

15. Custom Integrated Circuits

Academic and Research Staff

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

Graduate Students

R. Armstrong, D. Baltus, C. Bamji, L. Brocco, C. Hauck, Terry Hohol, C.-C. Huo, H. Jeong, S. Lin, A. Malamy, S. McCormick, K. Nabors, P. O'Brien, M. Reichelt, L. Seiler, C. Selvidge, W. Song, D. Standley, R. Wright, J. Wroclawski, C. Zukowski

15.1 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, Lynne Brocco, Charles Hauck, Sching Lin, Steven McCormick, Mark Reichelt, Larry Seiler

The goal of this project is the development of CAD techniques for the design of high-performance custom integrated circuits. Both architectural performance (parallelism) and circuit performance (speed and power dissipation) are considered, and a major theme is to coordinate these two aspects of performance in an optimal way.

In order to achieve these goals, two fundamental emphases are maintained. One is the search for insightful and effective representations at several levels, ranging from the functional view to architecture, logic, circuits, and layout. The aim of design is to coherently specify all the different levels of representation and to maintain a consistent alignment between them. Furthermore, there is a need to readily transform from one specification to another, and to propagate changes in one level of representation to all other levels of representation modified by this change.

Another major emphasis is the notion of design exploration. A completed design can be perceived as satisfying well-formedness constraints at all the different representational levels. For example, the logic level must be appropriately simulated, and layout must satisfy design rules. Nevertheless, there is still a great deal of design freedom, and it is desirable to be able to explore the alternatives consistent with these constraints, and to assess them in terms of the overall system requirements. Accordingly, a major emphasis in this project is the design of CAD tools which permit rapid exploration of design alternatives in a way that utilizes transformations between representations, but leaves the input desired functionality invariant.

Several projects are devoted to the study of the specific aspects of these themes. In earlier work,¹ the interaction of architectural exploration and circuit performance was carefully studied in the context of array multiplier design. In that work, it was realized that the use of retiming transformations allows a degree of pipelining at a granularity level which actually decreases performance, since too much time is devoted to latching results, and not enough is spent on combinatoric calculations. This perspective has now been extended in the Master's thesis of Charles Hauck, where the predictions of asymptotic complexity theory are contrasted with optimal design obtained through careful circuit simulation and layout. Initial results show that the bounds obtained in asymptotic complexity do not provide a realistic view of what can be obtained in highly optimized circuits. Nevertheless, a goal of this work has been to investigate the possibility of specialized asymptotic results that might serve as better predictors of area and delay in realistic designs.

For some time, major emphasis has been placed on CAD tools for the generation of layout corresponding to quasi-regular structures. The first accomplishment was a highly successful design-by-example PLA generator, which readily allows for a wide variety of circuit styles and technologies. This design approach has been generalized to form the regular structure generator,² which can generate a variety of regular structures including memory arrays, arithmetic logic units, and multipliers, as well as PLAs. This state-of-the-art tool has proven extremely powerful, and includes a design language which permits logic-dependent optimizations. A recent improvement utilizing this linguistic capability has been the inclusion of folding for PLA structures, as well as optimal packing for the interconnect in Weinberger arrays.

Two current projects seek to further enhance the capability of a regular structure generator. In one study, circuit modelling techniques are used to characterize the internal delay of system modules, as well as their port behavior, with respect to interconnection loading. Using this model, a composition program will permit accurate prediction of overall delay in a circuit composed of these constituent modules, including the effects of multiple nesting levels. This capability is highly useful for design exploration, since it permits the rapid assessment of the overall timing of complicated quasi-regular structures, without the need for comprehensive circuit simulation (particularly for large layouts).

Another area of enhancement for the regular structure generator has been the provision of compaction. Many currently available compaction strategies do not provide layouts that satisfy the human designer. Often, gratuitous jogs are introduced into the layout, and the resulting area is too large. A new cell compactor has been designed which not only minimizes the overall bounding box, but also minimizes the areas of the individual rectangles in a way that can be selectively controlled on differing layout layers. This program has been highly successful, and is now being extended to satisfy the needs of hierarchical compaction that is appropriate for the regular structure generator. For large and composed circuits, it is not desirable to compact the

entire circuit, but rather to compact the individual cells so that they can be connected with proper pitch alignment at the boundaries. Hierarchical compaction is intended to compact each constituent cell separately, but under the constraint of the previously mentioned pitch matching. This is an unsolved problem, but two efforts are currently devoted to achieving such a compactor.³ An initial result has been the development of a new representation for a module's periphery (called a donut). It has been shown that this representation contains all of the information needed by a hierarchical compactor. It also allows for the successful compaction of each cell's interior. The design of a complete hierarchical compactor will build on this result.

Two major projects are devoted to the circuit's relationship to layout, with substantial emphasis on parasitics. In one effort, a tool has been built to transform directly from a circuit representation to an efficient layout that will allow for optimized device sizing. When individual device sizes vary greatly, it becomes difficult to construct a satisfactory placement strategy that also permits efficient routing. In order to constrain this task, several target layout styles have been adopted, and transformation strategies have been coded to provide not only for compact layout, but also for optimal performance based on appropriate device size and minimization of parasitic elements leading to delay.

Another interesting project builds on our experience with high accuracy circuit extraction. In this project, the goal is to examine noise coupling between fine-line structures using highly aggressive design rules. Direct solution in the time domain, using electromagnetic theory, resulted in a substantial understanding of this problem. These results are now being modelled in the circuit domain with a perspective towards characterizing and controlling the noise coupling that is so abundant in these structures. It is contemplated that new circuit design styles resilient to these perturbations may be provided, as well as new fabrication processes that minimize the noise.

Finally, there is an extremely fundamental project aimed at providing a formal base for design exploration over several levels of representation simultaneously. Circuit and layout representations are represented as formal grammars that satisfy the constraints at each level of representation. By using grammar intersection techniques, a coupling grammar is obtained to unify these two levels of representation. Such a formal representation provides the needed representation to perform design exploration in a unified and coordinated way, while respecting the constraints at both the circuit and layout levels. This formalism is expected to provide the intellectual backbone for design representation and exploration. It also builds on extensive experience with formal languages and constraint representation.

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. S.L. Lin and J. Allen, "MINPLEX – A Compactor that Minimizes the Bounding Rectangle and Individual Rectangles in a Layout," Proceedings of the 23rd Design Automation Conference, June 1986.

15.2 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 new acceleration methods, called LASOR, which only use information locally available to each processor to optimally accelerate the convergence of these relaxation methods. 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 these algorithms over conventional SOR for spatially varying elliptic differential equations. We are also investigating multigrid methods and other techniques based on frequency domain analysis for quickly solving such equations using highly parallel machines.

Publications

Kuo, C.–C., B. Levy, and B. Musicus, "A Local Relaxation Method for Solving Elliptic PDE's on Mesh–Connected Arrays," *SIAM J. Stat. Sci. Comp.*, submitted for publication.

15.3 Architectures for Transactions Oriented Signal Processing

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

Bruce R. Musicus

Traditional research in signal processing architectures has focused on the lowest level of signal processing, in which raw waveforms or images are processed into "cleaner" data using

numerically intensive computations. These problems are interesting because of the staggering computational bandwidth required, and because the geometric regularity of the data and the processing allows a multitude of interesting, massively parallel solutions. Our latest focus, however, has been a more difficult class of problems which arise in signal interpretation tasks, where an abstract, structured description must be extracted from a measured waveform or image. A combination of numerical computation as well as rules and database techniques are often used to find the solution. The irregularity, data dependence, and unpredictable structure of the final signal description vastly complicates the use of extensive parallelism in the signal processor. Nevertheless, these applications often exhibit a transactions-oriented structure, in which large streams of data "packets" are processed through a pipeline of operations, with limited feedback and limited interaction between the processing of different packets. Our goal has been to examine applications in radar and image processing which fall into this area, and to develop guidelines for computer architectures to support this processing.

15.4 Parallel Algorithms and Architectures for Computer Vision

Analog Devices, Inc.

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 segmentation, 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.

To attack this problem, we have chosen a particular image processing problem. We are trying to devise algorithms for extracting the layers of a VLSI chip from an optical image of the chip. The advantage of this problem is that we can potentially develop good models of the VLSI fabrication process, and can exploit these models for interpreting the image. The features to be recognized

require relatively localized processing of data, and different regions may be processed independently. The difficulty is that the model is difficult to quantify; it mixes together numerical information about minimum feature sizes, with rules about how different VLSI layers behave. Approximate, brute force methods can be devised which replace the rules with numerical tests, but these tend to be non-robust and to perform poorly on images corrupted by noise, non-uniform lighting, and non-uniform focusing. Our research is aimed at combining numerical measurements and knowledge about VLSI into a model-based system which uses highly parallel algorithms to segment the image.

The key appears to be the development of adaptive thresholding and segmentation algorithms which can flexibly modify the boundaries of the neighborhood being processed, as well as adapting the parameters of the processing to the data found in the neighborhood. Rules about the behavior of edges, intensity, texture, and so forth must guide the search over the image and the labeling of pixels into homogenous regions.

To support this effort, we have written an architectural simulator which allows 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.

Publications

Wright, R., "Highly Distributed Image Understanding Algorithms Applied to Segmentation of VLSI Images," M.S. Thesis, Department of Electrical Engineering and Computer Science, M.I.T., February 1986.

15.5 High Performance Circuit Design

Defense Advanced Research Projects Agency (Contract N00014-80-C-0622)

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

Lance A. Glasser, John Wroclawski, Charles Zukowski, Mark Matson, William Song, Terry Hohol, Charles Selvidge, Adam Malamy

Mark Matson's optimization program, which we discussed last year, has been successfully completed, tried out, and validated on circuits of about 30 transistors. The goal of this work was to write a computer-aided design tool to optimally size the transistors in an nMOS circuit to meet speed specifications and dissipate minimum power. The optimizer is able to handle multiple critical paths.

RELIC, our reliability simulator, is now functional. The objective of this research effort is to demonstrate a simulator which predicts the reliability of MOS VLSI circuits. More precisely, we would like to predict and compare the reliability of different circuits with respect to those failure

mechanisms under the control of the circuit designer. In this project we are examining three failure mechanisms: metal migration, hot electron trapping, and time dependent dielectric breakdown (TDDB). The long term failure statistics of the circuit are predicted based on the "stress" accumulated during one cycle, assuming that this cycle repeats forever.

There are three parts to the simulator: the preprocessor, the models, and the post processor. The first two parts are now operational. The models are the central part of the program. We have added the TDDB, hot electron, and metal migration models to the underlying simulator, RELAX from Berkeley.

We are starting a new research project into the high frequency behavior of VLSI circuits. About 20 years ago it was shown that fundamental frequency domain quantities such as f_{\max} could be used to predict, for instance, the properties of ring oscillators. We are extending this classical theory in three directions. First, we have reformulated it in terms of Tellegen's theorem. Second, we are extending it to much more complicated and realistic models. And third, we are making some headway on a nonlinear formulation of the problem. It is too early to predict where all of this will go, but we hope to be able to say things about what sort of performance one could get out of networks made up of certain types of components assuming one were infinitely clever. A possible application for this work is cross comparisons of technologies where one would like to compare the cleverness of the device engineers rather than the circuit designers. Another possibility would be in the area of making technology tradeoffs in the design of CMOS devices. For instance, is it "better" to have a larger lateral diffusion with more overlap capacitance but shorter transistors or is it better to minimize the parasitic capacitance?

This fundamental limits work ties in to our on-going effort to understand the tradeoffs in digital circuits between noise margins, noise, and delay. In this work we have been able to show, under some very relaxed constraints, that there is a direct tradeoff between how immune a digital circuit is to noise and how slow it must run.

Publications

- L.A. Glasser, "A UV Write-enabled PROM," Proceedings of the 1985 Chapel Hill Conference on VLSI, pp. 61-65, May 15-17, 1985. Also M.I.T. VLSI memo. 85-239.
- L.A. Glasser, "Delay, Noise Margin, and Reliability in Digital Circuits," Proceedings of the 1985 Chapel Hill Conference on VLSI, pp. 309-328, May 15-17, 1985. Also M.I.T. VLSI memo. 85-241.
- C.A. Zukowski, J.L. Wyatt, Jr., L.A. Glasser, "Bounding Techniques and Applications for VLSI Circuit Simulation," Proceedings 1985 IEEE International Symposium of Circuits and Systems, Kyoto, Japan, June 1985. Also M.I.T. VLSI memo. 85-238.
- M. Matson, "Optimization of Digital MOS VLSI Circuits," Proceedings of the 1985 Chapel Hill Conference on VLSI, pp. 109-126, May 15-17, 1985.
- M. Matson, "Macromodeling of Digital MOS VLSI Circuits," Proceedings 22nd Design Automation Conference, June 1985.
- T.-A. Chu and L.A. Glasser, "Synthesis of a Self-timed Controller for a Successive-approximation A/D Converter," M.I.T. VLSI Memo No. 85-274, October 1985.

- L.A. Glasser, "Synchronizer Failure in A/D Converters," M.I.T. VLSI Memo No. 85-276, October 1985.
- T. Fountain, "A Computer Based Aid for VLSI Array Design," S.B. Thesis, Department of Electrical Engineering and Computer Science, M.I.T., May 1985.
- F. Park, "Serial Data Communication Link," S.B. Thesis, Department of Electrical Engineering and Computer Science, M.I.T., May 1985.
- G. Wright, "New Design for a UV Write-Enabled Memory," S.B. Thesis, Department of Electrical Engineering and Computer Science, M.I.T., May 1985.
- D. Stark, "An Associative Write Buffer for a LISP Microprocessor," S.B. Thesis, Department of Electrical Engineering and Computer Science, M.I.T., May 1985.
- K. Tognoni, "Metastable Problems in A/D Converters," S.B. Thesis, Department of Electrical Engineering and Computer Science, M.I.T., September 1985.
- D.R. Stiles, "Computer Aided Design and Verification of CMOS Integrated Circuits," S.M. Thesis, Department of Electrical Engineering and Computer Science, M.I.T., May 1985.
- P. Bassett, "A High-Speed Asynchronous Communication Technique for MOS VLSI Multiprocessors," S.M. Thesis, Department of Electrical Engineering and Computer Science, M.I.T., May 1985.
- M. Matson, "Circuit-Level Optimization of Digital MOS VLSI Designs," Ph.D. Thesis, Department of Electrical Engineering and Computer Science, M.I.T., January 1985.
- C. Zukowski, "VLSI Circuit Simulation Using Bounded Error," Ph.D. Thesis, Department of Electrical Engineering and Computer Science, M.I.T., May 1985.

15.6 Waveform Bounding for Fast Timing Analysis of Digital VLSI Circuits

Defense Advanced Research Projects Agency (Contract N00014-80-C-0622)

National Science Foundation (Grant ECS83-10941)

John L. Wyatt, Jr., Paul Penfield, Jr., Lance A. Glasser, Keith Nabors, David Standley, Peter O'Brien

The goal of this research is to enable rapid computation of signal propagation delay in monolithic 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 work takes as its starting point the original paper by Penfield *et al.*¹ Their work yields computationally convenient upper and lower bounds on the step response of linear RC tree networks, which are useful as electrical models for the branching interconnect lines on integrated circuits. We have extended these results to more general classes of RC networks that can serve as more realistic models in many cases. One extension is to nonlinear models² that can more accurately represent driver resistance and load capacitance in real circuits. Another extension we have completed is a method of bounding the response of RC meshes,³ 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 transistors, 2) whenever linear resistor models are used for transistors in logic gates or CMOS pass gates, as in Christopher Terman's¹⁰ program RSIM, and 3) to model interconnect networks with closed loops sometimes created by automatic routing programs.

Furthermore, we have succeeded in reducing the region of uncertainty in the original bounds for certain classes of networks of practical interest. 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.⁴⁻⁶

We are also engaged in a major effort to extend the original results,¹ intended for MOS applications, to include appropriate models for bipolar circuits.⁷ The whole project has been unified by our discovery that the mathematical core of this approach resides in the properties of a special class of matrices, called "M-matrices."⁸

References

1. J. Rubinstein, P. Penfield, Jr., and M.A. Horowitz, "Signal Delay in RC Tree Networks," IEEE Trans. Computer-Aided Design, CAD-2, 202-211, July 1983.
2. J.L. Wyatt, 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.
3. J.L. Wyatt, Jr., "Signal Delay in RC Mesh Networks," IEEE Trans. Circuits Syst. CAS-32, 507-510, May 1985.
4. J.L. Wyatt, Jr., Q. Yu, C. Zukowski, N.-H. Tan, and P. O'Brien, "Improved Bounds on Signal Delay in MOS Interconnect," Advances in Circuits and Systems-Proceedings China 1985 International Conference on Circuits and Systems, Beijing, China, June 1985, pp. 77-81. Also VLSI Memo 85-221, January 1985.
5. C. Zukowski, J.L. Wyatt, Jr., and L.A. Glasser, "Bounding Techniques and Applications for VLSI Circuit Simulation," Proceedings 1985 International Symposium on Circuits and Systems, Kyoto, Japan, June 1985, pp. 163-166. Also VLSI Memo 85-238, March 1985.
6. Q. Yu, J.L. Wyatt, Jr., C. Zukowski, H.-N. Tan, and P. O'Brien, "Improved Bounds on Signal

¹⁰Professor, Department of Electrical Engineering and Computer Science, M.I.T.

Delay in Linear RC Models for MOS Interconnect," Proceedings 1985 International Symposium on Circuits and Systems, Kyoto, Japan, June 1985, pp. 903-906.

7. P. O'Brien, and J.L. Wyatt, Jr., "Signal Delay in Leaky RC Mesh Models for Bipolar Interconnect," VLSI Memo 85-278, November 1985.
8. J.L. Wyatt, Jr., C. Zukowski, and P. Penfield, Jr., "Step Response Bounds for Systems Described by M-Matrices, with Application to Timing Analysis of Digital MOS Circuits," Proceedings 24th IEEE Conference on Decision and Control, Ft. Lauderdale, Florida, December 1985, 1552-1557. Also VLSI Memo 85-257, September 1985 and LIDS Paper LIDS-P-1492, September 1985.

Publications

Zukowski, C.A., "Bounding Enhancements for VLSI Circuit Simulation," Ph.D. Thesis, Department of Electrical Engineering and Computer Science, M.I.T., May 1985. Also VLSI Memo No. 85-249, June 1985.