



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

COMPUTER SYSTEMS LABORATORY

DEPARTMENTS OF ELECTRICAL ENGINEERING AND COMPUTER SCIENCE
STANFORD UNIVERSITY STANFORD, CA 94305



AD A142511

Research in VLSI Systems

Technical Progress Report

November 1983 - March 1984

Computer Science Department
Computer Systems Laboratory
Information Systems Laboratory
Integrated Circuits Laboratory

JUN 25 1984
A

DTIC FILE COPY

This work was supported by the Defense Advanced Research Projects Agency, contracts MDA903-84-K-0063, MDA903-83-C-0335, MDA903-80-C-0433, and MDA-903-80-C-0107.

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

84 06 22 057

Research in VLSI Systems

Progress Report for November 1983 - March 1984

**Department of Electrical Engineering
Department of Computer Science
Stanford University
Stanford, California 94305**

**Heuristic Programming Project and VLSI Theory Project
DARPA Contract No. N00039-83-C-0138
DARPA Order No. 4649**

**Principal Investigators: B. Buchanan, E. Feigenbaum, J. Ullman
Computer Science Department
Monitored by R. Ohlander**

**General Purpose VLSI-Based Multiprocessors
DARPA Contract No. MDA903-83-C-0335
DARPA Order No. 3773-6**

**Principal Investigator: John Hennessy
Computer Systems Laboratory
Monitored by P. Losleben**

**A Fast Turn Around Facility for Very Large Scale Integration (VLSI)
DARPA Contract Nos. MDA903-80-C-0432 and MDA903-84-K-0063**

**Principal Investigator: James Meindl
Integrated Circuits Laboratory
Monitored by P. Losleben**



✓
Letter in file

11

Table of Contents

1 Design Description, Analysis, and Synthesis	4
1.1 Regular Expression Compiler	4
1.2 Pascal-to-Silicon Compiler	4
1.3 TV - An nMOS Timing Analyzer	6
1.4 Control Compilation	8
1.5 Palladio: An Exploratory Environment for Circuit Design	9
1.6 Sticks Compaction	11
1.7 Cooperating Synchronous Systems	11
2 VLSI Processor Architecture	13
2.1 MIPS - A High-Speed Single-Chip VLSI Processor	13
2.1.1 Recent progress	13
2.1.2 The Optimizing Compiler and Benchmarks	15
3 Testing	15
3.1 A two-port JK flip-flop to simplify testing	15
3.2 Parametric Test System	16
3.3 Laser Activated Test Structures (LATS)	17
3.4 CMOS Test Stripe	18
3.5 Array Test Structures	19
3.5.1 Test Chip Composition	19
3.5.2 Chip Fabrication	20
3.5.3 Software Development	20
3.5.4 Measurement Results	23
3.6 MEDIUM Tester	26
3.7 ICTEST Testing System	27
3.8 Automatic Test Generation	27
4 Theoretical Investigations	28
4.1 Funneled Pipelining and VLSI-Oriented Algorithms	28
4.2 Performance Evaluation for Regular Expression Compiler	29
4.3 Multiprocessor Implementation Limits	29
4.4 VLSI Complexity of Functions with Special Local Properties	30
5 Fast Turn-Around Laboratory	30
5.1 Computer Automated Fabrication	30
5.1.1 Equipment installation	30
5.1.2 Language design	31
5.1.3 System design	32
5.1.4 Programming language	32
5.1.5 Fable interpreter	33
5.1.6 Graphics interface	33
5.2 Microlithography	34
5.2.1 MEBES Electron Lithography and Mask Making	34
5.2.2 Tri-Level Resists	34
5.2.3 Optical Lithography	34

5.2.4 Inspection	35
5.3 Processes, Devices, and Circuits	35
5.3.1 Fabrication of MIPS in 3.0 Micron nMOS	35
5.3.2 2.0 Micron CMOS Analog/Digital Gate Array	36
5.3.3 Laser Monitoring of Particulate Defects	36
5.3.4 Electrical End-Point Detection During Plasma Etching	37
5.3.5 Electrical vs. Physical Line-Width of Polycrystalline Silicon	37
5.3.6 Deep Trench Etching	38
5.3.7 Plasma Etching Diagnostics	39
5.4 Interconnections and Contacts	39
5.4.1 Sputtering Technology	39
5.4.2 Selective CVD of Tungsten	40
5.4.3 Multi-Layer Aluminum Alloy Interconnection	41
5.4.4 Fine Grain Polycrystalline PMOS/SOI Transistors	42
5.5 Cell Library	44
5.6 Packaging Technology	44

Abstract

This report summarizes progress in the DARPA funded VLSI Systems Research Projects from November 1983 to March 1984, inclusive. The major areas under investigation have included: analysis and synthesis design aids, applications of VLSI, special purpose chip design, VLSI computer architectures, signal processing algorithms and architectures, reliability studies, hardware specification and verification, VLSI theory, and VLSI fabrication. The major research problems are introduced and progress is discussed; the Appendix contains a list of published research papers from these projects.

Key Words and Phrases: VLSI, design automation, computer-aided design, special purpose chips, VLSI computer architecture, signal processing, routing, layout, memory reliability, VLSI theory, knowledge-based design systems, IC fabrication.

This work was supported by the Defense Advanced Research Projects Agency, contracts MDA903-84-K-0062, MDA903-83-C-0335, MDA903-80-C-0432, and MDA-903-80-C-0107.

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

Executive Summary

The major progress of note for this period is as follows:

1. *MIPS: A VLSI Processor.* MIPS (Microprocessor without Interlock between Pipe Stages) is a project to develop a high speed (> 1 MIP) single chip 32-bit microprocessor. During this period we managed to test both MOSIS 4μ parts and Stanford 3μ parts. The 4μ parts performed at the desired speed (2 MHz), while the Stanford parts had low yield and a full speed part has not yet been found. Large programs run on our special tester board confirmed that the part was fully functional at high speed.
2. *TV: An nMOS Timing Analyzer.* TV has been enhanced to include a new nMOS timing methodology. By correctly taking advantage of transparent latches in the absence of dynamic logic, performance improvements of 20-30% are possible. TV recognizes this more ambitious clocking methodology, correctly predicts performance, and guarantees that the methodology is used safely.
3. *Structured Random Logic Generation.* A system for creating one dimensional gate matrix style layouts (also called Weinberger arrays) from Boolean equations is being developed. This system includes a logic transformation system, placement by simulated annealing techniques, and automatic routing. Both nMOS and CMOS implementations are possible.
4. *Two-port JK flip-flop design* A two-port JK flip-flop that can be reconfigured into a shift register during test mode (for scan chains) has been designed.
5. *Palladio: An Exploratory Environment for Circuit Design.* Palladio is an environment for experimenting with design representations, design methodologies, and knowledge-based design aids. During the past five months a refined version of Palladio was implemented in Zetalisp on the Symbolics 3600 computer. The refined system is being used to investigate various hardware and software architectures for knowledge-based systems.
6. *Computer Support - Fable.* We have completed a prototype of a wafer fabrication description language called FABLE I [Ossher 83a] which will allow us to produce electronic run sheets which will guide a technician through the fabrication sequence or, ultimately, control an automatic fabrication facility. A key feature of this language is the separation to the high level process step specification from the equipment-specific detailed execution of these high level steps to enhance the protability of a process specification in space or time. Work is progressing on FABLE II and on the implementation.

7. *Electron Beam Lithography.* The Stanford MEBES machine has been used to routinely prepare masks for Fast Turn-Around Laboratory wafer fabrication including masks for two versions of MIPS with 3.0 μm minimum features. Incorporation of a 10 Mbit/sec Ethernet interface into the MEBES machine is nearing completion. This will *greatly* enhance our ability to accept mask information from outside sites.
8. *nMOS Wafer Fabrication.* The Fast Turn-Around Laboratory has completed fabrication of MIPS with 3.0 μm feature sizes.
9. *2 Micron CMOS.* We have developed a 2 μm mixed analog/digital CMOS gate array which includes poly- n^+ capacitors for switched capacitor filter applications. [Kuo 84] Design rules and SPICE parameters for this technology are being assembled and distributed.
10. *LPCVD Deposition of Tungsten.* Selective deposition of tungsten has been used as a contact metallurgy in both nMOS and CMOS processes. Additionally, experimental Schottky barrier source/drain PMOS transistors have been fabricated which demonstrate high transconductance and high resistance to latch-up.
11. *Fine Grain Polycrystalline PMOS Transistors.* Proton ion implantation has been used to produce silicon-on-insulator (SOI) PMOS transistors which demonstrate good ON/OFF current ratio characteristics. These devices are attractive as loads in SOI CMOS circuits.
12. *Area-time Bounds* We have extended Thompson's bisection technique for proving AT^2 bounds to multiple partitioning with certain average constraints. We have proved new (and more or less tight) AT^2 and AT bounds for error correcting codes, sorting, matrix multiplication, shifting, and restricted FFT.
13. *Test Generation for MOS.* We are devising test generation techniques for transistor switch faults in combinational nMOS circuits. Some initial experiments indicate that these techniques will achieve high fault coverage with only a modest increase in complexity when implemented as part of the D-algorithm.
14. *Tester.* We are beginning to distribute the first MEDIUM testers. We plan to supply at least one tester per contractor over the next 6 months or so.

Technical Progress

1 Design Description, Analysis, and Synthesis

1.1 Regular Expression Compiler

We are designing two enhancements to the compiler. First, we want to handle fork-join constructs, and will implement them by a preprocessor that converts pairs of expressions that must terminate at the same time into single expressions over expanded alphabets. Second, we are implementing a "friendly" front end that allows the user to talk in terms of while- and if- statements, for example, and to talk in terms of either input wires or abstract symbols, interchangeably. The result will be a high-level "microcode" language, but because our back-end heuristics for state assignment are demonstrably more powerful than the typical "look for sets of compatible states" heuristics, we expect to provide a system of high performance, with none of the problems that make our current language hard to use.

A related activity, evaluating the back-end heuristic, is reported in the Theory section.

Staff: E. Cohn, A. R. Karlin, H. W. Trickey, J. D. Ullman.

References: [Karlin 83]

1.2 Pascal-to-Silicon Compiler

Howard Trickey is continuing the design and implementation of Flamel, a translator of a Pascal subset into silicon. When completed, Flamel will read a Pascal program and produce data path and controller descriptions for a circuit with the same I/O behavior as the program. The goal is to have a compiler capable of producing designs with a variety of time/area tradeoff characteristics. What distinguishes this effort from similar projects is the extent to which the program may be transformed in an effort to find parallelism.

The control generation proceeds as follows: the program is read in and its basic blocks (straight-line code) and loops are identified. Within basic blocks, the initial microcode

schedules things as soon as data dependency requirements allow. As the data path is generated (see below), the schedule may change to reduce the resource requirement. This much has been implemented. The next step will be to rearrange the program structure, using conventional optimizing compiler techniques like code motion, loop unrolling, and expression height reduction to increase the number of things that can be done in parallel. Also, loops that can be pipelined will be identified, perhaps changing the microcode schedule to reduce the loop period. And sometimes, dataflow analysis will reveal that whole sections of program can be executed in parallel with other sections, using a fork/join control structure. How much of this is done depends on where the user wants to be in the time/area tradeoff.

The general scheme for the data path architecture is collection of functional units (adders, ALUs, registers, etc.) arranged in a bit-slice manner, with wiring tracks next to each slice used for local and global interconnection. A feature of Flamel that differentiates it from other compilers is that program variables are not fixed to registers. Rather, values may flow through the data path in a dataflow-like manner.

Flamel starts by assigning separate functional units to do every program operation and separate busses for every interconnection. Then it "folds" various resources together to reduce the cost. The cost is a rough estimate of the area that will be used; it takes into account such things as how many input and output multiplexors will be used, and how many things will be attached to busses. Folds of resources that aren't used in the same microcode cycle are preferred, but sometimes increasing the time requirement is the only way to get an acceptable area cost. This portion of Flamel has been implemented, and it seems to do a good job of assigning resources so that relatively few connections are needed.

A useful byproduct of the implementation effort has been a program called Gdraw. It converts a node-and-edgelist representation of a graph into a graphical representation, printable on a laser printer. Gdraw automatically positions nodes, routes edges, and typesets labels in a way that tries to avoid crossings. The program was developed to aid in debugging Flamel, but algorithms involved may also be of interest from the standpoint of VLSI placement and routing.

Staff: H. Trickey.

Related Efforts: MacPitts (Lincoln Labs), Shrobe (MIT), Agre (MIT), Plex (Bell Labs), CMU-DA system (Tseng and Siewiorek; Hitchcock and Thomas; Nagle, Cloutier, and Parker) (CMU), Palem, Fussel, and Welch (University of Texas, Austin), Organick (University of Utah), HINGE (University of Edinburgh), MIMOLA (University of Kiel, W. Germany), Bilgory (University of Illinois).

1.3 TV - An nMOS Timing Analyzer

TV and IA are timing analysis programs for nMOS VLSI designs. Based on the circuit obtained from existing circuit extractors, TV determines the minimum clock duty and cycle times. It calculates the direction of signal flow through all transistors before the timing analysis is performed, in contrast to combinations of designer-assisted and dynamic determination of signal flow, as in Crystal, being done at Berkeley. The timing analysis is breadth-first (block-oriented) and pattern independent, using only the values *stable, rise, fall*, as well as information about clock qualification. Its running time is linear in the number of nodes and transistors, and can analyze 4,000 transistors per minute of VAX 11/780 CPU time.

IA (TV's Interactive Advisor) allows the user to quickly experiment with ways to increase circuit performance. With the IA, the user can resize pull-ups and pull-downs or insert super buffers, and find out the effects of these changes on chip-wide performance interactively. By using information already computed by TV, it is able to propagate the effects of changes through 1,000 transistors per second of VAX 11/780 CPU time.

TV was heavily used in the MIPS project. When TV was run on the first version of MIPS, it predicted a cycle time four times longer than our original design goal. By making extensive modifications to the design we were able to reduce the cycle time to half the original prediction. Accuracies within 20% for most critica' paths compared to circuit simulation and fabricated chips have been achieved.

The major addition to TV has been the incorporation of a new high performance MOS timing methodology based on taking advantage of level-sensitive latches. Because D-type latches are used in nMOS design (instead of master-slave flip flops), signals can legally arrive (and correspondingly leave) a latch anytime the gating clock is high. Like other timing analysis systems, TV initially started and terminated all paths at rising edges of clocks; this also simplifies the analysis. Recently techniques have been developed and implemented which allow delays incurred while a clock is high to be charged to either that clock or the previous clock, so as to minimize and more accurately model the cycle time predicted for a design. We call this technique *borrowing*, because one clock phase essentially borrows time from the surrounding phases. These techniques still maintain the linear time of the previous algorithms. TV helped realize a 30% performance improvement in the MIPS cycle time.

With these tighter timing methodologies, it becomes more important to verify hold times. Two types of hold time checks are made. First, hold times are derived for all the inputs to a chip. Second, hold times are computed for all latches in the chip. If a user-specifiable safety margin is not met for latches within the chip, these latches are flagged as violations. This verification allows timing dependent two-phase clocking methodologies to be used (in contrast to strict two-phase designs which are guaranteed to work if the clocks are made slow enough), allowing for higher performance designs. These hold time checks also have running time linear in the number of nodes and transistors.

TV is currently being distributed by Stanford's Office of Technology Licensing for a nominal fee; contact Elizabeth Batson, Office of Technology Licensing.

Staff: N. Jouppi

Related Efforts: Crystal (Berkeley)

References: [Jouppi 83a, Jouppi 83b]

1.4 Control Compilation

A major focus of our design aid work is the creation of a control synthesis system. We found this portion of the MIPS design to be a major stumbling block, both in complexity and in the difficulty of meeting the desired performance without careful hand decomposition and tuning. Our goal is to automatically synthesize optimized control implementations from high level specifications that go beyond the capabilities of our earlier system, SLIM [Hennessy 81].

Our initial attack has been on two key problems: Decomposing PLAs into separate parallel PLAs and developing alternative back-ends for logic compilers.

We reported on our successful attacks on PLA partitioning in [Hennessy 83]. The algorithm used is a merge style algorithm; it is reasonably efficient and accurate at estimating PLA costs. Improvements in the range of 10-40 percent of the original area are standard. Current work is focusing on three subproblems that are important to enhancing the usefulness of PLA partitioning:

1. Placement estimation which will estimate the added routing area needed for a partitioning and modify the partitioning appropriately.
2. Delay-based decomposition that will decompose the PLAs according to timing constraints.
3. Alternative merging schemes that attempt to improve the results and running time obtained by the merging step.

Hennessy's algorithm starts by combining product terms into PLAs in a greedy fashion, then combining PLAs until no further area improvement can be achieved. There are three potential problems with this approach.

1. Greedy algorithms tend to get trapped in local minima that may not be close to optimal.
2. The iteration steps are very different in size.
3. The method does not lend itself easily to other types of objective functions (e.g., PLAs must be of more or less the same size).

We propose to overcome these problems by using an iterative algorithm such that in each iteration no more than one product term is moved from one PLA to another. This approach is similar in spirit to iterative improvement algorithms for module placement. It is also suited for applying simulated annealing. We have coded this algorithm. We expect to report some results in the next few months.

Our back-end efforts have concentrated on methods to generate structured, multi-level logic implementations. We are currently exploring *optimized* Weinberger array implementations. Current systems for generating Weinberger arrays cannot compete with PLA implementations. Our goal is to generate a Weinberger backend that will be more efficient than PLAs in some important cases. Some initial progress in exploring the use of simulated annealing to solve placement problems has been made. This work will be reported on in the conference on Simulated Annealing and Its Applications in Yorktown Heights, NY, May 1984.

Staff: C. Rowen, J. Hennessy, Y. Brandman, A. El Gamal

Related Efforts: Smile (Berkeley and IBM, Yorktown Heights), Lincoln Boolean Synthesizer (Lincoln Labs),

References: [Hennessy 83]

1.5 Palladio: An Exploratory Environment for Circuit Design

Palladio is an environment for experimenting with design representations, design methodologies, and knowledge-based design aids. It differs from other prototype design environments by providing the means for constructing, testing and incrementally modifying or augmenting design tools and design languages.

Palladio provides a testbed for investigating elements of circuit design that includes specification, refinement, simulation, and use of existing designs. It has facilities for conveniently defining models of circuit structure or behavior. These models, called perspectives, are similar to design levels; the designer can use them to interactively

create and refine circuit design specifications. Palladio provides an interactive graphics interface for displaying and editing structural perspectives of circuits in a uniform manner. A declarative logic behavioral language with an associated interactive behavioral editor is used for specifying a design from a behavioral perspective. Further, a generic, event-driven simulator can simulate and verify the behavior of a circuit specified from any behavioral perspective and can perform hierarchical and mixed-perspective simulation.

During the past five months we have completed a refined version of the the basic framework underlying Palladio. This refinement, which is implemented in Zetalisp on the Symbolics 3600 computer, addresses some of the efficiency problems of the original implementation. For example, the refined system includes the automatic translation of the behavioral specification of a circuit expressed as declarative logic assertions into Lisp procedures. The system uses the procedural form of behavior for efficiency in simulation of the circuit and the declarative form for reasoning about the circuit (e.g., for fault diagnosis or test generation).

The refined system serves as a common implementation environment for several research activities within the Heuristic Programming Project at Stanford and at the Fairchild Laboratory for Artificial Intelligence Research. For example, the Palladio system is being used by an advanced architecture project at Stanford to investigate the potential concurrency in an existing knowledge-based sonar signal interpretation program and by a group at Fairchild to investigate certain communication trade-offs in a proposed data-flow architecture for supporting knowledge-based systems.

The basic Palladio system is described in detail in [Brown 83].

Staff: H. Brown, G. Foyster, N. Singh (Stanford and Fairchild), C. Tong.

References: [Brown 83]

1.6 Sticks Compaction

Supercompaction refers to a set of 1-dimensional sticks compaction techniques. Our goal is to develop a highly predictable, high-quality 1-dimensional compactor to form the basis for further work on logic-to-layout compilation. Supercompaction is effectively a limited, efficient form of 2-dimensional compaction. A partially compacted layout is analyzed, then it is modified by introducing jogs and performing stretching in one dimension in order to improve cell pitch in the other. See our October 1983 site report for more information.

During this reporting period we have performed many more experiments and closer analysis of supercompaction techniques. We are now convinced that a combination of local search around critical-path components to find holes for them to fit into, combined with simple jog introduction heuristics, leads to computationally efficient, high-quality pitch minimization. Unlike other 1-dimensional compactors, supercompaction works better as the number of degrees of freedom increases. It finds good layouts even from machine-generated stick diagrams. We are now finishing the work on supercompaction and we will be able to return to the higher-level problem of logic-to-layout conversion.

Staff: W. Wolf, R. Mathews

References: [wolfICCAD83 83], [wolf84 84]

Related Efforts: CABBAGE(UCB), ALI(Princeton)

1.7 Cooperating Synchronous Systems

Digital systems of any significant size always have many independent clocks. In the future, a single VLSIC may have several clocks as well. In any event, as soon as two independent clocks are present in a system, metastability and hence, synchronization failure, become problems. In the same vein as our previous work on 2-phase clocking disciplines, we are seeking a notation and structuring rules for such systems to give safe but practical compositions of components whose structures are amenable to mechanical checking. We are also interested in analysis to display trade-offs, e.g., between

performance and reliability, for various structuring choices, and practical circuits to use when implementing ICs.

As an example of this line of work, consider preventing synchronization failures by using synchronizers that produce completion signals and a stopable clock. (See Chapter 7 of Mead & Conway for a discussion of this idea. See [newkirkLibrary83 83] for synchronizer and clock cells.) When is such a stretchable-clock system necessary, and how does it compare to a standard system using a fixed clock and a fast flip flop for synchronization?

Both theoretical calculations and measured data from actual parts show that a fixed clock is usually adequate. Problems develop only when you are trying to run a piece of a circuit fast with respect to its implementation technology, e.g., trying to get 5-micron nMOS to synchronize at a 10 MHz rate.

We can quantify the behavior of the two alternative systems in two ways. First, we can pick a desired reliability for the fixed-clock system and ask what amount of clock stretching occurs in the stoppable-clock system. As an example, for 4-5-micron nMOS, if the fixed clock rate is chosen such that the fixed-clock system suffers synchronization failures twice a second, the mean clock period of the stretchable-clock system increases by only approximately 10% percent. Second, we can analyze the effective clock rate versus the nominal clock rate for the stretchable-clock system. Since pushing its clock rate higher results in more clock stretching, what is the limiting clock rate? The answer for 4-5-micron nMOS is well in excess of 10-20MHz, so there is no practical barrier to a high clock rate.

We will discuss other results in greater depths as this work progresses.

Staff: D. Chapiro, R. Mathews

Related Efforts: Seitz's asynchronous systems work(CIT)

2 VLSI Processor Architecture

2.1 MIPS - A High-Speed Single-Chip VLSI Processor

MIPS (Microprocessor without Interlock between Pipe Stages) is a project to develop a high speed (> 1 MIP) single-chip 32-bit microprocessor. Like the RISC project at Berkeley, MIPS uses a simplified instruction set and is a load-store architecture.

The MIPS architecture is summarized in previous technical progress reports and is discussed in several publications.

2.1.1 Recent progress

The project history over the past year has been as follows:

- March 19, 1983 The design was submitted to MOSIS for fabrication using their 3μ and 4μ feature size nMOS runs.
- April 28, 1983 Fabrication began at the Stanford Fast Turn Around Facility on their 3μ feature size process.
- June 21, 1983 The Stanford line finished 8 wafers of 81 die each. Unfortunately, an implant problem resulted in enhancement and depletion thresholds that were about a volt too high. This made testing difficult but not impossible. The design faults were eventually isolated with these chips.
- June 24, 1983 Ten 3μ feature size parts arrived from MOSIS. They had been done with an experimental process and none of them worked.
- July 6, 1983 After a two week delay to set up the testing hardware, power was first applied to the design. Initial success was slow in coming as the threshold problems of the Stanford run were dealt with. The application of $-5V$ of substrate bias yielded the best results.
- July 16, 1983 The design had been shown to be mostly working but was acting oddly in some circumstances. A timing error explained the strange behaviour. The error caused the instruction register to be latched at twice the desired frequency.
- July 20, 1983 During the final stages of testing of the design, a minor logic problem was uncovered that prevented access to one bit of processor state.

- July 26, 1983 The 4μ feature size parts arrived from MOSIS. They corroborated the previous results of the Stanford run, and included the first known part with no fabrication defects. Complete small programs were eventually run on these parts.
- August 14, 1983 A second iteration of the design was submitted for fabrication to the Stanford facility.
- September 1983 Our high speed test board was sent out for wire wrap.
- November 8, 1983 Rerun of Stanford 3μ parts returned. The precharged design worked fully, but had severe poly resistance problems.
- February 3, 1984 High speed board was debugged and a simple program was run at full speed plus 10% on the MOSIS 4μ parts (2.2MHz clock).
- February 5, 1984 Stanford returned its third 3μ fabrication. Design worked (with very low yield) on functional tester. Fabrication-related problems restricted speed to 1/3 of predicted performance.
- February 20, 1984 MIPS executed Puzzle which has over 10 million dynamic instructions.

To circumvent the inability of the Medium Tester to do tests at speed, we have built a special purpose board to determine the maximum performance of the processor. A multibus board with 64K bytes of two-ported fast static RAM, the MIPS processor and some clock generation circuitry is inserted into a SUN workstation [BaskettBechtolsheim 82]. The M68000 loads the memory with a program which the MIPS processor then executes. The clock generation circuitry allows the M68000 to vary under program control the cycle time, the duty cycle and skew time of all the clocks.

Using this high speed test jig, we set out to test the MOSIS parts. One of the more interesting recent discoveries is that the bug in the first design, due to a race, disappears when the chip is run at over 1.5 MHz. The high speed board enabled us to run the MOSIS 4μ parts at their full predicted speed; fabrication difficulties prevented Stanford parts from running at full speed, though they fully operational at clock speeds up to 1.25MHz.

2.1.2 The Optimising Compiler and Benchmarks

We have recently completed the integration of the MIPS code generator with our UCode global optimizer. The results have been extremely rewarding in several ways. First, the optimizer performs quite well and enhances the performance on a set of kernel benchmarks by an average of almost 60%. Second, this average performance improvement exceeds the average improvement obtained on several other machines (The S-1, the DEC-10, and the M68000) by an average of about 15%. This confirms our initial design goal of designing the architecture as a good compiler target. We have also shown that a relatively small number of registers (11-12) are needed to support the active variables within a procedure.

Our continuing plans involve a series of instrumentation and measurements steps using the MIPS compilers. We have some preliminary data on this topic, our plans also include an in depth measurement of the effectiveness of a register file (a la RISC) versus a good register allocation strategy. We are also bringing up a LISP compiler (PSL) under the UCode system. This will provide numbers on MIPS performance for LISP; this compiler has just begun to produce runnable code.

Staff: J. Gill, T. Gross, J. Hennessy, N. Jouppi, S. Przybylski, C. Rowen, F. Chow, A. Agarwal.

Related Efforts: RISC (UCB), IBM 801 (IBM Yorktown), Cray-II (Cray Research).

References: [HennessyJouppi 81, HennessyJouppi 83, HennessyGross 83, GrossThomas 83, Przybylski 84a, RowenPrzy 84, Przybylski 84b, Chow 83, ChowHenn 84]

3 Testing

3.1 A two-port JK flip-flop to simplify testing

The two-port JK flip-flop can be regarded as a functional combination of a D flip-flop and a JK flip-flop controlled by their respective clock inputs. This work started by studying current flip-flop circuits used in the industry and then designing the two-port

version taking advantage of some existing designs. The aim is to design this flip-flop to be free of logic hazards and to be easily testable. Technology-specific features are also used whenever possible to achieve smaller and faster circuit implementation in four technologies: TTL, NMOS, CMOS and ECL. The design is completed and a paper describing this work is now under preparation.

[Staff] E. McCluskey, D. Liu

3.2 Parametric Test System

The Parametric test system consists of a HP desktop computer as controller with HP instruments and an automatic prober connected via an HP interface bus. The system software has been extended to include self-test capabilities and also now includes software for the CMOS test stripe. Wafer mapping has also now been included in the program.

A link to a Vax now exists from the Parametric Test System and data is also stored for processing on this computer. Software has been developed for sorting, formatting and displaying data on the Vax. The formatting is required to interface the data with the STATII package from NBS. This package is capable of generating statistics, histograms, wafer maps and correlation data.

Several other projects continue to take advantage of the useful arrangement of the system. The Stanford 40-pin probe card has been used with specifically developed software to measure a large series of contact resistance structures for a project under the supervision of Prof. R. Swanson. A multilevel interconnect project uses the system, again with software uniquely developed for this purpose, to determine such characteristics as Sheet resistance, shorts and opens, interlayer shorts and breakdown voltage. A bipolar project is currently making extensive use of the rapid matrix switching characteristics of the system to study contact resistance versus temperature. The sample is placed in a chamber and heated and cooled, several measurements are taken on each sample. Two other projects are planning to use the system as a testing tool. A GaAs test vehicle is currently being produced which will have many standard test structures on it, and another bipolar project is in the design phase.

3.3 Laser Activated Test Structures (LATS)

Recent studies have shown the feasibility of fabricating photoconductors with ultra-fast switching speeds and ultra-short turn-on times, when activated by a pulsed laser beam. There exists tremendous potential for application of these devices to function as optically switched sampling nodes for ultra-short time response characteristics of an electrical device under test.

A novel application of these devices lies in the field of functional testing. A typical Silicon photoconductor may be fabricated with an OFF resistance in the range of 10^9 ohms. The laser activated ON resistance of a non-optimised photoconductor element is estimated at approximately 10^4 ohms. Thus if a Laser Activated Test Structure (LATS) is placed between a circuit node and a sense amp., then turning on the optical switch will permit charge sharing between the node under observation and the input of the sense-amp. The Sense-amp will detect the presence or absence of charge. The technique could also be used to inject charge on a node by connecting a driver to a node via the optical switch. Turning on the switch permits charge sharing between driver and node. This permits programming of inputs. This technique can be extended to include a pre-determined series of nodes to be sensed or driven. This readily lends itself to automated programming of signal detection and injection.

This idea may be used in the following manner. After the wafer is fabricated the test for the die is designed. A Silicon layer (of suitable optical characteristics) is deposited and delineated. A contact lithography step defines the contacts of interest. This is followed immediately by a metal layer (usually the second on the die) which is to interconnect the optical switches, the critical nodes, sense-amps and drivers. The circuit is then ready for functional testing.

The main advantages of this method of testing will be:

1. The area consumed by the circuit will be less due to absence of the necessity to include monitoring devices such as LSSD.
2. The test pattern (that is the arrangement of optical switches and interconnects) may be determined at any stage after final circuit design, since circuit design is independent of the testing arrangement.

To date the investigation has proceeded along several convergent paths. It is required first to develop the technology for producing and testing these devices. An experimental set-up with a Nd:Yag laser to test the LATS devices in order to understand and characterize them is close to being completed. Processing is in progress to fabricate some LATS photoconductor switches. Several deposition systems are under scrutiny to determine suitability for this process. The most promising is a low temperature plasma system with chuck temperature control. This also aids LATS silicon switches because hydrogenation aids the particular conductive properties of interest and is readily available because of the low temperature decomposition of SiH_4 .

3.4 CMOS Test Stripe

A CMOS test stripe has been developed to enable access after processing to D.C. parametric data. The test stripe is completed and fits within the framework specified of an 8 x 40 probe pad array with 80 μm pads and 160 μm centers. The area consumed by the stripe is 3.12 mm x 1.2 mm. Alternatively, the stripe can be rearranged more conveniently as a horizontal structure to form a total die area of 6.48 mm x 0.64 mm. This is usually more convenient for interlacing with project chips.

The design was based on 2 μm design rules and a twin-tub process. The test structures cover the following important areas:

1. Sheet resistance and linewidths of each conductive layer.
2. Interlayer contact resistance for all combinations of material and dopant type.
3. Transistors with both thin and thick oxide gate regions to enable parametric and especially SPICE parameters to be extracted.
4. Capacitors to enable the normal C-V parameters to be extracted i.e., oxide thickness, threshold voltage etc.
5. Circuits include inverters, a ring oscillator and a transmission gate. The thin oxide transistor series was specially selected to enable SPICE parameter extraction.

3.5 Array Test Structures

The goals of this research are: 1) Create an integrated testing methodology for evaluating and characterizing VLSI circuits, 2) Determine, rigorously, how the VLSI fabrication process affects device and circuit performance, and 3) Statistically characterize design and/or fabrication process induced defects.

To obtain such goals requires that a method for acquiring a significant amount of experimental data be devised. This can be achieved through the use of microelectronic test chips which can be used to collect data for: 1) Extracting device and circuit parameters, 2) Characterizing fabrication process induced defects, 3) Extracting fabrication process parametric data, and 4) Optimizing layout rules. Through the use of these test chips, we can gather actual experimental data and create a database which can be used in computer aided manufacturing (CAM) of integrated circuits.

At the beginning of this period, the task of fabricating the array test chip was undertaken. Also, work began to provide support software for data acquisition once the chips had been fabricated. It was decided that the array test chip would be fabricated using the IC lab's standard NMOS process.

3.5.1 Test Chip Composition

This test chip would be used to extract device and circuit parameters, extract fabrication process parameters, and extract defect statistics. Layouts for this first generation microelectronic test chip were completed during this period. This first generation test chip has a die size of approximately 1.1 centimeters on a side. The test chip contains the following test structures and test arrays:

1. Metal over polysilicon step coverage arrays for statistical evaluation of metal continuity over polysilicon steps on gate and field oxides.
2. Metal-to- n^+ diffused region contact serpentine arrays for coarse statistical evaluation of metal-to- n^+ diffused region contact integrity.
3. Metal-to- n^+ polysilicon contact serpentine arrays for coarse statistical evaluation of metal-to- n^+ polysilicon contact integrity.
4. Individually addressable metal-to- n^+ diffused region contact ROM type

arrays for fine statistical evaluation of metal to diffused region interfacial contact resistance.

5. Interdigitated metal over polysilicon finger arrays for statistical evaluation of metal lithography over nonplanar surfaces.
6. Device and circuit parameter extraction test structures for collecting MOS device and circuit parameter data.

3.5.2 Chip Fabrication

The test chip was completely fabricated in the Stanford IC lab using the standard NMOS fabrication process. Photolithography masks were made by the MEBES mask manufacturing facility in the IC lab. These masks were 4" glass plates using chrome as the lithographically patterned layer. The 4" plates were to be used in a Canon step-and-repeat 4:1 fine lithography aligner.

The NMOS process is designed to yield a nominal transistor threshold voltage of 1 V for enhancement transistors and -3 to -4 V for depletion transistors. Transistor gate oxide thickness is targeted for 700 angstroms. The process is designed to obtain nominal sheet resistances; Metal - (.02-.04) ohm/sq., Poly - (15 - 25) ohm/sq., and Diffusion - (10 - 20) ohm/sq. The minimum metal, poly and diffusion line widths were 5 microns and were designed to evaluate the lithography and etching capability of the NMOS process.

The gate and field oxides were thermally grown and the dielectric isolation oxide was the standard low temperature deposited, phosphorus doped oxide. Contact metallurgy was standard Al/1% Si metal electron beam evaporated onto heavily doped single and polycrystalline silicon and annealed at 400 degrees centigrade to create an alloyed ohmic contact. No barrier metal was used in this contact metallurgy.

A lot of 8 device wafers and 8 test wafers were fabricated using this NMOS process.

3.5.3 Software Development

During this period, software was developed to control the equipment that would be used for data acquisition. Each test array and each set of test structures required a separate set of measurement routines for controlling the measurement equipment. A brief summary of each set of test software is given in the following paragraphs:

1. **Flnarray Test** - These subroutines are used to implement testing of the interdigitated metal finger array. The software allows the user the options of changing the measurement conditions or using the default conditions. The sequence of measurements are; 1) check the addressing circuitry, 2) check for continuity of each metal line and 3) check for a short or bridge between each line and its two neighboring lines. Each line is interrogated individually by selecting the appropriate address of that line.
2. **FETvt1 Test** - These subroutines implement testing of "single" enhancement or depletion mode transistors to determine the threshold voltage. "Single" implies that each transistor has no physical connection to any other device and all its terminals are brought out to separate I/O pads. Custom built hardware is used to stimulate the terminals to obtain a threshold voltage at a predetermined drain voltage and drain current set by the user. The software provides options for user selection of terminal current and voltage levels or a default set of values can be used.
3. **SRsWe Test** - These subroutines implement testing of specially designed test structures for measuring the sheet resistance and effective line width of a conductive layer. Several measurements are taken using multiple terminal configurations to eliminate parasitics in the measurement. The user, again, has the option to select the measurement current and voltage levels or use the default values.
4. **SRsWeFETvt Test** - These subroutines are a combination of the SRsWe Test and the FETvt1 Test. They implement the testing of a specially designed test structure which measures the sheet resistance and effective line width of a transistor gate and then measures the threshold voltage of the transistor. Essentially all the measurements that are made in the two separate tests are performed in this test. The software gives the user the options of selecting the terminal currents and voltages or using the default values.
5. **FETvt10 Test** - These subroutines implement the testing of a 10 x 10 transistor array to determine the threshold voltage of each transistor in the array. Each transistor is individually measured by selecting the appropriate address which connects its terminals to the I/O pads. The software allows the user to set the terminal currents and voltages or use the default values. The measurement sequence is the same as for the FETvt1 Test except for the addition of the addressing sequence.
6. **MscA Test** - These subroutines implement testing of the metal-over-polysilicon (on field and gate oxide) continuity arrays. An array is composed of metal lines of equal length connected as a serpentine and can be tested as one complete line or each line can be individually tested. The array is first

measured as a complete string and if it is discontinuous the software then sequences through the address of each "substring" and performs a continuity measurement. The software allows the user the options of selecting the terminal currents and voltage levels and whether the measurement is to check the address circuitry or perform an actual measurement.

7. **Mdcsa Test** - These subroutines implement testing of a metal-to-diffusion contact serpentine array. The sequence of measurements are the same as in the Msca Test and employs the same measuring strategy. The software allows the user to select terminal currents and voltages and whether the measurement is to check to address circuitry or to perform an actual measurement.
8. **Mpcsa Test** - These subroutines implement testing of the metal-to-poly contact serpentine array. The measurement strategy and the measurement sequence is the same as that of the Mdcsa Test. The software allows the user to select terminal voltages and currents and whether the measurement is to check the address circuitry or to perform an actual measurement.
9. **Mpsa** - These subroutines implement testing of a metal-over-polysilicon shorts array. This array is divided into eight sections of metal serpentine crossing poly fingers. The sequence of the measurement is to first check the metal serpentine for continuity. Next, a voltage is applied to the poly fingers and then a measurement for leakage current between the metal serpentine and the poly fingers is performed for each serpentine section. The software allows the user the option of selecting the terminal currents and voltages or using the default values.
10. **Rc Test** - These subroutines implement testing of special "single" contact resistance measurement test structures. Again, "single" implies that the test device terminals are all brought out to separate I/O pads. The software makes multiple measurements using different terminal configurations thereby eliminating parasitics when calculating the parameter value. The software provides the options of user selection of terminal currents and voltages or supplies default values.
11. **Rom1 Test** - These subroutines implement testing of an 8 x 128 metal-to-diffusion contact array. This test performs a true 4-point Kelvin measurement of each contact in the array. The software allows the user to first select the option to check the functionality of the addressing circuitry. Next, the software sequences through the addresses of each contact cell and performs a 4-point Kelvin measurement on that cell by connecting the terminals to the appropriate I/O pads. Multiple measurements using different terminal configurations are performed to eliminate the parasitics of the measurements. Again, the software allows the option for the user to select the terminal currents and voltages or to use the default values.

12. CAPTEST1 - These subroutines implement testing of large area field and gate oxide capacitors. The software performs a high frequency capacitance-voltage measurement on each test capacitor. The software allows the user to select the range of the voltage sweep and a guard ring voltage if desired. An initial parasitics measurement is performed before each measurement and is subtracted out of the measured value.

3.5.4 Measurement Results

10 x 10 Transistor Array

There are 2 array per die (1 enhancement and 1 depletion) and 18 arrays per wafer. There were a total of 1800 transistors on the wafer, but only 900 enhancement devices were measured. The threshold voltage was measured at 500 mV drain voltage and 5 microamp drain current. The designed threshold voltage was 1 V for enhancement transistors and the enhancement transistor arrays yielded an average threshold voltage of .95 V for 900 devices.

Metal-over-Poly on field oxide and gate oxide.

There are 4 arrays per die and a total of 36 arrays on each wafer. These arrays measure the continuity of metal lines crossing polysilicon steps. Two arrays are for metal-over-poly on field oxide and the other two are for metal-over-poly on gate oxide. For each set of arrays, one array has 10 micron wide metal lines and the other array has 5 micron metal lines. 36 arrays were measured. Preliminary results indicate that all the addressing circuitry for the arrays functioned properly and measurements yielded that all the metal lines were continuous. The total metal line length for each array was 18,000 microns and crossed 3720 poly steps. We conclude that metal deposition and etching of this process are adequate for metal features of 5 microns.

Metal-to-Diffusion Contact Serpentine Array

There are two of these arrays per die and a total of 18 per wafer. In each die, one array has 5 x 5 micron contact windows and the other array has 2.5 x 2.5 micron contact windows. An array consists of strings of contacts of equal length connected in a pattern

that forms a serpentine. The entire serpentine is 18,000 microns long and contains 1860 contacts. Preliminary results indicate that the addressing circuitry on all 18 arrays functioned properly. Measurements yielded that all the contact strings were continuous and none of them showed any high resistance. We conclude that the contact etching and metallurgy is adequate for contact windows as small as 2.5 x 2.5 microns.

Metal-to-Polysilicon Contact Serpentine Arrays

These arrays are identical to the metal-to-diffusion contact serpentine arrays except that the contacts are metal-to-poly. There are 18 arrays per wafer and each die has one array with 5 x 5 micron contact windows and another array with 2.5 x 2.5 micron contact windows. Preliminary results indicate that all the addressing circuitry functioned satisfactorily. Measurements yielded that all the contact strings were continuous and no high resistance strings were noted.

Interdigitated Metal Finger Array

This array contains 3 subarrays that have metal line widths and spaces of 10 and 7.5 microns, 12.5 and 5 microns, and 13.75 and 3.75 microns, respectively. All the metal lines perpendicularly cross poly lines which create an uneven topography. There are 9 arrays per wafer. Preliminary results indicate that all the addressing circuitry functioned satisfactorily on each array which has 128 addressable lines. Measurements yielded that all the metal lines were continuous and only approximately 10 lines out of 1252 were shorted. The 10 shorted lines were intentionally introduced in one of the arrays to determine operational integrity. We concluded that our metal lithography is capable of patterning a 3.75 micron metal spacing. The metal lithography process was carefully monitored to determine if this could be achieved.

Individual Transistor Threshold Voltage Extraction

Individual test structures were measured to extract threshold voltage for this process. There are 6 sets of individual transistor test structures, 3 sets of enhancement devices and 3 sets of depletion devices. In each of the 3 sets, the W/L's are: 300/20, 150/10 and

75/5 microns. A special hardware adapter was used to extract the threshold voltage of the devices at 500 mV drain voltage and 5 microamp drain current. Measurements yielded an average of .94 V for 45 devices with a standard deviation of 40 mV.

MOSFET Cross-bridge Sheet Resistor Combination

This is a unique test structure that combines the measurement of sheet resistance and effective line width with the measurement of threshold voltage. There are 8 structures per die, 72 per wafer. The transistor W/L is 75/5 microns. The same measurement procedure for measuring the individual transistor structures mentioned above was used here. Threshold voltage was measured at 500 mV drain voltage and 5 microamp drain current. Measurements yielded that the average sheet resistance for the gate material was 21 ohm/sq., the electrically measured line width that defined the gate was 4.5 microns for 5 micron drawn line width. The resulting transistor threshold voltage was .89 V.

Metal-to-Diffusion and Metal-to-Poly Contact Resistors

There are 4 individual contact resistor structures per die, a total of 36 per wafer. Measurements yielded an average contact resistance of 2 ohm for a 4 x 4 micron contact window for the metal-to-diffusion contacts and 8.6 ohms for a 4 x 4 micron contact window for the metal-to-poly contacts.

Metal, Diffusion and Poly Cross-bridge Sheet Resistors

There are 4 structures per die, 36 per wafer. The metal structures yielded an average sheet resistance of .03 ohm/sq. and an average line width of 5.5 microns for a 5 micron drawn line width. The diffusion structures yielded an average sheet resistance of 22 ohm/sq. and an average line width of 6.5 microns for a 5 micron drawn line width. The poly structures yielded an average sheet resistance of 19.7 ohm/sq. and an average line width of 4.6 microns for a 5 micron drawn line width.

From the preliminary data collected from the first test chip run, we concluded that to

obtain more defect statistics it would be necessary to scale down the dimensions of the entire chip. Therefore, plans include a second test chip with minimum feature sizes of 1 micron and new test arrays are to be added. A second fabrication run will be done to collect more data and data collection and analysis of this first run will continue.

3.6 MEDIUM Tester

The MEDIUM Tester is a simple, functional tester used for the great majority of functional testing done at Stanford. It is capable of driving and sensing pins serially at a rate of 100k pins/second, limited by the DMA speed of its host LSI-11 or VAX. As part of the ICTEST system it has been used to test over 100 designs, including MIPS. It is sufficiently fast that there have been no difficulties testing dynamic designs with normal storage times.

We are beginning to distribute MEDIUM testers now. The first five will be available in kit form at the Utah meeting (we hope); others will follow over the next six months or so. We are in the process of estimating the cost for the PC board work, for chip testing, and also for a box, a power supply, and the parallel interfaces. Current plans, still being firmed up with DARPA, involve distributing one (or maybe two) kits (or maybe complete testers) to each contractor. ISI will provide the SIEVE software for the testers.

The speed and temperature mysteries mentioned in the previous report turned out to be ratio or drive problems, and we have repaired them. We have characterized the DUT drivers on the pin-electronics chips, and incorporated the clock onto the tester-controller chip. The chip set is now in trickle production on successive MOSIS runs.

Staff: I. Watson, D. Chevert, C. Kendrick, R. Mathews, D. Chapiro

Related Efforts: SIEVE (ISI)

3.7 ICTEST Testing System

The ICTEST system is a unified system for functional testing and simulation of ICs. Tests are written in ICTEST, a superset of C extended to include testing primitives, and compiled to run against one of several simulators (ESIM, TSIM, or RSIM) or testers (MEDIUM tester, TEK S-3260).

Prompted by MIPS testing and other recent testing experience, we are beginning a round of back-end development for the ICTEST system. We are cleaning up the tester backends, preparatory to improving the S-3260 interface. Current plans call for developing a high-speed link to speed loading of test vectors and unloading of results from the S-3260; ironically, the S-3260, capable of broadside delivery of test vectors at 10MHz, is our slowest tester for typical tests! The ultimate objective of this work is to be able to probe and test large chips, e.g., MIPS, at speed.

We are also incorporating a new simulator back-end for MOSSIM. Because MOSSIM is written in Mainsail, this project will require dividing ICTEST into two cooperating processes. Our objective here is to provide a simulator back-end with more precise handling of pass-transistor problems.

Staff: I. Watson, S. Taylor, A. Salz, J. Newkirk

References: watsonCHERRY82

3.8 Automatic Test Generation

We have begun work on test generation for transistor switch faults in MOS combinational circuits. The classical fault model for gate logic assumes that faults express themselves as nodes stuck at 0 or 1. However, for MOS, switch faults, i.e., stuck-open and stuck-short, are both more natural models and are moreover unavoidable: if a gate of a transistor is stuck at 0, the transistor is effectively stuck open. Switch faults are particularly difficult for ATG because faulty combinational circuits in general become neither combinational nor digital. We are seeking computationally tractible test generation for such faults.

The basic solution techniques we are exploring involve a 2-phase test. For example, a stuck-open fault in a multiplexor leads to charge storage. The idea is to drive the affected node to a known value through a good path on the first step, then attempt to drive the node to the complementary value through the stuck-open transistor on the second step. If the transistor is indeed stuck open, the node value will not change. However, the test generation procedure must take care to validate the test in light of MOS circuit properties, in this case, potential charge sharing affecting the original known value.

The D-algorithm can be extended to generate valid tests for nMOS combinational circuits. The extensions entail only a modest increase in complexity, which is necessary due to the more complex nature of switch faults. We have applied the technique to some example circuits. For a 10-input, nMOS, ALU bit slice, 100% fault coverage of 56 switch faults is achieved with 12 test vectors.

Staff: H. Chen, J. Newkirk, R. Mathews

References: [chenATPG84 84]

4 Theoretical Investigations

4.1 Funneled Pipelining and VLSI-Oriented Algorithms

The paper on this subject, reported last time, has been published.

Staff: A. Siegel.

Related Efforts: Lipton-Valdes (Princeton).

References: [Hochschild 83]

4.2 Performance Evaluation for Regular Expression Compiler

We have obtained some theoretical results of the form that the new heuristic used for state coding in the regular expression compiler is guaranteed to perform within some small number of bits of optimal. For example, if the nondeterministic automaton that we are coding happens to be two deterministic automata running in parallel, then our method is guaranteed to produce a state code no more than two bits wider than the optimal state code.

Staff: A. R. Karlin, H. W. Trickey.

4.3 Multiprocessor Implementation Limits

A recent paper discusses the problems that one will face implementing multiprocessor systems on chips or wafers. First, it is argued that the ability to sort quickly will be essential for many important activities we might expect those systems to perform. However, any plane circuit that sorts quickly, whether a chip or wafer, must use area that grows proportionally to the square of the number of processors in the circuit. This result was known since the original VLSI lower bound work of Clark Thompson.

We propose that the problem can be mitigated somewhat by using nets of more than two processors to pass information. However, if m is the number of processors whose communication needs can be served by a single shared wire, then we can show that m must grow at least as the square root of the number of processors we wish to interconnect, if the circuit is to use area linear in the number of processors and yet be able to sort as fast as possible.

Staff: J. D. Ullman.

References: [Ullman 84]

4.4 VLSI Complexity of Functions with Special Local Properties

We show that the area A and computation time T for any circuit that computes an (N, M, l) -local function must satisfy $AT^2 = \Omega(\max[N, M] \text{ times } l)$. Several functions with this local property are investigated. For the t -Barrel Shifting function on n -sequences, a lower bound of $AT^2 = \Omega(tn)$ and a matching upper bound are proved. For sorting n integers each no less than $(1+\epsilon)\log n$ bits long, the lower bound proved is $AT^2 = \Omega(n^2 \log n)$. This result can also be generalized to only requiring the first t output integers. In this case, the lower bound is $AT^2 = \Omega(t \log t)$ for input integers more than $(1+\epsilon)\log t$ bits long. In the general case when the integers are only $(1+\epsilon) \log m$ bits long, the lower bound proved is $AT^2 = \Omega(mn2^m)$ and there is a circuit achieving $AT^2 = O(mn2^m \log^4 n)$. For the multiplication of an $m \times n$ binary matrix with an $n \times p$ binary matrix, we proved a lower bound of $AT^2 = \Omega(mnpq)$, where $q = \min(m, n, p)$. Upper bounds that differ by only a logarithmic factor from the lower bounds can also be obtained.

Staff: A. El Gamal, K. Pang

5 Fast Turn-Around Laboratory

5.1 Computer Automated Fabrication

Our efforts in this period have been much more administrative and bureaucratic than technical, though we have made good technical progress. Our biggest single stumbling block continues to be the difficulty of attracting qualified graduate students to the project; very few students who have the requisite computer science background are at all interested in IC fabrication.

5.1.1 Equipment Installation

After a 3-year delay the re-equipment money allocated to the computer automated fabrication effort has arrived, and we have purchased and installed a second VAX-11/750, *Cascade*. We have brought up 4.2BSD Unix on *Cascade* to avoid the overwhelming difficulty of a later conversion; this makes it relatively unavailable to the CAF staff, though it is being used by others for overnight computation.

We have installed a 10Mbit Ethernet in our laboratory, and have made electrical connections of Glacier, Mebes, Cascade, and the parametric tester to this net. Cascade is serving as a 10-3 gateway for the moment, providing IP connections between that net and the campus spine.

We have nearly completed the software necessary to transfer files between Mebes and our VAXes over the 10Mbit Ethernet. The Mebes is such a hostile programming environment and its availability for experimentation is so limited, and the documentation on every aspect of the Mebes computer system is so inadequate, that progress has been much slower than expected. Our system involves a TFTP implementation in FORTRAN running on Mebes, using absolute 10Mbit ether addresses to communicate with Glacier or Cascade, which run background TFTP server processes. This style of communication must be initiated from the Mebes end, which is an acceptable limitation.

We are postponing the purchase of more workstations until we have finished the system design (see 5.1.3, below).

5.1.2 Language design

We have made no changes to the FABLE I language definition yet. This delay has been caused by our efforts to get an evaluation of FABLE I by the computer science research community before embarking on FABLE II. The FABLE language is sufficiently beyond the current state of the art in structure representation and approach to parallelism that we have been very uneasy about proceeding without this evaluation. H. Ossher delivered a paper on FABLE I at the 1983 SIGPLAN conference, and a paper on the structuring mechanisms at the 1984 POPL conference. In general it was well received, although the complete lack of interest in IC technology on the part of those computer science researchers qualified to evaluate FABLE has been a major difficulty. We are now ready to proceed with FABLE II, though we do not yet have as much feedback as we would have liked.

5.1.3 System design

We have completed the block design for the FABLE interpreter and runtime system, and for the graphical interface to be used in clean rooms and offices. Our original plans had been to base the computation almost entirely on SUN workstations, but they have increased in price by a factor of 3 and have been shown to be relatively unreliable; this coupled with the poor-to-nonexistent support provided by Sun Microsystems has forced us to reconsider those plans. We now have a much more conservative design based on a VAX/750 and unspecified graphics workstations; we are actively investigating the suitability of Apple MacIntosh machines for that purpose.

We are basing our network communication on David Cheriton's V system, developed at Stanford under a separate ARPA/IPTO contract. This dependence on another's research for our well being has not been painless, but it is important to have the Stanford network environment be as homogeneous as possible, so we are continuing to use V.

5.1.4 Programming language

Computer science groups nationwide who are engaged in applications programming have been wrestling with the programming language problem; very few of the languages of acceptable quality have reliable commercial implementations, and we cannot afford to do our own implementation. After months of investigation, including attempts to use CLU, C, and Pascal, we have settled on the use of Modula-2. The new DEC Western Research Laboratory has produced a good Modula-2 compiler for the VAX and made it available to us. We have completed the installation of Modula-2 on both of our VAXes (Glacier and Cascade) and have produced perhaps 50% of the necessary interface specifications so that it can be used in our network environment. DECWRL Modula-2 is calling-sequence compatible with Berkeley C, so we will always have that fallback position. We do not yet have a working 68000 Modula-2 compiler, though several substandard compilers are available commercially.

5.1.5 Fable Interpreter

We have settled on the use of IDL as a vehicle of communication between the FABLE parser and the FABLE interpreter. IDL was used very successfully in the Diana intermediate form for Ada, and we do not expect great difficulties. We are delaying the production of a full parser until FABLE II is ready, so IDL representations (constructed either by parsing FABLE I code or by manual editing) provide the requisite layer of insulation from those changes. Block design for the interpreter is relatively complete; we expect to begin coding as soon as the Modula interfaces are complete (see 5.1.4).

5.1.6 Graphics Interface

Progress on the graphics interface is least satisfying to us. Although we are basing our system design on the V kernel, the graphics support that comes with V is wholly unsatisfactory for our purposes (too slow and too restrictive). We have explored many different alternatives, none of which has been completely fruitful. The graphics support from Sun Microsystems is even slower and more restrictive than V graphics. The various GKS kernel packages available are unsuitable for our text-oriented raster application. Last year P. Asente developed a raster-based high-speed window package compatible with V, but V has changed out from under it and will probably continue to do so.

Our needs to represent procedural information in graphics frames, to be able to store frames in a database so as to permit symbolic inter-referencing, and the all-pervasive need for fast response time, have led us to a commercial package called PostScript, offered by Adobe Systems. Unfortunately, Adobe is not yet ready to license PostScript to us in a way that is acceptable to the Stanford legal staff. We expect a resolution to this problem by late summer, and cannot make real progress until either it is licensed for our use or some competing commercial package becomes available. As a desperation measure we could produce our own implementation of the Xerox JaM language, which would be relatively satisfactory; that would require about a man-year of work which we would prefer not to spend.

Staff: B. Reid, H. Ossher, P. Asente, A. Bleiberg, L. Adams, M. Blatt, R. Perelman.

Related Efforts: Hodges and Katz (Berkeley), Gershwin (MIT).

References: [Ossher 83a, Ossher 83b, Ossher 84]

5.2 Microlithography

Our lithographic efforts in this period have been well under control for the service work that has been required.

5.2.1 MEBES Electron Lithography and Mask Making

Mebes has been routinely used for producing the 1X reticles for the Ultratech 900 stepper for our 3 μm and 2 μm NMOS and CMOS processes. Improvements in our E-beam resist processing are required to generate crisp 1 μm features for the next generation 1.25 μm CMOS technology. We are currently experimenting with AZ resists, which are exposed using 3 passes with the LaB_6 source in Mebes for reticles. These positive photoresists have much wider process tolerances than E-beam resists such as PBS and are capable of higher resolution. In addition, much more processing experience has been obtained on the characteristics of these resists and their interactions with the rest of processing, such as plasma etching.

The Mebes has been well characterized for specific tasks to be undertaken under the one-eighth micron development contract with Perkin Elmer.

5.2.2 Tri-Level Resists

Work is continuing on the etching procedures for the bottom organic planarization layer of the tri-level resist structure. By using a teflon cover plate on the active electrode of the MRC RIE etch system, the grass which was caused by back sputtering of metal from the electrode has been eliminated. With this plate, crisp lines and spaces with a 0.5 μm pitch in 0.4 μm thick resist were achieved.

5.2.3 Optical Lithography

The Ultratech 900 stepper has been well maintained and shown consistently satisfactory performances. During this period, it was used to complete three NMOS runs of the MIPS and a custom analog-digital NMOS run for the Stanford Linear Accelerator

Center. There has been continuing efforts on software development associated with the interface between the Ultratech and the Glacier. At present, we are working to transfer our Ultratech/VAX communications software package to Berkeley.

Multi-level resist structures to be used on the Ultratech are being considered. A tri-level structure would also use an organic planarization layer for the bottom level and spin-on glass as the intermediate layer. Because the Ultratech stepper uses dual wavelength exposure, standing waves in the resist are reduced; however, these resist structures may improve step coverage and critical dimension control.

5.2.4 Inspection

We have been investigating several real-time digital signal processing approaches for mask/wafer inspection. Our current project is to use digital transversal filters and finite-state machines to find "design rule violations" which may be caused by defects which change feature edge locations. The input to these filters will be the backscatter signal from a full wafer non-destructive SEM. Using the process design rules as a basis for inspection is a much more tractable problem than that of comparison of features derived from backscatter signals with CAD data.

Staff: R. F. W. Pease, D. Dameron, C-C. Fu, E. Crabbe.

5.3 Processes, Devices, and Circuits

5.3.1 Fabrication of MIPS in 3.0 Micron nMOS

During the present report period we have completed two runs of MIPS using 3.0 μm nMOS technology which results in a die size of 5.4 mm by 5.6 mm. The design has been shown to be completely functional as evidenced by the successful running of the Puzzle benchmark program. As detailed elsewhere in this report, the processor runs significantly slower than simulation would predict. From a fabrication standpoint, this is largely due to the fact that these devices displayed an abnormally high body-effect which significantly reduces the speed and noise margin of pass-gate logic. The causes of this high body effect are under investigation. For higher speed operation, it may be desirable to re-target the threshold shifting implants to be suitable for use with $V_{\text{Sub}} = -2.5 \text{ V}$ to

reduce both the body-effect of the transistors and to reduce the junction capacitance of the drain/source regions.

5.3.2 2.0 Micron CMOS Analog/Digital Gate Array

Our 2 μm process has been modified to include provision for high-quality MOS capacitors as would be required for switched capacitor filter applications. An additional n^+ implant is used early in the process sequence to produce the lower electrode of a MOS capacitor with a low voltage coefficient of capacitance. Electrical characterization of a switched capacitor filter using this process is in progress. [Kuo 84] Because the n^+ and p^+ source/drain regions in this process are quite shallow (0.3 μm and 0.55 μm , respectively) we have incorporated a selective deposition of tungsten in the contact regions to prevent junction leakage problems with the sputtered aluminum allow interconnections. A set of comprehensive design rules based on this process is being assembled and distributed to interested members of this community.

5.3.3 Laser Monitoring of Particulate Defects

We have arranged for the delivery of a Tencor Surf-Scan particle measurement system which will enable us to monitor the particle densities on bare wafers during each step of the wafer fabrication process. This system will be useful in the monitoring of particulate generating processes such as LPCVD to determine when the system is in need of repair and or cleaning. Prior to the advent of such systems, the determination of cleaning frequency was largely a subjective matter. More importantly, this piece of equipment allows us to dissect our process on a step-by-step basis to quantitatively determine which process steps and/or pieces of equipment are generating significant densities of particles.

While evaluating this particular piece of equipment, for example, we found that one of the most significant particle sources in our laboratory is the ion implanter -- a piece of equipment not normally considered to be a particle generator. Further investigation indicated that a poorly designed vacuum back-fill system was responsible for the generation of high densities of particles in this system. The availability of this wafer monitoring tool will greatly enhance our ability to control the defect density associated with particulate contamination in our laboratory.

5.3.4 Electrical End-Point Detection During Plasma Etching

Development of plasma etching techniques suitable for etching contact holes in the p-glass of our nMOS and CMOS processes continues. Because the selectivity of these processes is often no greater than 4 (i.e. relative etch rate of SiO_2 to that of either Si or photoresist is roughly 4) it is imperative to have an accurate means of determining when the etch is complete. The problem of overetching is particularly acute in the case of the 2 μm CMOS process in which we must etch contacts through 6000 Angstrom of SiO_2 without etching through a 3500 Angstrom deep drain/source junction.

Although optical spectroscopy of the emission emanating from the plasma has been used as a means of end point detection, the small area typically encountered during contact hole etching tends to reduce the utility of this technique because of a low signal-to-noise ratio. As an alternative, we have been monitoring the current which is induced by the plasma and flows through the sample itself. The magnitude of this current changes significantly when etching of the oxide (and, hence, exposure of the bare silicon to the plasma) is complete. By using a suitably anodized aluminum electrode, we have been able to use this signal for endpoint detection. Initial experiments indicate that the induced current between the electrodes doubles in magnitude as the contact holes open at the end of the etch. Refinement of this technique and an investigation of the effect of this current on device junctions is in progress.

5.3.5 Electrical vs. Physical Line-Width of Polycrystalline Silicon

Martin Buehler of JPL had previously reported differences between the electrical and physical line-widths of polycrystalline silicon in fully processed structures. [Buehler 83] Discussions since that time indicate that the source of this discrepancy may be due to enhanced oxidation along polycrystalline grain boundaries. Joint experiments with JPL and the DARPA process modeling program are being initiated to confirm this hypothesis.

5.3.6 Deep Trench Etching

Work has continued on developing deep trench isolation for the elimination of latch-up in CMOS. During this period this effort has concentrated on characterizing the electrical properties of polysilicon filled trenches. In particular, we have been measuring the interface properties of the trench wall/oxide structure to assess the probability of forming unwanted parasitic conduction channels. A principal problem with reported trench isolated CMOS devices has been parasitic channels in the n-channel devices. These channels are believed to be caused by high values of fixed interface charges, Q_f , associated with the growth oxide on the walls of the trench. To investigate these channels, a mask set was designed to measure Q_f on the bottom and sides of our trenches. Using a 700 Angstrom gate oxide grown on the walls of the trench and a 3000 Angstrom thick doped poly Si gate deposited on the walls, initial results indicate a low 10^{10} per cm^2 fixed charge density. This Q_f value is significantly below reported values and is probably due to lower ion energy associated our use of "plasma" mode etching as opposed to the reported trench etching using the higher ion energy RIE mode etching.

Combined F/Cl anisotropic etching of silicon - Advanced isolation schemes for VLSI processes require controlled anisotropic etching of Si. Fluorine based plasma etching offers high etch rates, high selectivity and safer non-toxic gas handling. However, F based etch processes such as with CF_4 or SF_6 tend to etch isotropically except at low pressures and high ion energies, where rates and selectivity are low. Whereas, chlorine based etching tends to be more anisotropic but suffers from low selectivity. We have been investigating combined F and Cl plasma etching using a mixture of C_2ClF_5 and SF_6 . It was found that near anisotropic etching can be obtained at relatively high rates and high selectivity. The key to obtaining this anisotropic etching appears to be a reaction between activated Cl species and the photoresist masking material to form an inhibiting layer on the sidewalls. If an gaseous substitute for the role of the resist can be found, then it should be possible to switch back and forth between isotropy and anisotropy etching and thus be able to program the shape of a side wall.

5.3.7 Plasma Etching Diagnostics

Effects of plasma etching on contact resistance - Plasma etching of contact holes offer small geometries and tighter size control over wet etching. However, to get good selectivity over Si the etching processes are set up the deposit polymers on exposed Si while etching Si. The removal of these polymers is critical to obtaining low contact resistances. After examining a number of both dry and wet methods of removing these polymers, it was found that a short Si plasma etch using C_2ClF_5/SF_6 gave the best contacts. Auger measurements of Si surfaces after using this etch, gave low C concentrations agreeing the contact results.

Correlation of emission and etching non-uniformities - A spectrometer system is being set up to locally measure optical emission from different points in the plasma volume with the aim of correlating concentration variations with etching characteristics.

Staff: J. D. Shott, J. P. McVittie, J. R. Pfister, K. C. Saraswat, S. H. Goodwin, L. Lewyn, J. D. Plummer, B. Bakoglu.

Related Efforts: Oldham (Berkeley).

References: [Pfister 84, Lewyn 84, Kuo 84]

5.4 Interconnections and Contacts

5.4.1 Sputtering Technology

During this period we have installed and characterized a Balzers sputtering system which allows us to sequentially sputter or co-sputter a variety of target materials including aluminum, copper and silicon alloys of aluminum, tungsten, titanium, and silicon. This piece of equipment is central to much of our multi-level interconnections and contacts research activity. Initially this system has been used to replace E-beam evaporation for metalization of our nMOS and CMOS devices. The greatly improved step coverage provided by sputtering relative to that of E-beam evaporation greatly reduces the need for glass reflow and/or tapered contact holes elsewhere in the process. Additionally, the uniform grain size provided by sputter deposition results in much greater line width

uniformity during wet chemical etching of these films. This, in turn, eliminates the creation of "hot spots" due to current crowding in the metal and makes it possible to continue to use wet chemical etching at smaller feature sizes than are achievable with E-beam evaporated films.

5.4.2 Selective CVD of Tungsten

Low pressure chemical vapor deposition of W in a hot wall reactor and its applications to the VLSI technology have been investigated. W has been deposited selectively on Si, Al, WSi_2 and PtSi. Selectively deposited W makes reliable and low resistance ohmic contacts to n^+ and p^+ shallow diffusions. This process appears to remove the residual native oxide on a Si surface and thus provides atomically clean interfaces for low resistance contacts. W acts as a diffusion barrier between Al and Si. Contact resistance of W to n^+ and p^+ shallow junctions with doping density between 10^{19} to $2 \times 10^{20} \text{ cm}^{-3}$ has been characterized. At a doping density of 2×10^{20} to specific contact resistivity to n^+ Si was about $2 \times 10^{-8} \text{ ohm cm}^2$ and to p^+ Si was about $2 \times 10^{-7} \text{ ohm cm}^2$. This technology was used to fabricate the MIPS.

Schottky contacts to N type Si show excellent I-V characteristics with a barrier height of 0.62 eV. Schottky contact PMOS have been fabricated and incorporated in CMOS structures for latch-up immunity. WSi_2 was formed by ion implanting As in W deposited on Si and should prove useful for silicidation of source-drain junctions. Thick layers of W show excellent step coverage due to the fundamental nature of the LPCVD process and should be very useful for interconnection applications in a multilayer technology. Selective deposition technology shows promise for planarization of vias by refilling them.

Selective low-pressure chemical vapor deposition of tungsten (W) has been incorporated into routine fabrication of $3 \mu\text{m}$ nMOS and $2 \mu\text{m}$ CMOS technologies as a contact metallurgy for shallow junctions. The tungsten film provides both a low resistance ohmic contact to the source/drain regions and a diffusion barrier between Si and Al to eliminate junction spiking during annealing operations.

Selective deposition technology is also being investigated as a means of planarizing the vias in a multi-layer metal system. Vias with vertical sidewalls will be etched and then refilled by selectively depositing W in the contact regions resulting in a planar surface. If the thickness of the selective deposition can be increased sufficiently the step coverage during subsequent Al alloy depositions can be greatly improved.

Finally, selective deposition technology has been used to form Schottky barrier drain/source regions of PMOS transistors. By adding a lightly doped boron region (similar to the "tip implants" common in the fabrication of small geometry nMOS transistors) to this Schottky drain structure we have been able to maintain the latch-resistance of the Schottky PMOS device while preserving the high transconductance of the conventional implanted PMOS transistor.

5.4.3 Multi-Layer Aluminum Alloy Interconnection

A detailed study of aluminum/titanium alloys has been made in hopes of finding a replacement for the more conventionally used aluminum/copper alloys. Sputtered aluminum copper alloys reduce the hillock growth but are undesirable from a plasma etching standpoint because the copper halides are not volatile compounds. Aluminum/titanium alloys, on the other hand, show good resistance to hillock formation and yet possess good etching properties because of the volatile nature of titanium halides.

Homogeneous alloys and layered structures (e.g. AlTi/AlTi ...) of aluminum with titanium, tungsten and copper have been investigated for use in a multilayer interconnection technology. Aluminum is preferred over most other metals because of its low resistivity and silicon compatibility, but it has problems with hillock growth and electromigration. Aluminum/ copper is usually used to solve these problems, but it is difficult to dry etch. In addition, hillocks are not completely eliminated.

An investigation of hillocks, resistivity and interlayer shorts has been undertaken. The films were prepared by either cosputtering aluminum and other elements or by sputtering from composite targets. To test for hillocks, samples were exposed to 450 °C

annealing. When testing for interlayer shorts, a low temperature CVD oxide was used as the insulator.

There are two different types of hillocks observed in aluminum alloys. First, there are the hillocks similar to what one observes on pure aluminum films and second, a new type of hillock which can be two to three times higher than the film thickness. This second type of hillock or *pillar* has been observed in homogeneous films of aluminum with titanium and tungsten. X-ray analysis revealed that these pillars are composed of the same material as the surrounding film. In addition, the concentration of pillars depends on the concentration of impurities in the aluminum film and can be as low as only a few pillars per square centimeter. When a film of pure aluminum is annealed for 30 minutes at 450 °C, the surface roughness is typically 500-1000 Angstroms with occasional small 0.5 to 1 μm pillars. An equivalent aluminum/titanium film has a surface roughness as small as 50 Angstroms (with 6.4 atom.% titanium), but can have 1.5 μm pillars. Aluminum/tungsten films can have surface roughness below 20 Angstroms, but also have pillars as high as 2 μm . It has been found that these pillars can be eliminated by using alternating layers of aluminum with titanium rather than a homogeneous mixture.

Another problem with using aluminum alloys is the increased resistivity of the films. Even homogeneous aluminum films with 2% copper exhibit higher resistance and films of Al/Ti and Al/W can exhibit more than twice the resistivity of pure aluminum films. After anneal, aluminum alloy resistivity drops by 25% making it still better than tungsten films. It has been found that the resistivity of layered aluminum/titanium and aluminum/tungsten films as mentioned above is about the same as pure aluminum films.

5.4.4 Fine Grain Polycrystalline PMOS/SOI Transistors

Small geometry PMOS transistors have been fabricated in fine grain polycrystalline silicon. The use of proton ion implantation has been shown to be an effective means of passivating the grain boundaries in the devices which greatly increases the ON/OFF current ratio which is, in turn, a measure of their suitability as loads in a static RAM cell. [Singh 83] Efforts are under way to incorporate this fine grain technology into a complete CMOS technology which will feature nMOS devices in single crystal silicon and PMOS/SOI loads in the polycrystalline layer.

This work reports results from a study including the weak inversion behaviour of P-channel MOS transistors fabricated in polycrystalline silicon. The devices have a wide range of channel dopings, with channel lengths and widths down to $1.25 \mu\text{m}$. The use of very thin polysilicon enables the gate to modulate the channel conductivity of devices in fine-grain polysilicon with gate voltage excursions of under five volts. The devices have a slow turn-on and exhibit an extended weak-inversion region. Weak inversion currents increase with applied drain voltages up to about 5 volts, with long channel devices also showing this phenomenon. Short-channel effects are seen as the channel length approaches $2 \mu\text{m}$ and are mitigated by using a higher channel doping. Device currents drop sharply below a channel width of $2 \mu\text{m}$, although the narrow-width effect is not as pronounced as the short-channel effects.

Hydrogenation of these devices is being investigated. Ion implantation is being used as the technique of hydrogenation. The initial results indicate that the leakage current in the OFF state is reduced by one order of magnitude, the ON current is increased by two orders of magnitude, the weak-inversion turn-on slope is improved, and mobility is increased.

The hydrogenated PMOS transistors fabricated in fine grain poly-Si appear to be an excellent candidate for loads in static RAM applications. In comparison to poly-Si resistor loads they should offer reduced power consumption and higher speed and in comparison to bulk CMOS they should offer better latch-up immunity.

Staff: K. C. Saraswat, J. P. McVittie, D. Gardner, T. Michalka, B. Bakoglu

Related Efforts: Trotter (Miss. State.).

References: [Saraswat 84, Bakoglu 84, Singh 83]

5.5 Cell Library

We have now remedied an oversight in our initial publication arrangements for the Cell Library with Addison-Wesley, and we are empowered to distribute the CIF for the library over the ARPAnet to DARPA VLSI contractors on a no-fee, no-redistribution basis. Beware of this copyright trap! Contact rob%helens@Score if you would like to obtain a copy of the library this way. (Warning: the library is large, and many mailers have trouble with mobygrams; you may need to make arrangements so that we can use ftp to send the library to you.)

Staff: R. Mathews, J. Newkirk

References: [newkirkLibrary83 83]

5.6 Packaging Technology

An outgrowth of a class design project requires a package with a very large pin count. We are investigating sources, viability, and testure fixturing for 144-pin pin grid and other similar large packages. We will report any workable solutions we find.

Staff: J. Duluk, M. Santoro, J. Newkirk

References

- [Bakoglu 84] H. Brian Bakoglu, and J. D. Meindl.
Optimal Interconnect Circuits for VLSI.
In *1984 ISSCC Digest of Technical Papers*, pages 164-165. IEEE,
February, 1984.
- [BaskettBechtolsheim 82] Bechtolsheim, A., Baskett, F., and Pratt, V.
The SUN Hardware Architecture.
Technical Report 229, Computer Systems Laboratory, Stanford
University, March, 1982.
- [Brown 83] Brown, H., Tong, C. and Foyster, G.
Palladio: An Exploratory Environment for Circuit Design.
Computer 16(12), December, 1983.
- [Buehler 83] M. G. Buehler.
JPL Site Report.
In *DARPA Semi-Annual Report*. DARPA VLSI Contractor's Meeting,
November, 1983.
- [chenATPG84 84] H. Chen, J. Newkirk, R. Mathews.
Test Generation for MOS Circuits.
In *Proc. 1984 ATPG Workshop*. IEEE, Arlington VA, February, 1984.
- [Chow 83] Chow, F.
*A Portable Machine-Independent Global Optimizer - Design and
Measurements*.
PhD thesis, Stanford University, December, 1983.
- [ChowHenn 84] Chow, F.C. and Hennessy, J.L.
Register allocation by priority-based coloring .
In *Proc. of 1984 Compiler Construction Conference*, pages . ACM,
Montreal, Canada, June, 1984.
- [GrossThomas 83] Gross, T.
Code Optimization Techniques for Pipelined Architectures.
In *Proc. Compeon Spring 83*, pages 278 - 285. IEEE Computer Society,
San Francisco, March, 1983.
- [Hennessy 81] Hennessy, J.L.
A Language for Microcode Description and Simulation in VLSI.
In Seitz, C. (editor), *Proc. of the Second CalTech Conference on VLSI*,
pages 253-268. California Institute of Technology, January, 1981.

- [Hennessy 83] Hennesst, J.L.
Partitioning Programmable Logic Arrays.
In *Proc. Int. Conf. on Computer-Aided Design*. IEEE, Santa Clara,
Ca., September, 1983.
- [HennessyGross 83]
Hennessy, J.L. and Gross, T.R.
Postpass Code Optimization of Pipeline Constraints.
ACM Trans. on Programming Languages and Systems 5(3), July,
1983.
- [HennessyJouppi 81]
Hennessy, J.L., Jouppi, N., Baskett, F., and Gill, J.
MIPS: A VLSI Processor Architecture.
In *Proc. CMU Conference on VLSI Systems and Computations*, pages
337-346. Computer Science Press, October, 1981.
- [HennessyJouppi 83]
Hennessy, J., Jouppi, N., Przybylski, S., Rowen, C. and Gross, T.
Design of a High Performance VLSI Processor.
In *Proc. of the Third Caltech Conf. on VLSI*. Calif Institute of
Technology, Pasadena, Ca., March, 1983.
- [Hochschild 83] Hochschild, P., Mayr, E., and Siegel, A.
Techniques for Solving Graph Problems in Parallel Environments.
In *Proc. 24th FOCS Symposium*, pages 351-359. IEEE, 1983.
- [Jouppi 83a] Jouppi, N.
TV: An nMOS Timing Analyzer.
In Randal Bryant (editor), *Proceedings 3rd CalTech Conference on
VLSI*. California Institute of Technology, March, 1983.
- [Jouppi 83b] Jouppi, N.
Timing Analysis for nMOS VLSI.
In *Proceedings 20th Design Automation Conference*, pages 411-418.
1983.
- [Karlin 83] Karlin, A. R., Trickey, H. W., and Ullman, J. D.
Experience With a Regular Expression Compiler.
In *Proc. ICCD*, pages 656-665. IEEE, 1983.
- [Kuo 84] J. B. Kuo, O-H. Kwon, D. C. Galbraith, F-C. Shone, J. D. Shott, J. T.
Walker, R. W. Dutton, and J. D. Meindl.
A 2 μ m Poly-Gate CMOS Analog/Digital Array.
In *1984 ISSCC Digest of Technical Papers*, pages 262-263. IEEE,
February, 1984.

- [Lewyn 84] L. Lewyn and J. D. Meindl.
Physical Limits of VLSI DRAMs.
In *1984 ISSCC Digest of Technical Papers*, pages 160-161. IEEE,
February, 1984.
- [newkirkLibrary83 83] J. Newkirk, R. Mathews.
The VLSI Designer's Cell Library.
Addison-Wesley, Reading MA, 1983.
- [Ossher 83a] Harold L. Ossher and Brian K. Reid.
FABLE: A Programming Language Solution to IC Process Automation
Problems.
In *Proceedings of SIGPLAN '83*. ACM, June, 1983.
- [Ossher 83b] H. L. Ossher and B. K. Reid.
*FABLE: A Programming-language Solution to IC Process Automation
Problems*.
Technical Report 248, Computer Systems Laboratory, Stanford
University, September, 1983.
- [Ossher 84] Harold L. Ossher.
Grids: a New Program Structuring Mechanism Based on Layered
Graphs.
In *Proceedings of POPL '84*, pages 11-22. ACM, Salt Lake City, Utah,
January, 1984.
- [Pfiester 84] J. R. Pfiester, J. D. Shott, and J. D. Meindl.
Performance Limits of NMOS and CMOS.
In *1984 ISSCC Digest of Technical Papers*, pages 158-159. IEEE,
February, 1984.
- [Przybylski 84a] Przybylski, S.A.
Design Verification and Testing of MIPS.
In Paul Penfield, Jr. (editor), *Proc. Third MIT Conference on
Advanced Topics in VLSI*, pages . Mass. Inst. of Tech., Cambridge,
Mass., January, 1984.
- [Przybylski 84b] Przybylski, S., Gross, T., Hennessy, J., Jouppi, N. and Rowen, C.
Organization and VLSI Implementation of MIPS.
Journal of VLSI and Computer Systems 1(3), Spring, 1984.
- [RowenPrzy 84] Rowen, C., Przybylski, S., Jouppi, N., Gross, T., Shott, J., and
Hennessy, J.
MIPS: A High Performance 32-Bit NMOS Microprocessor.
In *Digest of International Solid-State Circuits Conf.* IEEE, San
Francisco, Ca., February, 1984.

- [Saraswat 84] K. C. Saraswat, S. Swirhun, and J. P. McVittie.
Selective CVD of Tungsten for VLSI Technology.
In *165th Meeting of the Electrochemical Society*. The Electrochemical Society, May, 1984.
To be published.
- [Singh 83] H. J. Singh, K. C. Saraswat, J. D. Shott, J. P. McVittie, and J. D. Meindl.
Scaling of SOI/PMOS Transistors.
In *1983 IEDM Digest of Technical Papers*, pages 67-70. IEEE Electron Devices Society, December, 1983.
- [Ullman 84] Ullman, J. D.
Some Thoughts About Supercomputer Organization.
In *Proc. COMPCON 84*, pages 424-432. IEEE, 1984.
An expanded version appears as STAN-CS-83-987, Dept. of CS, Stanford Univ.
- [wolf84 84] W. Wolf.
Compaction Strategies for Integrated Circuit Layout.
PhD thesis, Stanford University, April, 1984.
- [wolfCCAD83 83] W. Wolf, R. Mathews, J. Newkirk, R. Dutton.
Two-Dimensional Compaction Strategies.
In *Proc. International Conf. on Computer-aided Design*. IEEE, September, 1983.

Publications

- [Bakoglu 84] H. Brian Bakoglu, and J. D. Meindl.
Optimal Interconnect Circuits for VLSI.
In *1984 ISSCC Digest of Technical Papers*, pages 164-165. IEEE,
February, 1984.
- [Chow 83] Chow, F.
*A Portable Machine-Independent Global Optimizer - Design and
Measurements.*
PhD thesis, Stanford University, December, 1983.
- [Foyster 84] Foyster, G.
A Knowledge-Based Approach to Transistor Sizing.
Technical Report HPP-84-3, Stanford University Heuristic
Programming Project, 1984.
- [Gross 83] Gross, T.R.
Code Optimization of Pipeline Constraints.
PhD thesis, Stanford University, August, 1983.
- [Hochschild 83] Hochschild, P., Mayr, E., and Siegel, A.
Techniques for Solving Graph Problems in Parallel Environments.
In *Proc. 24th FOCS Symposium*, pages 351-359. IEEE, 1983.
- [Karlin 83] Karlin, A. R., Trickey, H. W., and Ullman, J. D.
Experience With a Regular Expression Compiler.
In *Proc. ICCD*, pages 656-665. IEEE, 1983.
- [Kuo 84] J. B. Kuo, O-H. Kwon, D. C. Galbraith, F-C. Shone, J. D. Shott, J. T.
Walker, R. W. Dutton, and J. D. Meindl.
A 2 μ m Poly-Gate CMOS Analog/Digital Array.
In *1984 ISSCC Digest of Technical Papers*, pages 262-263. IEEE,
February, 1984.
- [Lewyn 84] L. Lewyn and J. D. Meindl.
Physical Limits of VLSI DRAMs.
In *1984 ISSCC Digest of Technical Papers*, pages 160-161. IEEE,
February, 1984.
- [Ossher 83a] Harold L. Ossher and Brian K. Reid.
FABLE: A Programming Language Solution to IC Process Automation
Problems.
In *Proceedings of SIGPLAN '83*. ACM, June, 1983.

- [Ossher 83b] H. L. Ossher and B. K. Reid.
FABLE: A Programming-language Solution to IC Process Automation Problems.
Technical Report 248, Computer Systems Laboratory, Stanford University, September, 1983.
- [Ossher 84] Harold L. Ossher.
Grids: a New Program Structuring Mechanism Based on Layered Graphs.
In *Proceedings of POPL '84*, pages 11-22. ACM, Salt Lake City, Utah, January, 1984.
- [Pfiester 84] J. R. Pfiester, J. D. Shott, and J. D. Meindl.
Performance Limits of NMOS and CMOS.
In *1984 ISSCC Digest of Technical Papers*, pages 158-159. IEEE, February, 1984.
- [Przybylski 84a] Przybylski, S., Gross, T., Hennessy, J., Jouppi, N. and Rowen, C.
Organization and VLSI Implementation of MIPS.
Journal of VLSI and Computer Systems 1(3), Spring, 1984.
- [Przybylski 84b] Przybylski, S.A.
Design Verification and Testing of MIPS.
In Paul Penfield, Jr. (editor), *Proc. Third MIT Conference on Advanced Topics in VLSI*, pages . Mass. Inst. of Tech., Cambridge, Mass., January, 1984.
- [RowenPrzy 84] Rowen, C., Przybylski, S., Jouppi, N., Gross, T., Shott, J., and Hennessy, J.
MIPS: A High Performance 32-Bit NMOS Microprocessor.
In *Digest of International Solid-State Circuits Conf.* IEEE, San Francisco, Ca., February, 1984.
- [Singh 83] H. J. Singh, K. C. Saraswat, J. D. Shott, J. P. McVittie, and J. D. Meindl.
Scaling of SOI/PMOS Transistors.
In *1983 IEDM Digest of Technical Papers*, pages 67-70. IEEE Electron Devices Society, December, 1983.
- [Ullman 84] Ullman, J. D.
Some Thoughts About Supercomputer Organization.
In *Proc. COMPCON 84*, pages 424-432. IEEE, 1984.
An expanded version appears as STAN-CS-83-987, Dept. of CS, Stanford Univ.


```

-- 1 .OF 3
-- 1 - AD NUMBER: A137398
-- 48 - SBI SITE HOLDING SYMBOL:      NFL
--                                     NUC
-- 70 - FIELDS AND GROUPS: 9/1
-- 71 - ENTRY CLASSIFICATION: UNCLASSIFIED
-- 72 - CORPORATE AUTHOR: STANFORD UNIV CA
-- 73 - UNCLASSIFIED TITLE: RESEARCH IN VLSI SYSTEMS
-- 74 - TITLE CLASSIFICATION: UNCLASSIFIED
-- 75 - DESCRIPTIVE NOTE: TECHNICAL PROGRESS REPT.
-- 11 - REPORT DATE: NOV 1983
-- 12 - PAGINATION: 7P
-- 13 - CONTRACT NUMBER: MDA903-79-C-0680, MDA903-82-C-0335
-- 20 - REPORT CLASSIFICATION: UNCLASSIFIED
-- 21 - SUPPLEMENTARY NOTE: SPONSORED IN PART BY CONTRACTS NO-833-RCHD-
-- 0432 AND MDA903-83-C-0107
-- 23 - DESCRIPTORS: *MICROELECTRONICS, *INTEGRATED CIRCUITS,
-- *PARAMETRIC ANALYSIS, TEST METHODS, PERFORMANCE(ENGINEERING), SOLID
-- STATE ELECTRONICS, CHIPS(ELECTRONICS), BREAKDOWN(ELECTRONIC
-- THRESHOLD), VOLTAGE, CAPACITORS, COMPUTER AIDED DESIGN,
-- MINICOMPUTERS
-- 24 - DESCRIPTOR CLASSIFICATION: UNCLASSIFIED
--      <<P FOR NEXT PAGE>> OR <<ENTER NEXT COMMAND>>
-- 25 - IDENTIFIERS: *VLSI(VERY LARGE SCALE INTEGRATION)
-- 26 - IDENTIFIER CLASSIFICATION: UNCLASSIFIED
-- 27 - ABSTRACT: THE PARAMETRIC TEST FACILITY PROVIDES TWO ESSENTIAL
-- SUPPORT FUNCTIONS FOR THE FAST TURN-AROUND LABORATORY - FEEDBACK OF
-- PROCESS-CONTROL INFORMATION TO THE FABRICATION LINE, AND
-- CHARACTERIZATION OF ACTIVE DEVICES AND PARASITICS FOR CIRCUIT
-- DESIGN. THE BASIC PARAMETRIC TESTER, CONSISTING OF A RUCKER AND
-- KOLLS MODEL 1032 STEPPING WAFER PROBER, AN HP 4145 SEMICONDUCTOR
-- PARAMETER ANALYZER, AND AN HP 8942A SWITCHING MATRIX CONTROLLED BY
-- AN HP 9845 CALCULATOR, HAS BEEN UPGRADED BY ADDITION OF AN HP 3400
-- SYSTEM, DUM, AN HP 5315 FREQUENCY COUNTER, AND AN HP 4871
-- CAPACITANCE BRIDGE. THE DUM GIVES INCREASED SYSTEM FLEXIBILITY
-- WHILE THE FREQUENCY COUNTER PROVIDES FOR THE DIRECT MEASUREMENT OF
-- RING OSCILLATOR PERFORMANCE. CUSTOM HARDWARE HAS BEEN DESIGNED AND
-- BUILT FOR THE RAPID EXTRACTION OF MOS TRANSISTOR THRESHOLD AND
-- CAPACITOR BREAKDOWN VOLTAGES. THE SYSTEM INCLUDES THE NECESSARY
-- SOFTWARE TO MEASURE SUCH PROCESS CONTROL PARAMETERS AS LINE WIDTHS,
-- SHEET RESISTANCES, AND THRESHOLD AND BREAKDOWN VOLTAGES. SOFTWARE
-- HAS ALSO BEEN DEVELOPED FOR THE STATISTICAL ANALYSIS OF THESE DATA
-- ON THE HP 9845 COMPUTER. THIS SOFTWARE CAN PRODUCE PLOTS OF
-- PARAMETER VARIATIONS ACROSS A WAFER AS WELL AS THE USUAL MEANS,
-- STANDARD DEVIATIONS, AND HISTOGRAMS. (43709)
--      <<P FOR NEXT PAGE>> OR <<ENTER NEXT COMMAND>>
-- 30 - ABSTRACT CLASSIFICATION: UNCLASSIFIED
-- 31 - INITIAL INVENTOR: 12
-- 33 - LIMITATION CODES: 1
-- 35 - SOURCE CODE: 332550
-- 36 - DOCUMENT LOCATION: NTIS
-- 40 - GEOPOLITICAL CODE: 0612
-- 41 - TYPE CODE: 1

```

END

FILMED

8

1/1