MIT Conference

Dr. Paul Penfield, Jr., MIT; Dr. William Howard, Motorola; Lynn Conway, Xerox; Dr. Paul Gray, MIT.

The M I T Conference on Advanced Research in Integrated Circuits

The Massachusetts Institute of Technology hosted a conference on Advanced Research in Integrated Circuits on January 28-30, 1980. The chairman was Prof. Paul Penfield, Jr., Department of Electrical Engineering and Computer Science. To maximize interaction during the conference, MIT chose to limit the number of registrants to approximately 250. The conference featured 28 speakers on topics ranging from process and bipolar transistor modeling to automatic compilation of integrated circuit layouts. There were no parallel sessions so everyone had a chance to attend and share in this broad range of talks. Many of the speakers, especially those concerned with processing and devices, made a special attempt to explain their work to the systems-oriented people in the audience. There seemed to be mixed reaction to this interdisciplinary approach, but on the whole the effect was quite positive.

To promote discussion and minimize proprietary information problems, none of the speakers submitted written papers. The only publication of the conference was a set of abstracts of the talks. LAMBDA has obtained permission from the authors to reprint most of these abstracts. Unfortunately there are five for which we were not able to obtain reprint clearance. The intent in printing them is to encourage understanding and communication in this field and maximize the awareness of the activities of the groups involved. If you find some of the abstracts particularly interesting, please contact the authors for more information.

There has been no specific commitment by a university to host next year's conference, but Stanford University is known to be considering it. We hope to have a formal announcement of next year's conference in the July issue of LAMBDA.

The abstracts are reprinted in the order presented.

INCREMENTAL DEPENDENCY SYSTEMS AND THE DAEDALUS GRAPHICS EDITOR

Howard E. Shrobe, Artificial Intelligence Laboratory, MIT, Room NE43-814, Cambridge, MA 02139

Daedalus is an experimental graphics editor which facilitates structured design of large scale integrated digital systems. It is implemented on the MIT LISP machine and has been used so far in two design projects.

Daedalus combines an interactive graphics capability with a procedural design language, offering greater flexibility than would be available from either representation alone. When a cell layout is worked out in the graphics editor, Daedalus automatically codes a representation of the cell in the design language as well. In this form it may be edited, extended and parameterized using the features of the LISP machine's sophisticated programming environment. Complex cell designs which have been specified procedurally may be invoked using a simple menu selection protocol within the graphics editor.

Incremental dependency systems were developed at the MIT Artificial Intelligence Laboratory as a technique for supporting evolutionary design; their use has been investigated so far in the context of analog circuit analysis and software engineering. The main research goal of the Daedalus project is to investigate the applicability of incremental dependency systems to LSI design.

An incremental dependency system contains two components: The first of these is responsible for deducing new information; in Daedalus this is implemented as a set of local propagation rules which relate the various parameters of a component (e.g., the channel lengths and widths and the pull-up/pull-down ration of an inverter). These rules can propagate information bi-directionally so that given the values of any of several subsets of parameters, the values of the remaining parameters may be deduced. The second component of an incremental dependency system is responsible for remembering how something was deduced. This second form of information is used for maintaining consistency as a design evolves. Using these techniques, Daedalus can automatically move or stretch one component (e.g., a wire) to make room for a second component (e.g., an inverter) which has had a parameter changed. Research in the future will explore the applicability of these ideas to larger design contexts.

TOWARD A THEORY OF SYSTOLIC ALGORITHMS FOR VLSI

M.J. Foster and H.T. Kung, Computer Science Department, Carnegie-Mellon University, Pittsburgh, PA 15213

The design of a special-purpose chip must begin with an algorithm design, which specifies conceptually the overall structure of the chip. The algorithm design is an integrated effort that includes the data-flow and geometry of the overall system, as well as a specification of the function of each cell. In the past several years, researchers at CMU and elsewhere have designed a large number of systolic algorithms that are suitable for direct VLSI implementation. Typically, in these algorithms several data streams move at constant speeds over fixed paths in a network, interacting at cells where they meet. Is there a theory behind these seemingly complicated algorithms? In this paper we establish a theory concerning one class of systolic algorithms—linear pipelines with bidirectional, multispeed data flow. We view this as the first step toward a general theory of algorithms for VLSI.

Let the inputs be the sequences $a_0 a_1$,...and $b_0 b_1$,.... We identify three types of computations that can be naturally mapped onto linear pipelines. These are typified by three examples: pattern matching, convolution, and searching. In these computations, terms of the outer product $a_0 b_0$, $a_1 b_0$, $a_0 b_1$,... are combined in which the difference between indices, the sum of indices, and one index, respectively, are held constant. Table 1 lists some computations of each type. Pipeline algorithms with various throughputs and latencies can be constructed automatically for each problem in the table.

In some computations, outer product terms must be combined with additional intermediate terms to produce the results. Integer multiplication, for example, is a constant index-sum computation in which each digit must be combined with carries from lower-order digits. In this paper, we present general methods for analyzing the outer product combinations required by computation and for

developing appropriate velocities and spacings for the data streams in the pipeline. By applying these methods to the pattern matching problem, we discovered that error-tolerance in matching can be achieved using "carries" analogous to those of integer multiplication.

We are in the process of extending this work to two-dimensional systolic algorithms. Our methods already provide automatic transformations of certain one-dimensional algorithms into two-dimensional algorithms of greater modularity. For example, if the cell function on two words can be represented as an "inner product" of their bit representations, a linear pipeline of words can be transformed into a two-dimensional pipeline of bits. Further work should lead to a fairly complete characterization of systolic algorithms.

Table 1. Linear pipeline algorithms

Constant Sum of Indices	
	- Convolution
	- Multiplication of polynomials
	- Serial-parallel multiplication of integers
Constant Difference of Indices	
	- FIR filtering
	- String matching with wild-card characters of error tolerance**
Constant Index	
	- Polynomial evaluation of many points (including DFT)**
	- Decomposable searching problems (nearest neighbor searching, range queries, etc.)**

**new result.

THE LIMITS TO PRECISION IN MONOLITHIC ANALOG CIRCUITS

Barrie Gilbert, Analog Devices, Inc., Semiconductor Division, 2801-B Pacific Avenue, Forest Grove, OR 97116

The precision of monolithic analog circuits continues to climb. Many commercial products, such as 12-bit D/A and A/D converters and analog multipliers with guaranteed 1/4 percent accuracy were completely unavailable five years ago. Improvements in linear-circuit technology and manufacturing have contributed significantly to this progress. Often, highly innovative design approaches have overcome barriers previously believed to be fundamental.

What determines the limits to progress in this direction? Is it fundamentally possible to construct a monolithic 20-bit converter under any conditions? What extra constraints to precision are imposed when high-speed operation is also required? Or if the cost must be severely bounded? Will the market volume for precision analog circuits decline as digital signal-processing techniques take over? Or rise as demand for their particular advantages increases? What is the analog equivalent of a ROM? What marriages of analog and digital monolithic circuit techniques promise to increase usable precision still further? Questions of this type will be raised; some answers will be offered.

ARTIFICIAL MICROSTRUCTURES RESEARCH AND ITS IMPACT ON FUTURE INTEGRATED ELECTRONICS

Henry I. Smith, Lincoln Laboratory, Massachusetts Institute of Technology, Lexington, MA 02173, and
Dept. of Electrical Engineering/Computer Science, MIT Cambridge, MA 02139

Current research on artificial microstructures will impact on future integrated electronics by providing methods for fabricating devices and circuits with much finer features and greater complexity, and by providing the means for manipulating and improving materials properties. Recent progress in improving the precision and tolerances of microlithographic and processing techniques will be described. Graphoepitaxy is an example of how artificial microstructures can be used to manipulate materials properties. Uniformly oriented films of silicon were obtained by crystallizing amorphous silicon over surface relief gratings in fused silica. The crystallographic and electrical properties of this material will be described together with recent results on devices. With further development, we believe graphoepitaxial silicon will replace silicon-on-sapphire and may lead to three-dimensionally integrated circuits.

COMPARISON OF COMPUTER PREDICTED AND MEASURED BEHAVIOR OF AN IC BIPOLAR TRANSISTOR

David H. Navon, Electrical and Computer Engineering Department
University of Massachusetts, Amherst, MA 01003

As bipolar microcircuit chip design becomes more complex there is an increasing need for a better understanding of the detailed nature of the electric carrier and potential distributions in the devices that comprise these integrated circuits. Numerical techniques can be applied to the simultaneous solution of the Poisson, and electron and hole carrier continuity equations for a precise prediction of the time-dependent carrier flow in a two-dimension bipolar transistor structure of arbitrary design. Computer-aided analysis permits the simultaneous inclusion of effects such as the lattice, impurity, temperature, field, and carrier-carrier scattering on carrier mobility; doping- and temperature- dependent carrier recombination and generation rates are also incorporated, both of the Shockley-Read-Hall variety as well as the Auger type. The rigid-band approximation of bandgap narrowing due to heavy doping in the emitter is also modeled. The Franz-Keldysh effect in the emitter space-charge region may also be included.

The results of these calculations have yielded predictions for a modeled transistor's current gain over seven orders of magnitude of current and at three different temperatures. Caution needs to be exercised in choosing an appropriate grid structure for the difference equation technique. Measurements of the current gain of transistors with an internal structure corresponding to the device simulated were made over these current temperature ranges in order to carefully assess the validity of the modeling procedures. The agreement was encouraging except at very low currents, at a low temperature.

Questions related to more accurate device modeling, particular for sub-micron device structures are addressed. Techniques for the incorporation of the results of these calculations in the prediction of the behavior of complex integrated circuits are discussed.

MACHINE COMPOSITION OF VLSI SYSTEMS

Carver Mead, Computer Science Dept., California Institute of Technology, Pasadena, CA 91125

The only known method of designing a large system is to decompose it into nearly independent subsystems. This process is applied recursively until the submodules are small enough to be implemented directly. The successful reassembly of modules is dependent upon rigid observation of composition rules. Without proper composition rules, problems are introduced into an ensemble of modules which cannot be attributed to any single module. Once a proper set of rules is adopted, machine aids can be used to eliminate most of the errors and design time associated with composing modules into a system. Examples of composition rules will be described and an overview of Dave Johannsen's latest chip compiler will be given.

A PARALLEL SEARCH TABLE FOR LOGARITHMIC ARITHMETIC

Andreas Bechtolsheim and Thomas Gross, Computer Systems Laboratory, Stanford University Palo Alto, CA 94305

A structure for searching a table of ordered values in parallel will be presented. The structure contains one comparator for each stored value, which is performing a digit-sequential comparison. Since the table is ordered, the best-matching value can be found by local inspection; the delay incurred is therefore independent of the size of the table and only proportional to the number of digits.

The structure appears to be most useful for table-lookup of functions which have a large input resolution compared to their output resolution. For this application, pairs of corresponding input and output values are stored. Finding the output value for an

arbitrary input is done by searching the set of input values in parallel for a matching interval.

A test version of a read-only parallel search table has been built in NMOS technology. In density this implementation is comparable to a PLA, since the same basic cells are used. A parallel search table can thus offer significant area advantages over approaches based on partitioned ROMs.

In the second part of the paper, the application of a parallel search table to logarithmic arithmetic will be discussed. In logarithmic arithmetic, numbers are represented by their logarithms. Addition and subtraction in this number system are performed by table-lookup. The respective functions lend themselves to implementations as parallel search tables.

Logarithmic arithmetic could be rather useful for high-speed signal processing applications. Whereas logarithmic multiplication requires just one addition, logarithmic addition and subtraction involve one subtraction, one addition, and the table-lookup. Assuming that these operations have comparable time complexity, logarithmic addition or subtraction will take three steps in time when logarithmic multiplication or division take one. With the parallel search table, a one-chip implementation of a 16-bit logarithmic arithmetic unit appears practical. With linear interpolation, precision can be increased to 32-bits, at an additional cost in time.

ASAP: ADVANCED SYMBOLIC ARTWORK PREPARATION

Kent Hardage, Hewlett-Packard Co., 1900 Garden of the Gods Road Colorado Springs, CO 80907

ASAP, an integrated circuit structured design methodology, has been developed to meet the needs for high performance custom LSI designs. It is the result of some "fresh thinking" to bring about some new concepts and to provide a new perspective from which to view LSI design.

The first of these concepts is the mapping from the "real plane" (mask level design rules) into the "symbolic plane" for IC layout. The result is symbolic abstractions that represent components and interconnect on a coarse working resolution providing the same density IC layout at the local cell description as handcrafted rectangle layout.

The second concept is that through the use of (1) symbolic regular structures, (2) a hierarchical design approach, and (3) attention to wiring management or global design, LSI designs can be achieved quickly with no area penalty and with a high success rate on first cuts at custom IC designs. The regular structures have a multilevel library description of circuit, logic, and physical.

Major topics in the ASAP CAD cycle to be discussed are on-line device modeling, multilevel libraries, design and layout verifications, and the generation of artwork from an ASAP layout. The results are in—approximately a dozen LSI designs have been completed with this methodology. Data will be presented to show that while this is not the ultimate VLSI CAD implementation, very key concepts have been demonstrated at the LSI level to establish criteria for VLSI design systems.

MOSFET CIRCUIT DESIGN FOR VLSI

Peter W. Cook, Stanley E. Schuster, James T. Parrish, Victor DiLonardo, and Darryl R. Freedman, T.J. Watson Research Center , IBM, P.O. Box 218, Yorktown Heights, NY 10598

Increasing capabilities of lithography and device processing have resulted in a rapid advance in the level of integration attainable in principle with MOSFET devices. This talk will review the basics of MOSFET scaling, and present some of the problems in scaling. Alternatives to "random" logic design approaches will be discussed, in particular the PLA Macro method, which will be compared with random logic in a microprocessor environment. The actual implementation of these PLA based parts in an advanced 1 um. MOSFET technology will be described. In that technology, individual logic gate (NOR) delays of 1.1 nS (loaded) are attainable, and PLA ALUs with byte add times of 11.5 nS have also been obtained. The circuitry used in these results will be described in detail.

MANAGING THE COMPLEXITY OF VLSI

A. Richard Newton, Dept. of Electrical Engineering and Computer Sciences, University of California, Berkeley, Berkeley, CA 94720

The design of VLSI circuits involves the management and interpretation of vast amounts of data from the initial architectural-level specification through to the final integrated circuit processing. At each stage of the design there is a potential for introducing errors or failing to detect existing errors in the circuit. To make effective use of the computer in achieving error-free, high performance circuits, it is essential that the various data representations of the design be chosen carefully and an integrated set of design programs be available to the design engineer. A number of approaches to the representation and storage of VLSI designs will be described and some examples of the use of these representations with design programs will be presented.

RECENT RESULTS WITH THE PERMEABLE BASE TRANSISTOR

C.O. Bozler, G.D. Alley, D.C. Flanders, R. A. Murphy, and W. T. Lindley, Lincoln Laboratory MIT Lexington, MA 02173

A new transistor suitable for both microwave and logic integrated circuits is being developed. This device has the potential for operation at frequencies above those achieved by any other solid state three-terminal device. The most important feature of this transistor is the thin tungsten grating of submicrometer period which has been embedded inside a single crystal of gallium arsenide. The embedded metal grating, which forms a Schottky barrier with gallium arsenide, is the base of the transistor and can be used to raise and lower a potential barrier in the semiconductor between the grating lines. The name given to this device is the permeable base transistor (PBT). Even at this very early stage of development, devices have been fabricated with a noise figure of 3.5 dB , an associated gain of 9 dB at 4 GHz, and a maximum frequency of oscillation of 17 GHz. Wafers that are currently being fabricated are expected to be better in performance. This new device has many advantages over the gallium arsenide FET besides the expected higher performance. For example there are no high field regions near exposed surfaces and a semi-insulating substrate is not required. The details of the device structure, fabrication techniques and performance will be described..

VLSI MATRIX COMPUTATION ARRAY PROCESSORS

Sun-Yuan Kung, Dept. of Electrical Engineering Systems University of Southern California Los Angeles, CA 90007

A key issue in VLSI high-concurrency designs is of developing a versatile interconnection network and timely data-movement on it. The reason is that in VLSI systems, communication consumes most of the time, area and energy, while logical operations are virtually free, due to the availability of "millions of gates" offered by modern device technology.

Aiming at the issue, H.T. Kung first proposed "systolic arrays" computing structures that efficiently exploit the high concurrency available with VLSI. There are also a number of signal processing applications described in [1]. However, all the schemes have very strict timing requirements. Therefore they can be frustrated by internal computing uncertainties. More crucially, they are all restricted to be data-independent and, therefore, are not applicable to data-dependent algorithms, e.g. "pivoting" in solution of linear equations.

To overcome these difficulties and also to gain some other flexibilities very much desired in practical applications, we propose a new data-flow type (data-dependent and self-timed) of array processors. In this paper we shall emphasize several of the basic design conceptions. First, it is strongly favored to have a natural and simple mapping between the topological properties of computation algorithms and that of computing structures. Also the tendency is to transfer complexity from the high-level software end to the localized hardware processor by adding local controls and local memories. Finally, the data-movement timing takes place in accordance with a natural "computational wavefront" mechanism, with computations

flowing according to a sort of Huyghen's principle: wavefronts are locally determined and cannot intersect. Each wavefront pipelined on the computing network corresponds precisely to one recursion step in the computation algorithm.

Based on this methodology, we have developed special-purposed array processors for—solving linear equations with pivoting—solving null-space problems—solving least-square solutions (via Givens' method—computing eigenvalue decompositions

Several other options, e.g., block matrix operations (i.e., subdivisions of algorithms) and pipelined arithmetic units (i.e., low-level pipelining) can also be accommodated in the array processor designs.

SUBMICRON PHOTOLITHOGRAPHY USING INORGANIC RESIST/POLYMER BI-LEVEL SCHEME

K.L. Tai,* R.G. Vadimsky,* and M. Oliveria**, *Bell Laboratories, Murray Hill, NJ 07974 **Dept. of Electrical Engineering and Computer Science, MIT, Cambridge, MA 02139

Using the inorganic resist (Ge-Se)/polymer bi-level scheme, we have defined submicron patterns including details as fine as 0.5 um. lines and spaces. The exposure tool is a 10:1 reduction projection printer with an f/1.4 projection lens and utilizing light of wavelength 4300A.

These high resolution features in the thin, inorganic resist layer are transferred faithfully into the underlying 2.5 um. thick polymer layer by reactive ion etching, Producing high-aspect ratio, vertical walled features. This bi-level scheme overcomes the two most difficult problems in conventional photolithography. First, the standing wave effect is eliminated by the high absorbency of the Ge-Se layer (2.5x10 cm at 400 nm). Second, the thick polymer layer covers steps and other irregularities of wafer topography so that the photosensitized Ge-Se layer is perfectly smooth.

As a result, high resolution thick polymer patterns having vertical walls can be printed on any wafer surface, including these which are stepped and/or highly reflecting. The new method offers the possibility of submicron lithographic technology based on conventional photolithography carried out using commercially available exposure machines.

HIGH RESOLUTION METHODS FOR MICROFABRICATION

Alec N. Broers, T.J. Watson Research Center, IBM P.O. Box 218 Yorktown Heights, NY 10598

Optical, electron beam, ion beam, and X-ray methods are competing to expose the resist patterns for future VLSI devices. This talk will discuss the present state of each of the methods and describe approaches being pursued to improve their performance. Possibilities for combining the different techniques in an optimum manner will also be discussed.

A CONTEMPORARY PERSPECTIVE ON INTEGRATED CIRCUIT DESIGN

Jonathan Allen, Dept. of Electrical Engineering and Computer Science, and Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, Ma 02139

The IC design process can be viewed as a succession of transformations between representations keyed to particular design foci such that the original algorithmic intent is preserved. At each of these levels (e.g. architecture, circuit, layout), well-formedness criteria can be established, and an important research goal is to determine the optimal interplay between the human designer and the computer-based design system so that a valid design meeting these criteria results, which is capable of delivering the desired level of performance in space and time. Design decisions must be seen as carrying with them this burden of well-formedness which in some cases may be so great as to leave little room for structural innovation. The nature of design tools, both synthetic and analytic, is examined from this perspective, and the current status and future prospects of design system capability is assessed.

MODELS FOR STRUCTURED INTEGRATED CIRCUIT DESIGN

I. Buchanan, Department of Computer Science University of Edinburgh United Kingdom

J. P. Gray, Department of Computer Science California Institute of Technology Pasadena, CA 91125

Traditional design tools based on geometric representations do not provide an adequate base from which to construct and verify silicon implementations of complex systems. More comprehensive structural, physical, and behavioral descriptions must be developed from the appropriate representations. This paper proposes models which may be used to construct unified and consistent descriptions of the structural, physical and behavioral attributes of a design and also discusses a method of capturing these descriptions using a textual representation embedded in an object oriented programming language.

APPLICATIONS OF SCANNING CW LASERS AND ELECTRON BEAMS IN SILICON TECHNOLOGY

J. F. Gibbons, Stanford Electronics Laboratories Stanford, California 94305

Scanning cw laser and electron beams provide an extremely convenient tool for rapidly heating the surface of a semiconductor to a precisely controlled temperature. As a result they can be used for a variety of semiconductor processing operations, including growth of surface oxides, reduction of Qss in deposited oxide films, the annealing of ion implanted layers and the improvement of the electrical properties of both metal silicides and deposited silicon films. Experiments performed to explore the device applications of these processes will be discussed in this paper.

THE SCHEME-79 CHIP

Jack Holloway, Guy Lewis Steele Jr., Gerald Jay Sussman, Artificial Intelligence Laboratory M I T Cambridge, MA 02139

Alan Bell Xerox Palo Alto Research Center 3333 Coyote Hill Road Palo Alto, CA 94304

We have designed a single-chip microcomputer (which we call SCHEME-79) which directly interprets a typed-pointer variant of SCHEME, a dialect of the language LISP. To support this interpreter, the chip implements an automatic storage allocation system for heap-allocated data. It also provides an interrupt facility for user interrupt routines implemented in SCHEME. In this paper we describe why SCHEME is particularly well suited to direct implementation of a LISP-like language in hardware, how the machine architecture is tailored to support the language, and the design methodology by which the hardware was synthesized. We develop an interpreter for SCHEME written in LISP which may be viewed as a microcode specification. We describe how this specification is converted by successive compilation passes into actual hardware structures on the chip. In the process we develop a language embedded in LISP for describing layout artwork. This allows us to procedurally define architectural elements. The architectural elements are generalized components. The generators accept parameters to produce the specialized instances used in a particular design. In conclusion, we discuss the expected performance of the current design and directions for improvement, both in the circuit performance and in the algorithms implemented by the chip.