minimization, and (4) implementation as a disjoint cover.

In Method 1, the circuit is restructured for maximum speed, and then the power supply voltage V_{DD} is lowered while increasing the delay of the circuit to the maximum allowed delay. Power consumption for the reduced V_{DD} circuit is lowered because power is proportional to V_{DD}^2. In Method 2, the delay of the circuit is maintained constant, but portions of the circuit that are not on the critical path are collapsed into two levels of logic, reducing power dissipation. In Method 3, observability don't-cares are used to minimize the internal gates in the circuit. Minimization alters the signal probabilities at the gate outputs in a manner such that power consumption is reduced. In Method 4, the circuit is implemented as a disjoint two-level cover, where each AND gate is disjoint from all other AND gates—this is a minimal power realization.

Experimental results on a variety of examples are extremely encouraging. Method 2 resulted in the circuits with the lowest power dissipation.

Sequential Logic Synthesis

We are exploring methods to optimize sequential circuits for low power consumption. In particular, we are developing methods for the state assignment of finite state machines that result in low power dissipation. For instance, if a state s_2 can be reached by many input vectors from state s_1, to minimize power, a good heuristic is to assign s_1 and s_2 uni-distant codes. The combinational logic of the state-encoded machine can be optimized using techniques described in the previous section to obtain a minimal power circuit realization.

1.10.3 Formal Verification Techniques

Pipeline Verification

It is difficult to verify a behavioral specification against a logic-level implementation since the logic-level implementation can be serial, parallel, or pipelined implementation of the behavioral description. We are investigating automatic methods for verifying pipelined implementations against unpipelined specifications. These methods ensure that each data transfer taking place upon the execution of any instruction in the unpipelined circuit also occurs in the pipelined circuit. In the case where data dependence hazards exist, verification requires checking that pipelined circuit contains bypass logic, as well as checking that this logic results in data transferring correctly. Similarly, branch instructions may require purging of the pipeline or disabling of register/memory writes.

A symbolic simulation method is being developed that will efficiently verify pipelined μ-processors against instruction set specifications.

Free Binary Decision Diagrams

We are developing a set of algorithms for (1) the probabilistic construction of free Binary Decision Diagrams from multilevel combinational logic circuits and (2) the Boolean manipulation of free Binary Decision Diagrams.

With the aid of a previously shown result, namely that the equivalence of two free Binary Decision Diagrams can be decided probabilistically in polynomial time, we have developed a set of polynomial-time algorithms to reduce, i.e., remove redundant nodes and subfunctions from a free Binary Decision Diagram. The algorithms have a time complexity of $O(n^k|G|\log(|G|))$ where n is the number of primary inputs to the free Binary Decision Diagram G, $|G|$ is the number of nodes in G,

and $k \geq 0$ is a constant. We have shown that reduced, ordered Binary Decision Diagram construction methods that exploit the strongly canonical form can be modified to apply to free Binary Decision Diagrams.

A *concat* algorithm that computes a free Binary Decision Diagram corresponding to the Boolean AND of two free Binary Decision Diagrams has been developed and is currently being implemented. This *concat algorithm* along with the k-reduce algorithm provides the basis for a Boolean manipulation package that can be used to solve many verification and testing problems.

Asynchronous Circuit Verification

We have addressed the problem of verifying that the gate-level implementation of an asynchronous circuit, with given or extracted bounds on wire and gate delays, is equivalent to a specification of the asynchronous circuit behavior described as a classical flow table.

We have developed a procedure to extract the complete set of possible flow tables from a gate-level description of an asynchronous circuit under fixed or bounded wire delay models. Given an extracted flow table and an initial flow table specification, we have implemented procedures to construct a product flow table to check for machine equivalence, under various modes of operation.

We have considered fundamental and non-fundamental modes of operation. We have taken into account single output change and multiple output change flow tables as well as single input change and multiple input change flow tables. Flow table extraction and equivalence checking procedures have been tailored for each mode of operation and each type of asynchronous flow table.

Under the discretized bounded delay model, where each gate can take on integer delay values within a range, we have developed implicit enumeration methods for verification. These Binary Decision Diagram-based methods do not require explicit representation of each possible flow table corresponding to the circuit with bounded gate delays or explicit enumeration of states in the circuit. These methods are more efficient than the corresponding ones for a continuous delay model.

1.10.4 Publications

Ashar, P., A. Ghosh, and S. Devadas. "Boolean Satisfiability and Equivalence Checking Using General Binary Decision Diagrams." *Integration: VLSI J.* 13(1): 1-16 (1992).

Ashar, P., S. Devadas, and K. Keutzer. "Gate-Delay-Fault Testability Properties of Multiplexor-Based Networks." *Formal Meth. VLSI Des.: Int. J.* Forthcoming.

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. Simul. Test.* Forthcoming.

Bryan, M.J., S. Devadas, and K. Keutzer. "Necessary and Sufficient Conditions for Robust Stuck-Open Fault Testability in Multilevel Networks." *IEEE Trans. Comput.-Aided Des.* 11(16): 800-803 (1992).

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.* Forthcoming.

Devadas, S. "Comparing Two-Level and Ordered Binary Decision Diagram Representations of Logic Functions." *IEEE Trans. Comput.-Aided Des.* Forthcoming.

Devadas, S., and K. Keutzer. "Validatable Nonrobust Delay-Fault Testable Circuits Via Logic Synthesis." *IEEE Trans. Comput.-Aided Des.* 11(12): 1559-1573 (1992).

Devadas, S., K. Keutzer, S. Malik, and A. Wong. "Computation of Floating Mode Delay in Combinational Logic Circuits: Practice and Implementation." *IEEE Trans. Comput.-Aided Des.* Forthcoming.

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

Devadas, S., K. Keutzer, and S. Malik. "A Synthesis-Based Test Generation and Compaction Algorithm for Multifaults." *J. Electron. Testing: Theory Appl.* Forthcoming.

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

Ghosh, A., S. Devadas, and A.R. Newton. "Heuristic Minimization of Boolean Relations Using Testing Techniques." *IEEE Trans. Comput.-Aided Des.* 11(9): 1166-1172 (1992).

Ghosh, A., S. Devadas, and A.R. Newton. "Sequential Test Generation and Synthesis for Testability at the Register-Transfer and Logic Levels." *IEEE Trans. Comput.-Aided Des.* Forthcoming.

Meeting Papers

Camposano, R., S. Devadas, K. Keutzer, S. Malik, and A. Wang. "Implicit Enumeration Techniques Applied to Asynchronous Circuit Verification." Paper presented at the 26th Hawaii International Conference on System Sciences, Wailea, Hawaii, January 1993.

Devadas, S., H-F. Jyu, K. Keutzer, and S. Malik, "Statistical Timing Analysis of Combinational Logic Circuits." Paper presented at the International Conference on Computer Design: VLSI in Computers and Processors, Cambridge, Massachusetts, October 1992.

Devadas, S., K. Keutzer, S. Malik, and A. Wang. Certified Timing Verification and the Transition Delay of a Circuit." Paper presented at the 29th Design Automation Conference, Anaheim, California, June 1992.

Devadas, S., K. Keutzer, S. Malik, and A. Wang. "Verification of Asynchronous Circuits with Bounded Wire Delays." Paper presented at the International Conference on Computer-Aided Design, Santa Clara, California, November 1992.

Ghosh, A., S. Devadas, K. Keutzer, and J. White. "Estimation of Average Power Dissipation in Combinational and Sequential Circuits." Paper presented at the 29th Design Automation Conference, Anaheim, California, June 1992.

Ghosh, A., A. Shen, S. Devadas, and K. Keutzer. "On Average Power Dissipation and Random Pattern Testability of Combinational Logic Circuits." Paper presented at the International Conference on Computer-Aided Design, Santa Clara, California, November 1992.

Van Aelten, F., S. Liao, J. Allen, and S. Devadas. "Automatic Generation and Verification of Sufficient Correctness Properties of Synchronous Processors." Paper presented at the International Conference on Computer-Aided Design, Santa Clara, California, November 1992.

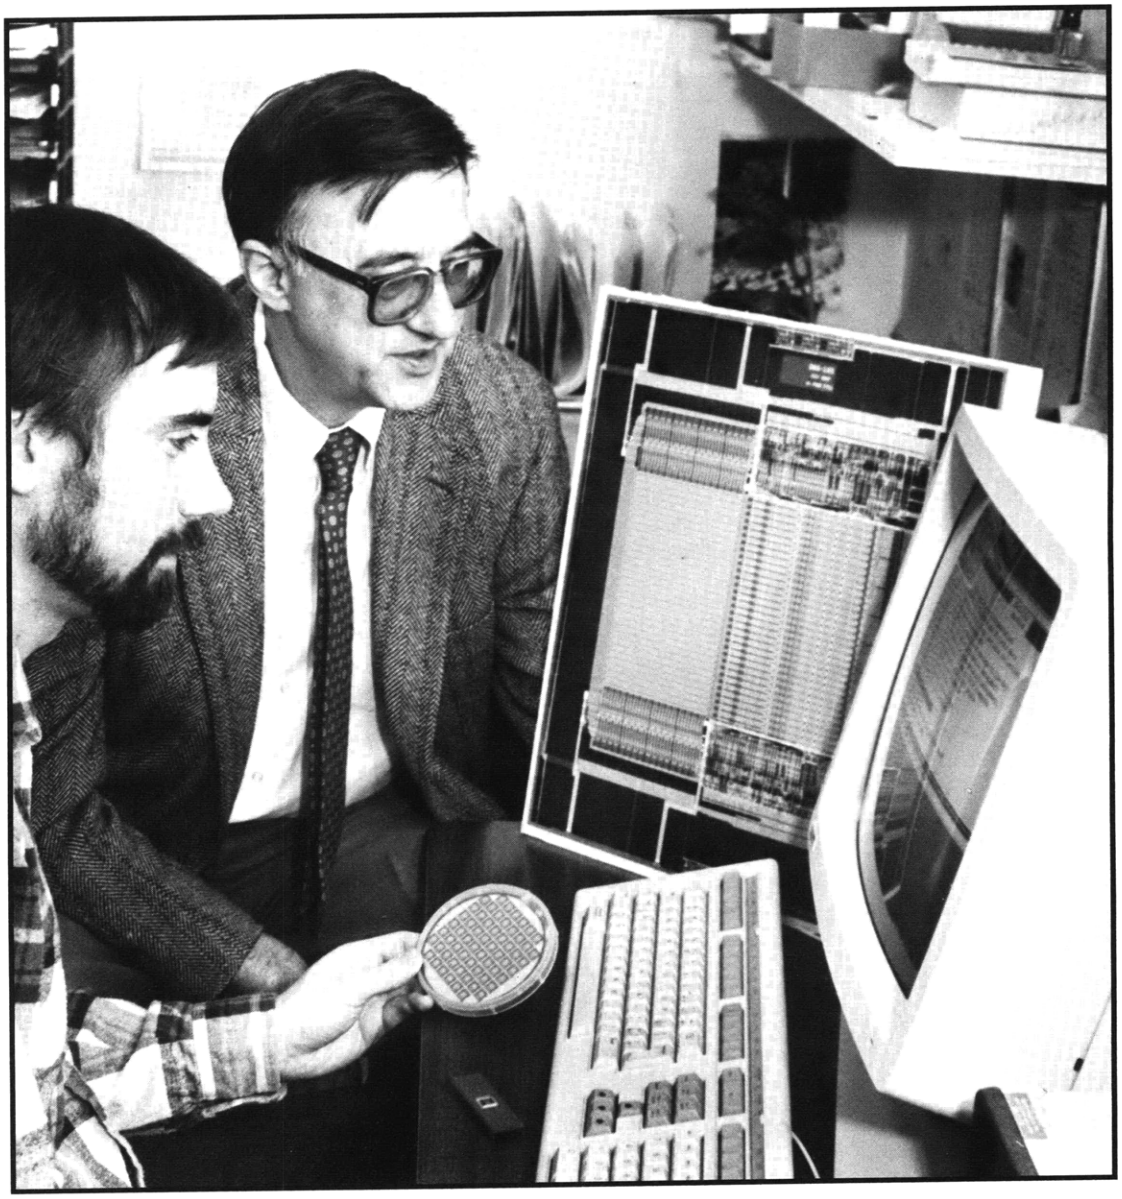

Professor Donald E. Troxel