

14. Custom Integrated Circuits

Academic and Research Staff

Prof. J. Allen, Prof. L.A. Glasser, Prof. B. Musicus, Prof. P. Penfield, Jr., Prof. J.L. Wyatt, Jr., Dr. M. Matson

Graduate Students

R. Armstrong, I. Bain, D. Baltus, C. Bamji, C. Hauck, H. Jeong, C.-C Kuo, S. McCormick, K. Nabors, P. O'Brien, L. Seiler, W. Song, D. Standley, H.-N. Tan, R. Wright, Q. Yu, J. Wroclawski, C. Zukowski

14.1 Conversion of Algorithms to Custom Integrated Circuits

U.S. Air Force - Office of Scientific Research (Contract F49620-84-C-0004)

Jonathan Allen, Robert Armstrong, Donald Baltus, Cyrus Bamji, Charles Hauck, Steven McCormick, Larry Seiler

This project is devoted to the study of fundamental techniques used in the development of both algorithms and representations for the conversion of functional specifications to custom integrated circuit layout specifications. Included in this area are five major areas of investigation. These are the development of functional specification languages, the development and utilization of architectural exploration techniques, methods for the generation and composition of macro cells, the characterization and enforcement of circuit performance constraints, and the development of special purpose architectures for both digital signal processing and custom VLSI design. In addition, some topics related to architectural exploration and the development of special purpose architectures are considered in sections 14.2-14.4 of this report and are under the supervision of Professor Musicus. Some topics in circuit performance characterization and constraints are also dealt in section 14.7 under the supervision of Professors Glasser and Wyatt. In this section, we describe the remaining activities being carried out under this contract.

At the level of functional specification languages, we continue to build on earlier experience with the language SRL which was devised in a doctoral thesis by Gary Kopec as a means to characterize digital signal processing algorithms in a semantically rigorous way. Current work aims at the provision of specifications in the form appropriate to data flow languages, such that compilation procedures can automatically reveal the latent parallelism in these algorithms. This is a long term fundamental investigation with the intent of providing a natural coupling from the input functional specification to a subsequent phase of architectural exploration whereby the designer can explore performance variations through functionally invariant transformations.

A variety of techniques have been developed for such architectural exploration, and we have been utilizing the retiming technique for the systematic exploration of pipeline array multiplier performance.¹ In this work, we have first used techniques of asymptotic complexity to reveal the opportunity for optimal pipelining. Following the development of this perspective, we have gone on to use retiming to show how to obtain all possible versions of pipelining within an array multiplier, the degree of pipelining being characterized by a single parameter that allows the designer to naturally accommodate to the circuit performance of the target technology being employed. Once a particular point in the design space is obtained, it is then possible to generate that particular array multiplier using a new tool developed by us called the regular structure generator.² Using the regular structure generator, we are able to compose large circuits from a set of constituent cells by means of algorithmically utilizing three design specification files. The first file characterizes the constituent leaf cells to be used in the design. The second file characterizes the overall architecture which utilizes these leaf cells to produce the final design, and the third file specifies all possible allowed local interconnects between the leaf cells. This tool has been used to generate many regular and quasi-regular designs, and is very well suited to the generation of pipelined array multiplier configurations. Once a particular layout is generated, its performance can be characterized by the techniques described below.

In order to gain further understanding of the techniques needed to build important macro cells, we have focused on the development of procedural techniques for the generation of high performance floating point adders and multipliers. An important part of this activity has been to emphasize new techniques to insure the correctness of the resulting design achieved through such a composition process. In the past, the main checks that were used for this purpose involved only design rule checks and possible computation of proper bus sizing. By introducing a new rigorous formalism, we have been able to characterize cells that do not involve any races and hazards.³ Most importantly, we have been able to show that it is possible to specify a rigorous technique for composing these well formed cells into larger cells which in turn are guaranteed to be free from races and hazards. In this way, a fundamental design discipline is provided that frees the designer from the need to examine a large circuit for possible race and hazard conditions, which are very difficult to detect. A new emphasis has also been placed on the development of design techniques under circuit performance constraints. At present, designers of custom circuits must extract the circuit from the final layout in order to characterize its performance. In this project, we are aiming to provide top down design techniques that will allow the manipulation of a design subject to circuit performance constraints. In order to allow this degree of flexibility, it is necessary to constrain the final layout methodology, and we have provided one such constrained layout technique for CMOS, and two for NMOS which allow us to be able to predict circuit performance at the architectural level without the need for complete instantiation of the layout. We expect to build on macro modelling techniques developed under this contract which have been characterized in section 14.5 by Professor Glasser.

A long time emphasis of this project has been the comprehensive construction of a placement and interconnect program, called PI, which is capable of placing and routing arbitrary size rectangular modules orthogonally related with interconnect on all four sides in minimum area. During this contract period, new techniques for power bus routing⁴ and compaction of the constituent cells have been completed.

A major part of the effort under this contract has been focused on techniques for providing and characterizing circuit performance. Professors Glasser and Wyatt have been focusing on techniques for macro modelling, path delay optimization, and waveform bounding techniques, and a major new program called EXCL has been completed as a Master's thesis by Steven McCormick.⁵ In this program, an input layout specified in CIF is converted to a circuit representation by high precision extraction of resistances, capacitances, and transistors from the layout in a form suitable for SPICE simulation. This program has been very successful and is widely used both here at MIT and at other laboratories. A default mode of this program is also able to extract the fundamental topology of the circuit in a form appropriate for unit delay logic simulation. Finally, in a thesis project by Wai Lee, a very high performance 32 bit gallium arsenide multiplier is being designed, providing an opportunity for the utilization of novel circuit techniques in a leading edge technology utilizing gallium arsenide MODFETS.

The last major emphasis of this contract has been on the development of special purpose architectures for both computer aided design of integrated circuits and digital signal processing tasks. A doctoral thesis by Larry Seiler⁶ has been completed on the design and construction of a special purpose accelerator for the computation of design rule checking. In this technique, the layout is rasterized and design rule checking is performed through pattern matching techniques utilizing four distinct custom integrated circuits which were designed for this accelerator. Speed ups of the order of 100 over conventional single sequence machines have been obtained, and addressing efficiencies have also been developed which avoid the necessity for instantiating large parts of the layout at any one time. This work also includes extensive performance measurements of existing design rule checkers as well as the new accelerator architecture. In section 14.2 in this section Professor Musicus has described a number of investigations under this contract focusing on special purpose architectures for signal processing, and we are also starting to study the design of a highly efficient but general architecture to serve as a base for linear algebra computations. It has been observed that many important digital signal processing operations, as well as circuit simulation algorithms, can be represented in terms of linear algebra operations, so that such an architecture can be expected to be exceedingly useful. There is a great deal of parallelism that can be exploited, but there is also the necessity to deal realistically with problems of sparse matrix representation. We expect this work to be a continuing major emphasis of this project.

During this contract period the Principal Investigator completed two overviews of the relation

between computer architecture and digital signal processing.^{7,8}

References

1. C.E. Hauck, C.S. Bamji, and J. Allen, "The Systematic Exploration of Pipelined Array Multiplier Performance," Proceedings of the International Conference on Acoustics, Speech, and Signal Processing, March 1985.
2. C.S. Bamji, C.E. Hauck, and J. Allen, "A Design-by-Example Regular Structure Generator," Proceedings of the 22nd Design Automation Conference, June 1985.
3. R. Armstrong, "Procedural Design of a High Speed Floating Point Arithmetic Unit," M.S. Thesis, Department of Electrical Engineering and Computer Science, M.I.T., March 1985.
4. A.S. Moulton, "Routing the Power and Ground Wires on a VLSI Chip," M.S. Thesis, Department of Electrical Engineering and Computer Science, May 1984; also VLSI Memo 84-184, June 1984.
5. S.P. McCormick, "EXCL: A Circuit Extractor for IC Designs," Proceedings of the 21st Design Automation Conference, June 1984; expanded version available as M.S. Thesis.
6. L.D. Seiler, "A Hardware Assisted Methodology for VLSI Design Rule Checking," Ph.D. Thesis, Department of Electrical Engineering and Computer Science, M.I.T., February 1985.
7. J. Allen, "Computer Architecture for VLSI Speech Processing," Proceedings of SpeechTech '84, April 1984.
8. J. Allen, "Computer Architecture for Digital Signal Processing" IEEE Proc. 73, 5, 852-873 (1985).

14.2 Automatic Design of Systolic Arrays

U.S. Air Force - Office of Scientific Research (Contract F49620-84-C-0004)

Bruce R. Musicus, Hong Jeong

A major area of research which we have been pursuing is in the relationship between signal processing algorithms and VLSI architectures. Digital signal processing is an application which is particularly well suited to custom VLSI implementation. These algorithms are often composed of elementary calculations which are decomposable into highly parallel and independent processing streams organized in a very regular structure. As a result, it is often possible to implement these algorithms with a highly decentralized parallel processing architecture.

We are working on new techniques for systematically mapping arithmetic calculations onto deeply pipelined, and highly parallel systolic architectures. We capture the flow and timing of data in an algorithm through a new matrix polynomial description which strongly resembles linear state space equations. Inputs are combined into various internal states after some delays, and then outputs are formed from the inputs and the states. These internal states may be mapped to a physical array of processors in a variety of time-independent or time-varying ways, constrained only by certain causality restrictions. Each such mapping yields a new pattern of data flow through the physical processor net, and thus results in a new "architecture" for executing the algorithm. Retiming techniques suggested by Charles Leiserson can be shown to correspond to diagonal similarity transforms of the system. More complex geometric transformations which

remove broadcasting, alter fan-in and fan-out, or radically change data flow behavior, can also be derived by systematically rotating the physical processors through the algorithm state variables. We view this research as a first step towards a more general examination of the feasibility of parallel networks, including analysis of I/O bottlenecks, and memory use.

Publications

Jeong, H., "Modelling and Transformation of Systolic Network," M.S. Thesis, Department of Electrical Engineering and Computer Science, M.I.T., May 1984.

14.3 Parallel Algorithms and Architectures for Solving Elliptic Partial Differential Equations

U.S. Air Force - Office of Scientific Research (Contract F49620-84-C-0004)

Bruce R. Musicus, Bernard Levy, Chung-Chieh Kuo

Solving spatially varying elliptic partial differential equations is an especially difficult and numerical intensive problem. Fortunately, the structure of these equations lends itself to decentralized, parallel solution. Once discretized on a regular grid, with a separate processor allocated to each node, simple iterative schemes such as Successive Over-Relaxation (SOR) only require communication between processors at adjacent nodes. Unfortunately, these schemes tend to converge at a slow geometric rate.

We have developed a new acceleration method, called LASOR, which only uses information locally available to each processor to optimally accelerate the convergence of the relaxation method. Spatially varying relaxation parameters are chosen in such a way as to force the lowest and highest frequency modes of the iteration into the form of nearly pure sinusoids, having eigenvalues which are relatively easy to calculate, and which enjoy relatively fast convergence rates. Computer simulation has demonstrated the superior convergence characteristics of this algorithm over conventional SOR for spatially varying elliptic differential equations.

Publications

Kuo, C.J., B.C. Levy, and B.R. Musicus, "The Specification and Verification of Systolic Wave Algorithms," in VLSI Signal Processing, (IEEE Press 1984) pp. 271-281.

14.4 Parallel Algorithms and Architectures for Computer Vision

U.S. Air Force - Office of Scientific Research (Contract F49620-84-C-0004)

Bruce R. Musicus, Hong Jeong, Rosalind Wright

Classical image processing algorithms such as contrast stretching, color correction, histogram

construction, and so forth, rely on uniform point-by-point processing of some predefined region in the picture. Interaction between processing steps is minimal, and the flow of data through the memory and processor is completely predictable. As a result, it is easy to decompose such algorithms into highly parallel and regular computations. Unfortunately, more sophisticated algorithms, such as boundary finding or object recognition, do not appear to decompose so nicely. These algorithms typically alternate between block data processing steps and decision making steps. The block data processing steps often involve operations such as filtering, contrast stretching, or table-lookup, which can be easily sped up in a parallel architecture. After processing a small neighborhood of data, however, these algorithms usually stop to analyze the progress of the computation in order to decide what to do next. Such adaptive, data-directed processing results in unpredictable bursts of data access and computation. Building highly parallel or pipelined systems to accelerate such systems is quite difficult.

We are investigating three different approaches to this problem. The simplest approach starts with conventional algorithms designed for sequential machines. These programs can be partitioned into more or less parallel streams in various ways. For example, individual processors may be responsible for sub-blocks of the picture, and then additional processing is used to "seam" together the blocks. The disadvantage of these ideas is that much effort is wasted on trying to combine partial information from multiple processors, and the quality of the result may suffer from the lack of global viewpoint.

It is interesting to try to find new algorithms for line and object finding which are more inherently parallel, and which do not force the use of arbitrary boundaries between processing entities. One technique is to try to redesign the algorithms so that several processors are used to "crawl" around the picture, each examining a local section of the image around their current position, and building up a global model of the image in memory as they travel. By leaving partial models tagged onto pixels, it may be possible to efficiently coordinate multiple processors, prevent duplication of effort, and efficiently combine results. A more speculative idea is to search for highly distributed relaxation-like algorithms for pixel-per-processor architectures, in which processors agree through distributed protocols on where object boundaries lie, how to fit lines through these boundaries, and so forth. We are working towards algorithms in which lines and possibly even objects will "crystallize" out of the morass of image data through consensus of large numbers of processors.

To support this effort, we are writing an architectural simulator which will allow us to test out parallel algorithms on simulated parallel architectures. Such detailed examination of the algorithmic ideas will help to isolate the communication and synchronization bottlenecks in the algorithms, and also test their performance on image data.

14.5 Very Large Scale Integrated Circuit Research

National Science Foundation (Grant ECS81-18160)

U.S. Air Force - Office of Scientific Research (Contract F49620-84-C-0004)

Lance A. Glasser, Mark D. Matson

We are continuing our work in modeling and optimization of digital MOS VLSI circuits. We have developed sophisticated models describing the power consumption, output response delay, output response waveshape, and input capacitance of MOS combinational logic gates. The models account for the effects of input waveform shape on logic gate behavior. We have also created a general macromodeling software support package, and have interfaced our models to this package. The model predictions are typically within several percent of SPICE estimates, a major improvement over the resistor-capacitor models used in most current timing simulators. Moreover these high accuracies come at a small price in computation time because we model the response of the entire logic cell, rather than using a more sophisticated transistor model and then having to compute the transistor's interactions to obtain the cell's response. The accuracy and computational speed make the models pertinent to simulation and optimization applications.

Our optimization research has produced a CAD tool for optimizing digital MOS designs. This tool finds the transistor sizes that minimize circuit power consumption subject to constraints on signal path delays. The optimization algorithm exploits properties of the digital MOS domain to convert the primal optimization problem into a dual form which is much easier to solve. The principal advantage of this approach is that it permits a divide and conquer strategy where the original problem can be broken into many smaller subproblems which can then be solved almost independently. As a result the CAD tool can optimize a circuit in roughly the amount of time needed to perform a transistor level simulation of the circuit.

At present we are also investigating the application of our optimization techniques to the design of digital MOS circuits subject to constraints on immunity to electrical noise.

Publications

Glasser, L.A. and L. Hoyte, "Delay and Power Optimization in VLSI Circuits," Proceedings of the 21st Design Automation Conference, June 1984.

Matson, M.D., "Macromodeling of Digital MOS VLSI Circuits," Proceedings of the 22nd Design Automation Conference, June 1985, accepted for publication.

Matson, M.D., "Optimization of Digital MOS VLSI Circuits," Proceedings of the Chapel Hill Conference on VLSI, May 1985.

14.6 A Circuit Theory for Digital VLSI Systems

National Science Foundation (Grant ECS81-18160)

Lance A. Glasser, Paul Penfield, Jr., John L. Wyatt, Jr., Isaac Bain, William Song, John Wroclawski, Charles Zukowski

Analysis techniques appropriate for use on very large scale integrated (VLSI) circuits were examined as part of this research effort. The key obstacle to the effective analysis of VLSI circuits is the huge numbers of transistors which need to be examined. Our goal was to discover specialized techniques that exploit the way these circuits were typically designed in order to make analysis and simulation computationally tractable. This objective lead us in two major directions. The first is the study of "waveform bounding," a technique in which the exact response of a network is not examined but rather bounds on the response are achieved. The other research direction was the exploitation of the design methodology to help provide early error detection.

Significant results were achieved in the area of waveform bounding. A number of new theorems increased our understanding of the fundamentals of this domain. Central among the results was a discovery of the connection between waveform bounding and "waveform relaxation," another powerful technique receiving research attention today. Equally important for the eventual use of this technique in practical computer-aided design (CAD) tools, the range of applicability and tightness of the bounds were both increased. A prototype simulator was developed to test the theory.

In the area of exploiting the self-imposed rules that a designer uses to build a complex microcircuit, we developed a working prototype CAD tool to check designs as they are entered. First, the engineer describes, to the computer, the self-imposed rules he or she plans to obey, and then, as the design is done, those rules are checked interactively. Specialized languages (embedded in LISP) for describing methodologies and circuits were developed.

Publications

- Wroclawski, J.T., "DTOS, An Integrated VLSI Design System," S.M. Thesis, Department of Electrical Engineering and Computer Science, M.I.T., September 1984.
- Glasser, L.A. and C.A. Zukowski, "Communication Density Considerations in Multiprocessors," J. VLSI and Computer Systems, submitted for publication.
- Bain, L.I. and L.A. Glasser, "Methodology Verification of Hierarchically Described VLSI Circuits," IEEE Trans. Computer-Aided Design, submitted for publication.
- Wyatt, J.L., C.A. Zukowski, P. Bassett, P. Penfield, Jr., and L.A. Glasser, "The Waveform Bounding Approach to Timing Analysis of Digital MOS IC's," VLSI Memo 83-148, July 1983.
- Song, W. and L.A. Glasser, "Power Distribution Techniques for VLSI Circuits," VLSI Memo 83-156, November 1983.
- Zukowski, C.Z. and L.A. Glasser, "Bounding the Response of One-Way MOS Logic Gate Models," VLSI Memo 84-176, May 1984.

Bain, I., "Methodology Verification of Hierarchically Described VLSI Circuits," VLSI Memo 84-195, July 1984.

14.7 Waveform Bounding for Fast Timing Analysis of MOS VLSI Circuits

Defense Advanced Research Projects Agency (Contract N0014-80-C-0622)

National Science Foundation (Grant ECS83-10941)

John L. Wyatt, Jr., Paul Penfield, Jr., Lance A. Glasser, Charles Zukowski, Quinjian Yu, Han-Ngee Tan, Keith Nabors, David Standley, Peter O'Brien

The goal of this research is to enable rapid computation of signal propagation delay in monolithic MOS logic circuits for purposes of timing analysis. Our approach is intended as an alternative to the two methods that are currently standard practice: "exact" numerical simulation and the use of approximate delay formulas. It is an attractive alternative in many cases because of problems inherent in both those approaches.

For "exact" numerical solution the computation time becomes quite large even for small logic circuits with, say, 200 transistors: it is utterly prohibitive for modest sized chips or major portions of chips with 10,000–500,000 transistors. Exact solution not only requires more computer time than the digital circuit designer can normally afford, it also gives more detailed answers than he can normally use.

At the other extreme, approximate delay formulas require relatively little computer time and yield a minimum of detailed information. They can be useful guides to a designer, but there is usually no known absolute limit to the error that can result from their use.

Our research over this past year has had two focuses, one narrow and one broad. The narrowly focused work has been concentrated on bounds for signal delay in linear RC models for on-chip interconnect. Taking as a starting point the work of Rubenstein, Penfield and Horowitz,¹ we have been able to include more general networks than RC trees driven by voltage step inputs. We have also succeeded in reducing the region of uncertainty in the original bounds for certain classes of networks of practical interest.

One extension we have completed is a method of bounding the response of RC meshes, which are more general than RC trees in that resistor loops are allowed. These networks are important in practice 1) as models for the gates of large MOS pad driver resistors, 2) whenever linear resistor models are used for transistors in logic gates of CMOS pass gates, as in Chris Terman's program RSIM, and 3) to model interconnect networks with closed loops sometimes created by automatic routing programs. Another successful extension is to networks with resistive paths to ground, which are appropriate models for, e.g., interconnect to bipolar logic gates.

Tighter bounds have been achieved for unbranched lines and certain classes of RC trees by exploiting slew rate limits on the node voltages and exploiting the spatial convexity of interconnect voltage during transients in a novel way.

The more broadly focused and ambitious project, carried out by Charles Zukowski in his doctoral research, has been to consider bounds for the behavior of general digital MOS circuit models.

He has shown that typical digital MOS circuits have monotonicity properties sufficient to guarantee that the response of tightly coupled subcircuits can be bounded by the response of simpler models, e.g., linear or piecewise linear circuits. He has also shown that Waveform Relaxation techniques developed for the efficient exact analysis of digital MOS circuits can be extended to include the relaxation of bounds.² As a result, bounds on the responses of tightly coupled subcircuits can be used iteratively to efficiently produce rigorous bounds on the responses of entire VLSI circuits. More powerful monotonicity properties have been shown for intermediate-complexity restoring logic gate models, leading to more efficient bounds in this important special case.³

In the broad project, bounding is considered as a framework within which to combine the latest algorithms for VLSI simulation in a form where uncertainty can be measured and therefore managed efficiently. The use of relaxation techniques allows partitioning to exploit latency and achieve a computation time that scales roughly linearly with circuit size. The use of simplified models to calculate response bounds allows efficient approximating algorithms, such as piecewise linear simulation, to be exploited. Bounding work for linear circuits can also be incorporated as one step in generating bounds on the responses of more complex circuits.

Extensive examination of uncertainty accumulation in digital MOS circuits, when using the proposed bounding strategy, has shown that the behavior of a large portion of a typical VLSI circuit can be feasibly bounded. A limited amount of experimental evidence gathered for combinational logic circuits has confirmed this conclusion. Bounding simulators can potentially be used to find the remaining small portions of a circuit requiring exact analysis, allowing entire circuits to be considered in one bounding framework. Possible additional uses for bounding algorithms such as worst case design and input-independent analysis have also been considered in some detail.

References

1. J. Rubinstein, P. Penfield, Jr., and M.A. Horowitz, "Signal Delay in RC Tree Networks," *IEEE Trans. Computer-Aided Design*, CAD-2, 3, 202-211 (1983).
2. C.A. Zukowski, "Circuit Simulation with Iterating Bounds," VLSI Memo 84-159, January 1984.
3. C.A. Zukowski and L.A. Glasser, "Bounding the Response of One-Way MOS Logic Gate Models," VLSI Memo 84-176, May 1984.

Publications

- Zukowski, C.A. and J.L. Wyatt, Jr., "Sensitivity of Nonlinear 1-Port Resistor Networks," IEEE Trans. Circuits, Syst., CAS-31, 1048-1051, December 1984; also RLE Memo 84-3, August 1984.
- Wyatt, J.L., Jr., "Monotone Sensitivity Theorem for Nonlinear, Nonuniform RC Transmission Lines, with Application to Timing Analysis of Digital Integrated Circuits," IEEE Trans. Circuits Syst., CAS-32, 28-33, January 1985; also RLE Memo 84-2, August 1984.
- Wyatt, J.L., Jr., "Signal Delay in RC Mesh Networks," IEEE Trans. Circuits Syst., CAS-32, June 1985, to be published; also VLSI Memo 84-196, August 1984.
- Wyatt, J.L., Jr. and Q. Yu, "Signal Delay in RC Meshes, Trees and Lines," Proceedings of the IEEE International Conference on Computer-Aided Design, Santa Clara, California, November 1984, pp. 15-17; also VLSI Memo 84-198, August 1984.
- Wyatt, J.L., Jr., Q. Yu, C. Zukowski, N.H. Tan, and P. O'Brien, "Improved Bounds on Signal Delay in MOS Interconnect," Proceedings of the IEEE International Conference on Circuits and Systems, Beijing, China, June 1985, accepted for publication; also VLSI Memo 85-221, January 1985.
- Yu, Q. and J.L. Wyatt, Jr., "Improved Bounds on Signal Delay in RC Trees Using Inequalities on the Derivatives of Node Voltages," VLSI Memo 84-197, August 1984.
- Tan, N.H. and J.L. Wyatt, Jr., "Time Optimal Trajectories Associated with Voltage Bounds in RC Tree Networks," VLSI Memo 84-205, October 1984.
- Zukowski, C.A., "Bounding Enhancements for VLSI Circuit Simulation," Ph.D. Thesis, Department of Electrical Engineering and Computer Science, M.I.T., 1985.
- Zukowski, C.A., "Relaxing Bounds on Linear RC Mesh Circuits," VLSI Memo 84-166, April 1984.

