

12

Semiannual Technical Summary

AD-A148 145

Restructurable VLSI Program

31 March 1984

Lincoln Laboratory

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

LEXINGTON, MASSACHUSETTS



Prepared for the Defense Advanced Research Projects Agency
under Electronic Systems Division Contract F19628-85-C-0002.

Approved for public release; distribution unlimited.

DTIC
ELECTE
NOV 30 1984
A

DTIC FILE COPY

84 11 29 007

The work reported in this document was performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology. This work was sponsored by the Defense Advanced Research Projects Agency under Air Force Contract F19628-85-C-0002 (ARPA Order 3797).

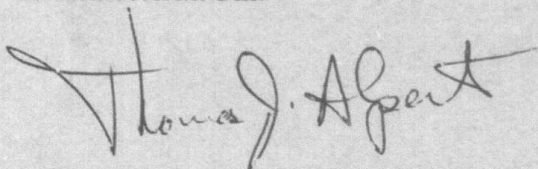
This report may be reproduced to satisfy needs of U.S. Government agencies.

The views and conclusions contained in this document are those of the contractor and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the United States Government.

The Public Affairs Office has reviewed this report, and it is releasable to the National Technical Information Service, where it will be available to the general public, including foreign nationals.

This technical report has been reviewed and is approved for publication.

FOR THE COMMANDER

A handwritten signature in dark ink, reading "Thomas J. Alpert". The signature is written in a cursive style with a large initial "T" and a long horizontal stroke at the end.

Thomas J. Alpert, Major, USAF
Chief, ESD Lincoln Laboratory Project Office

Non-Lincoln Recipients

PLEASE DO NOT RETURN

Permission is given to destroy this document
when it is no longer needed.

**MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LINCOLN LABORATORY**

RESTRUCTURABLE VLSI PROGRAM

**SEMIANNUAL TECHNICAL SUMMARY REPORT
TO THE
DEFENSE ADVANCED RESEARCH PROJECTS AGENCY**

1 OCTOBER 1983 — 31 MARCH 1984

ISSUED 15 OCTOBER 1984

Approved for public release; distribution unlimited.

**DTIC
SELECTE
NOV 30 1984
S D
A**

LEXINGTON

MASSACHUSETTS

ABSTRACT

This report describes work performed on the Restructurable VLSI Program sponsored by the Information Processing Techniques Office of the Defense Advanced Research Projects Agency during the period 1 October 1983 through 31 March 1984.



Accession For	
DDIC	<input checked="" type="checkbox"/>
DDIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

TABLE OF CONTENTS

Abstract	iii
List of Illustrations	vi
I. PROGRAM OVERVIEW AND SUMMARY	1
A. Overview	1
B. Summary of Progress	2
1. Design Aids for RVLSI	2
2. Applications	3
3. Testing	4
II. DESIGN AIDS FOR RVLSI	7
A. MACPITTS	7
1. Organelle Design	7
2. Layout Routines	7
B. Lincoln Boolean Synthesizer	8
C. Chip Assembler	8
D. Linking Shell (LSH)	9
E. Restructurable Wafer Editor (RWED)	10
III. APPLICATIONS	11
A. Dynamic Time Warping System	11
B. Digital Integrator	15
IV. TESTING	17
A. Tester on Chip (TOC)	17
B. Optical Probe	17
1. Light Source Assembly	18
2. Microscope	19
3. TV Camera and Monitor	19
4. Detector Circuit	19
5. Optical Probe Driver	19
C. Test Vector Generation	19
D. Testing Strategies	22
References	25

LIST OF ILLUSTRATIONS

Figure No.		Page
1	DTW Distance and Path Computer	12
2	DTW Level Building Processor	13
3	Packaged RVLSI Integrator Wafer	14
4	Assembled Optical Prober	18
5	Electrical and Optical Characteristics of a Laser Source	20

RESTRUCTURABLE VLSI PROGRAM

I. PROGRAM OVERVIEW AND SUMMARY

A. OVERVIEW

The main objective of the Lincoln Restructurable VLSI Program (RVLSI) is to develop design methodologies, architectures, design aids, and testing strategies for implementing wafer-scale systems with complexities approaching a million gates. In our approach, we envisage a modular style of architecture comprising an array of cells embedded in a regular interconnection matrix. Ideally, the cells should consist of only a few basic types. The interconnection matrix is a fixed pattern of metal lines augmented by a complement of programmable switches or links. Conceptually, the links could be either volatile or nonvolatile. They could be of an electronic nature, such as a transistor switch, or could be permanently programmed through some mechanism such as a laser. The RVLSI Program is currently focusing on laser-formed interconnect.

The link concept offers the potential for a highly flexible, restructurable type of interconnect technology that could be exploited in a variety of ways. For example, logical cells or subsystems found to be faulty at wafer-probe time could be permanently excised from the rest of the wafer. The flexible interconnect could also be used to circumvent faulty logic and tie in redundant cells judiciously scattered around the wafer for this purpose. Also, the interconnect could be tailored to a specific application in order to minimize electrical degradations and performance penalties caused by unused wiring and links.

Further, the testing of a particular logical subsystem buried deep within a complex wafer-scale system poses a very difficult problem. A properly designed restructurable interconnect matrix could be temporarily configured to improve both the controllability and observability of internal cells from the wafer periphery. In this way, each component cell or a manageable cluster of cells could be tested in straightforward manner using standard techniques. With an electronic linking mechanism, it is possible to think in terms of a dynamically reconfigurable system. Such a feature could be used to alter the functional mode of a system subject to changes in the operating scenario, or it could be used to support some degree of fault tolerance if the system architecture was suitably designed.

Several major areas of research have been identified in the context of the RVLSI concept:

- (1) System architectures and partitioning for whole-wafer implementations.
- (2) Placement and routing strategies for optimal utilization of redundant resources and efficient interconnect.
- (3) Assignment and linking algorithms to exploit redundancy and flexible interconnect.
- (4) Methods for expediting cell design with emphasis on functional level descriptions, enhanced testability, and fault tolerance.
- (5) Methods for testing complex, multiple-cell, whole-wafer systems.

Complementary work on the development of various link and interconnect technologies as well as fabrication/processing technology is being supported by the Lincoln Air Force Line Program, and results are reported under the Lincoln Laboratory Advanced Electronic Technology Quarterly Technical Summary.

B. SUMMARY OF PROGRESS

Work for this period is reported under three headings: Design Aids for RVLSI (Section II), Applications (Section III), and Testing (Section IV). The following paragraphs summarize progress to date.

1. Design Aids for RVLSI

During the last reporting period significant changes were made to the basic floor plan of the MACPITTS silicon compiler target architecture. These were made possible by the substitution of a channel router for the old rivers and ring routers. The new router makes it possible to position I/O pads on all four edges of the layout thereby reducing the area devoted to connections between the internals of the chip and the periphery. To further improve layout efficiency and decrease chip size, the flag generator was rewritten to produce more compact structures. Super-buffers are used for clock distribution and the basic flag cell contains fewer transistors. The cell was designed with an input pitch matched to the capabilities of the unconstrained channel router, and has an unstretched height of about 80λ . The ability to position flags on either side of the internal routing (between the data path and control section) was also incorporated.

The new router has also made it possible to eliminate gate ordering constraints previously imposed on the Weinberger array. This further improves layout efficiency in the control section and is expected to improve the speed of the ordering process. Additionally, improvements were made in the layout of the power distribution grid. Full installation of the new router is in progress and remains to be completed.

The Lincoln Boolean Synthesizer (LBS) is being retargeted for $3\text{-}\mu\text{m}$ CMOS. New input and output pads have been designed and some changes have been made in the basic layout of the array, since the new design rules involve more than a simple scaling of the old JPL CMOS rules. Testing of sample layouts is under way using the design rule checker and CMOS node extractor. Further refinements include an output option allowing LBS to produce files necessary for its use in conjunction with the chip assembly tool, and some modifications to the min-cut ordering routine to increase running speed.

The Chip Assembler (HCA) has been installed for general usage and documentation has been prepared. Routing routines were improved after substantial experimentation and evaluation, and programs for appropriate merging of output files have been completed. Features include:

- (a) Manual design of cells using CAESAR and direct passing of cells from LBS.
- (b) Usage of these cells to construct new, more complex cells in a hierarchical manner.
- (c) Automatic interconnection of cells from a net list.

- (d) Manual interconnection to modify, complete or augment the automatic routing.
- (e) Expanded CIF output of the entire finished design.
- (f) A table-driven mechanism for inputting all "technology dependent" parametric information.
- (g) A minimal set of usage prompt and error messages allowing operation of the system by naive users.

It remains to generate actual test circuits for submission to MOSIS thereby validating the fully operational status of the system.

An effort is under way to modify the generalized linker (LSH) to support an incremental zap-and-test restructuring strategy. The objective is to map the unordered set of laser cuts and zaps which implement the desired net list into an optimum sequence (possibly including extra cuts and zaps) such that as each net is formed, laser probing can be used to verify integrity. This is accomplished by forming a "test net," using only unassigned tracks, and connecting to at least one wafer I/O pin. To achieve this, the "zap" program was modified to group together all zaps and cuts belonging to the same signal net. Also, a strategy was developed to construct the test net. This has necessitated some modifications to LSH, most of which are complete at this time. It is expected that the new option will be available soon and will be exercised in restructuring the next integrator wafer.

The Restructurable Wafer Editor (RWED) has been ported to a 68000-based microcomputer system, which supports operation of the laser table facility and coordinates testing functions while performing in an iterative zap-and-test mode. RWED has been interfaced to the test control routines, and is now able to control laser power and perform other functions relating to a coordinated, automated zap-and-test capability.

2. Applications

Integrator wafer i4w6 has been successfully restructured and demonstrated at the tester limit of 27 MHz, thereby exceeding the 25 MHz design goal. To implement a fully operational wafer-scale system, 64 of the available 192 cells must be functional. In this case 81 good counter cells were found, 58 of 60 input amplifiers were good, and all 32 output amplifiers were functional. Only 26 of the 568 wafer-length interconnect tracks were defective, which greatly simplified the task of *assignment and linking*. 1876 laser connections and 137 cuts were required to implement the system.

Some minor clearance problems with the link layout were observed and have been corrected for subsequent production runs. Two cells were found to have failed after the linking process but were easily replaced by nearby neighbors. The cell failures were inadvertently undetected at wafer-probe time and were not induced by the linking process itself.

A portable demonstration unit was built which exercises the restructured integrator in a realistic manner. The wafer has been successfully demonstrated at speed via this mechanism with simultaneous, asynchronous reads and writes.

Development continues of a systolic wafer-scale array to implement a form of the Myers-Rabiner dynamic time warping/level-building algorithm for connected word recognition. An architectural variation on the original wafer design is being pursued to improve throughput and simplify external control hardware. For improved throughput, a parallel arithmetic distance computer was designed to calculate a squared euclidean metric using table lookup techniques. A 16-element metric can be calculated in $1 \mu\text{s}$ which matches the expected speed of the present path computer design. To simplify external control, consideration has been given to adding a level building processor element array to the basic processor array. The dynamic logic array portion of the new distance computer cell has been designed in detail using CAESAR, has been node extracted and switch level simulated, and has been submitted to the 1 March $3\text{-}\mu\text{m}$ CMOS run.

Also during this period the development of a flexible, "C" language, non-real-time functional simulation of the DTW wafer was undertaken. The system was designed to accurately model (on a bit-by-bit basis) the actual computations performed by the DTW wafer. This system will allow final decisions to be made regarding numeric accuracy and choice of distance measure.

3. Testing

A new optical probe source has been built for application to dynamic CMOS circuits. A simplified mechanical assembly has been designed which comprises a low-power laser diode, packaged with a collimating lens in a cylinder, which mounts easily in the eyepiece of a binocular microscope. As was expected, discharging of dynamic nodes was encountered due to migration of carriers generated deep within the silicon substrate. By shortening the optical input to a single pulse of a few microseconds duration, it was possible to probe successfully without destroying dynamically stored state information. Further experimentation is in progress to establish interaction distances and to determine operating parameters necessary to guarantee the integrity of state.

None of the Tester-on-Chip (TOC) chips received from the M37A MOSIS $3\text{-}\mu\text{m}$ NMOS run was found functional. After careful study it was determined that an excessive voltage drop in the power distribution to the tri-state I/O pads was suffered due to an inordinately long diffusion run. The MACPITTS pad library has been modified to correct this layout anomaly and the new pad designs have been substituted directly into the TOC chip layout file. This new design will be submitted to the next $3\text{-}\mu\text{m}$ NMOS run, which is currently scheduled for May 1984.

A new computer program for test vector generation has been implemented called BANDITS (Boolean Analysis of Digital Timeless Systems). This program accepts the gate-level description of a logic circuit and produces a reduced set of input patterns that will drive the circuit outputs to the logic 0 and 1 states. The program generates excitation patterns for solving a network in the steady state, and ignores the effects of propagation delays on a network's response as well as rejecting patterns that will cause oscillatory conditions. The current version can accommodate circuits up to 2000 gates and having up to 30 primary inputs and/or state variables.

An effort has been undertaken to extend current scan-set techniques for use in testing very complex (e.g., wafer-scale) VLSI components with built-in testing resources. The approach uses a prime linear shift register (LFSR) as a source of pseudorandom bit patterns which are shifted along the

system scan path (i.e., the system latches configured as a single long shift register). The scan path strategy reduces the testing problem to one of testing only the combinational parts of a system. The combinational logic is partitioned in "cones" connected along the length of the scan register whose outputs are hashed into a signature register. Extensive probabilistic analyses have suggested that very good coverage can be obtained by hashing the signature register with the scan-set register thereby eliminating the explicit need for a separate source LFSR. This amounts to implementing a very sophisticated random number generator which analysis indicates is capable of providing acceptable test vector coverage for all logic cones. Considerable theoretical effort has been devoted to developing the constraints and bounds on expected fault coverage.

II. DESIGN AIDS FOR RVLSI

A. MACPITTS

Work on improvements to the MACPITTS silicon compiler has progressed in several areas. Once completed, these additions will have a significant effect on layout efficiency and circuit performance.

1. Organelle Design

The flags (one-bit registers) used by the compiler were redesigned to obtain a smaller layout, employ less transistors, and incorporate super-buffers for the clock lines.

The pitch of the input lines now allows the use of a channel router (instead of the river router) for the interconnection between the data-path and Weinberger array sections. These new flags can be placed to the right of the data-path (as before) or to the right of the Weinberger array, resulting in a decrease in the horizontal dimension of the chip.

A manual describing the L5 layout language has been published.¹ This layout language is used by MACPITTS and LBS.

2. Layout Routines

a. The channel router described in the previous Semiannual Technical Summary² is being incorporated into MACPITTS. This router will replace the river router used to generate the interconnection between the data-path and control sections. The use of this channel router will allow for unconstrained optimization of the columns in the Weinberger array and units in the data-path, as certain connection points will not need to maintain their linear relative positions as required when a river router is employed. The new router will also support placement of the flags in two locations, namely to the right of the data-path and right of the control section, thereby reducing both the vertical and horizontal dimensions of the layout.

b. A new scheme for pad placement is being developed to allow placement of the pads on the four sides of the chip. This will require modifications to the ring router used to connect the pads to the interior of the layout.

In the course of designing a very complex chip with MACPITTS, a problem was discovered in the design of the multiplexers. In a very special case, the voltage at the input may not be correct due to an unacceptable voltage drop along a polysilicon line. The situation can be easily corrected manually but the solution remains to be incorporated into the compiler.

A graduate student at M.I.T. has used MACPITTS to produce a set of chips to simulate a neural model related to the human auditory system. These chips will be back from MOSIS shortly and give us the opportunity to do some further electrical testing and performance evaluation.

B. LINCOLN BOOLEAN SYNTHESIZER

A new version of LBS has been installed. This version is implemented in the 3- μm CMOS technology that has become the new MOSIS standard for CMOS. As LBS operates in two phases (i.e., a technology independent logic generator and placement optimizer, followed by a layout generator), it was necessary to modify only the second section of the program.

The new design rules cannot be characterized as a simple rescaling of the previous ones, thereby forcing changes to some aspects of the layout. In any case, the differences between the old 5- μm JPL CMOS and the new 3- μm CMOS layouts are minor.

There are new input and output pads, using the same circuit as before. The rule on extensions around the contacts is different: extensions are not the same in all directions for poly-to-metal-1 connections. This forced several changes in the layout of ground connections in addition to the actual poly-to-metal-1 connections.

This LBS implementation does not exploit a second metal layer as two-metal layer technology is offered only in a restricted number of MOSIS runs.

C. CHIP ASSEMBLER

A first version of the Chip Assembler is now available for general use.

The Chip Assembler allows:

- (1) Design of basic cells manually using CAESAR as the graphic editor, and/or directly from LBS;
- (2) Use of these basic cells to build new cells by manual placement;
- (3) Automatic routing of signals as specified in a net list. The routing is done in two steps, a global routing phase followed by a channel router acting on independent channels (all the routing is done in two layers);
- (4) Manual intervention to modify, complete, and add to the automatic routing (this must be done for the power and ground lines);
- (5) Expand CIF output of the whole design;
- (6) All the technology dependent information necessary for routing (e.g., minimum widths of lines and clearances) is read from files, allowing its use with any two-layer technology. At this moment, we support 4- μm NMOS, 5- μm and 3- μm CMOS. The routing is done using the polysilicon and first-metal layers.

Some improvements were added to the global and channel routers once we were able to examine the effects of the routines on a good number of examples. The placement routines for determining the channel crossing points in the global router were expanded to consider some special cases that result in a smaller number of jogs in the final interconnection. The channel router was modified to solve some problems due to the possible appearance of fixed terminals along the four sides of a

channel whose dimensions cannot be altered. This situation is different from the usual one encountered in the channel router problem, where terminals are located only on two opposite sides and the width of the channel is variable.

The usage and error messages were expanded to constitute a minimal set that allows its utilization by a nonexpert user. A brief internal user's manual and simple illustrative examples are available.

D. LINKING SHELL (LSH)

An effort has been initiated for automating the testing of metal line interconnections among the cells of a restructurable wafer. The object is to map the unordered set of laser cuts and zaps that comprise the interconnections onto an optimum sequence (possibly including additional cuts and/or zaps) so that as each net is formed, laser probing can be used to verify integrity. To achieve this, a "test-net" is formed on the wafer, utilizing only those track segments and programmable links which have not been reserved (by the routing program) for building the "signal-nets." An initial signal-net is assumed to be provided by the user. This net can be as simple as a pair of horizontal and vertical tracks which are linked together as well as being linked to a wafer I/O pin, called the test-pin. Testing can then be achieved by (1) forming the signal-net, (2) linking it to the test-net, (3) using the laser probe to excite each cell pin tied to the signal-net, in turn, to observe the photocurrent at the test-pin, and finally (4) separating the signal- and test-nets from each other. The most critical step in this approach is step (2), where the two nets must be linked together with minimal utilization of the wafer's resources. This is necessary since cutting track segments and links, being an irreversible action, may hamper future linking of the cells if future linking runs become necessary to modify the wafer. For the same reason, it is best to generate the linking commands for the entire wafer before generating the commands for testing.

The algorithm for linking a signal-net to the test-net aims at finding a programmable link between a pair of vertical and horizontal track segments, one from each net. If found, that link is used to join the two nets together. After testing, the link is cut, leaving both the signal- and test-nets unaffected. If such a link could not be found, the algorithm seeks to find a single track segment that intersects (i.e., has a programmable link to) a pair of horizontal and vertical track segments that belong to the signal-net and the test-net, respectively. If successful, the links are formed and testing is done. After the test, only the link to the signal-net is destroyed. This way, the signal-net is left unaffected but an additional track segment is added to the test-net. Thus, as each signal-net is tested, the test-net may "grow," making it easier to find the necessary links/track segments for testing subsequent signal-nets.

To implement the above approach, first it was necessary to modify the "zap" program so that all zap/cut commands that belong to the same signal-net are grouped together. This has been achieved and the new version of the "zap" program has been released. The next significant task has been to modify/extend the internal data storage formats of LSH so that information necessary to implement the test algorithm can be conveniently stored. This task and the implementation of the algorithm itself are nearing completion. A write-up describing how to use the new test option in LSH has already been prepared and included in the LSH documentation.

E. RESTRUCTURABLE WAFER EDITOR (RWED)

RWED, the program which controlled the laser table from the VAX through an Apple microprocessor, has been ported to a 68000 microprocessor. Command files are transferred from the VAX and report files written back. The advantages are several: first, programs which used to reside on the Apple and were difficult to maintain due to limited Apple availability can now be written and debugged on the VAX; second, response time is less sensitive to VAX loading; third, and most important, there can be a closer integration of the restructuring and testing processes since the same computer controls both. RWED commands can now initiate both laser-probe and functional testing and test results can influence the progress of restructuring. These changes have been checked out but not yet used in an actual restructuring exercise.

A reference manual describing the RWED was published.³

III. APPLICATIONS

A. DYNAMIC TIME WARPING SYSTEM

Several changes have taken place in the DTW architecture, prompted by results indicating that previously reported performance criteria may have been too stringent in areas which directly affect recognition throughput.⁴ Also, a rethinking of the original architecture (diagonal array) led to a more efficient implementation of the system as a row of processors. The new system architecture allows the simultaneous execution of node processing and system I/O. A modified system design incorporating contributions from both alternatives is now being pursued.

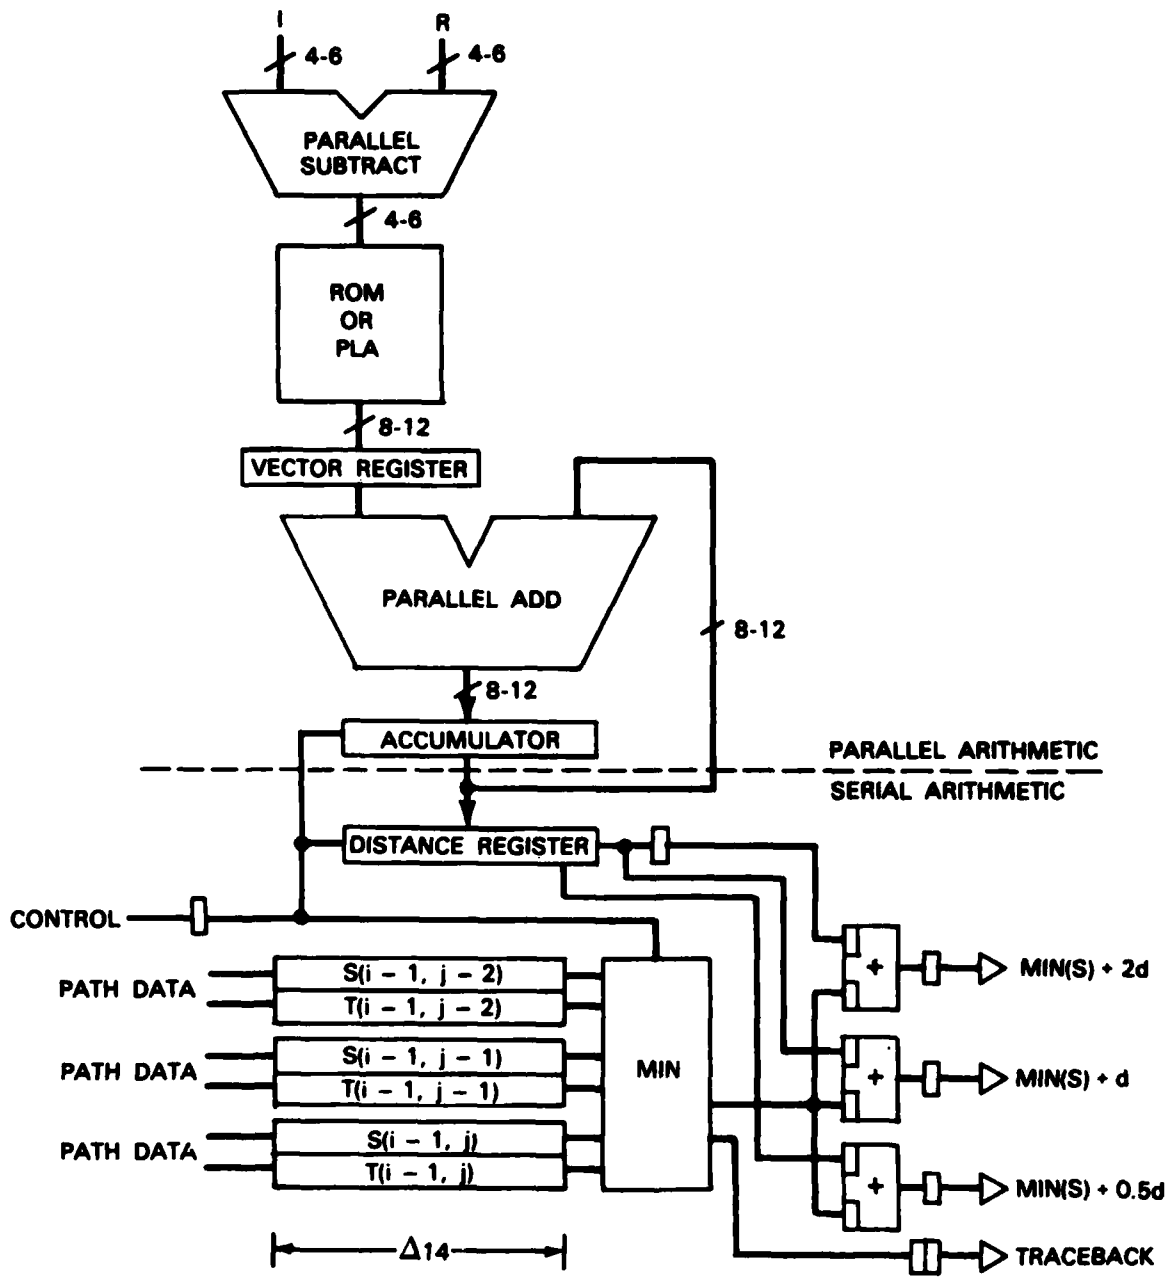
The new architecture is still targeted at implementing the Myers-Rabiner dynamic time warping (DTW) level building algorithm.⁵ However, instead of calculating distances using bit-serial arithmetic, the new distance computer involves a parallel subtraction followed by a nonlinear operation such as squaring or absolute value formation. The nonlinear operation is performed using a programmable logic array (PLA) on the result produced in the subtraction. This reduces the time necessary for completing a distance calculation by a factor of 10, thus allowing full utilization of the bit-serial path computer (Figure 1). A portion of the distance computer has been designed and submitted for fabrication to a 3- μ m CMOS MOSIS run.

Another modification in the proposed architecture is the inclusion of a processor for performing the higher-level functions of level building normally handled by an external general-purpose processor (Figure 2). This significantly reduces the performance requirements of the external controller to simply handling memory management and communication functions. However, this requires that the reference templates be normalized to a predetermined length before use and would require investigating new front-end techniques for handling frame data, such as downsampling and template pre-warping, as well as their effects on performance.

To substantiate the performance estimates of the proposed system, a high-level language, non-real-time simulation has been written in "C." This simulation includes the capability of varying several system parameters, such as distance metric or row width, while gathering statistics on both word and string errors for several input utterances. Following is a list of the current variable system parameters:

- (1) Row Width (number of processors)
- (2) Distance Metric
- (3) Path Constraint
- (4) Endpoint Relaxation Parameters (see Reference 5):
 - (a) Delta R1
 - (b) Delta R2

The parameters are altered by specifying a starting value, a step size (auto-increment), and a limit. The specifications are entered through a command file by way of assignment statements. If a



140807-N

Figure 1. DTW distance and path computer.




Figure 2. DTW level building processor.




Figure 3. Packaged RVLSI integrator wafer.

parameter is not specified, a default value is used. A default step size is assigned if none is given, but a limit must always be specified when using the auto-increment feature. The command file also allows overriding of default file names used for finding the reference vocabulary and the input test utterances. These files contain a list of the names of the parameter files to be used in the matching process.

At the conclusion of each input template match, the resulting reference string is compared to the input word string and all resulting word errors are classified and recorded in an error record specifically maintained for this system configuration. All word substitutions are retained in a vocabulary substitution list to identify commonly confused words. If a word exists in the input which is not included in the reference vocabulary, it is added to the working vocabulary and tagged for identification. This should not be confused with adding it to the actual reference vocabulary, as it is merely a means of identification.

Still to be added is a package for interpreting the error statistics gathered in the above process. It should be easy to implement a powerful set of graphics routines for plotting error rate vs one or two of the system parameters which were varied. This would allow easy identification of trade-offs in hardware complexity and system throughput to actual system performance.

These simulations will provide concrete answers to questions surrounding the performance of any proposed architectural changes or algorithmic variations. It will also provide us with the means for both obtaining conclusive results in those areas of disagreement among the research community, and exploring the DTW algorithm's performance in other environments and applications. Once these answers have been obtained, we will be in a good position to immediately finalize the detailed design of an appropriate wafer-scale system implementation.

B. DIGITAL INTEGRATOR

Integrator wafer i4w6 has been successfully laser restructured and operates at 27 MHz; it is our first restructured wafer (Figure 3).

Of the 192 counter cells this wafer had 81 which passed wafer-probe tests on the Tektronix S3260 tester. Sixty-four good cells are required to make an integrator. Fifty-eight of the 60 input amplifiers and all the 32 output amplifiers were good. Only 26 of the 568 wafer-length tracks were defective. With such a high interconnect yield there were no problems in doing an automatic assignment and linking of the wafer with the LSH programs. Likewise, the fully automatic generation of laser command files was done without difficulty. Considerable effort was required to generate some special command files which the LSH program was not able to create. At this point in the development of the linking process, it was deemed desirable to test each laser connection and cut by laser probing an active device on the wafer and sensing the presence of a photocurrent through a laser connection or the absence of a photocurrent due to a cut at a package pin. This was easy to do for the bus signals because all internal pins are accessible and no special connections are required. It is not true, however, for the data and select signals which pass from cell to cell. For these signals we first generated a net which connected together all signals in a logical column and then to a package pin so that laser connections could be tested. Then cuts were made working from

the end so as to form the final nets and allow testing of the cuts in the process. This required human intervention to generate a strategy for each of the eight logical columns and create correctly ordered command files, a process which was time consuming and error prone. To make the system connections, 1876 laser connections and 137 cuts were required. The wafer has about 40,000 link positions. Not a single problem was observed with the laser connections, but many problems were encountered with cuts. The laser link layout did not include adequate space for cuts, especially with the positioning errors in our laser table. To compensate, a cutting procedure was developed which made several cuts but unfortunately often created a high-impedance connection through charred polyimide. These connections could be detected with laser probing, and additional cutting would open them up. Therefore, the additional effort for testing the cut links was essential to success. A cutting procedure has since been developed which is much improved over the one used here.

Our strategy with this wafer was to link up one logical column at a time, and then laser probe the nets and functionally test the cells. Functional testing was done with the wafer on the laser table using special test circuitry controlled from the VAX computer through a 68000 microprocessor controller. There was only one interconnect problem, and that was related to bonding pad spacing. During restructuring two cell problems occurred: one cell had two defective counters, and one cell suffered a slow output on one of its four counters. These cells were not replaced during the initial structuring. The input amplifiers were not initially connected so that all internal nets could be measured from package pins. The completed wafer operated correctly except for the cells noted above, but only at slow speed due to the absence of drivers for the high-speed clock lines.

The fault in the cell with two defective counters can be accounted for by one internal line stuck low. The wafer-probe test sequence was deficient in not covering this particular condition and it has been corrected. The slow output is apparently due to a weak pull-up transistor internal to a cell. The corrected test program should now catch this fault also.

The cell with two defective counters was replaced with a spare cell by cutting off the defective one, connecting the spare, and rerouting several lines using both manual and automatic methods. We will not replace the other cell since it does not prevent correct operation. Input amplifiers were inserted using the new cutting procedure.

The input data shift registers are specified to operate at 25 MHz. The completed wafer operated up to 27 MHz on the Tektronix S3260 tester. No link trimming has been done on this wafer, that is, only the cuts necessary to isolate signals have been made. Therefore some signal lines may have excess capacitance due to unused links. These refinements will be made on a future wafer. In the packet radio system incrementation of the counter can take several microseconds, and the 256 counters must be read out at a maximum rate of about 0.8 MHz. There is no problem with the increment rate. The readout rate requirement is satisfied, but with less than optimum margin due to the one slow signal, untrimmed lines, and a minor cell design error which makes the output signals especially sensitive to track capacitance.

A small stand-alone exerciser has been built which demonstrates operation of the circuit at a 25-MHz write rate and with overlapped, asynchronous read and write similar to a real application.

IV. TESTING

A. TESTER ON CHIP (TOC)

Thirteen TOC chips were received from the M37A MOSIS 3- μ m NMOS run in November. They were tested on the Tektronix S3260, using a pattern file generated by the nl simulation used to verify the design. A systematic problem was immediately discovered. The logic high output on the tri-state pads was about 2 V. It was later determined that this was the result of a combination of increased transistor conductivity and higher diffusion sheet resistance, peculiar to the Hewlett-Packard process. The power supply voltage thus dropped across the side of the pad on the way to the pad buffers.

However, with the tester thresholds compensated appropriately, one of the TOC chips was about 90 percent functional. The rest failed completely. Optical inspection of some chips revealed that they did not suffer from the incomplete metallization etching problems seen on the previous run.

This run also included four of the simple test designs. They also suffered from the low-voltage-output problem. Initially, none of them worked, even with the tester thresholds set low. Later, these and the TOC chips were retested at $V_{dd} = 6.5$ V. Three of the test chips passed the functional test (at 5 MHz). The one TOC chip that almost worked only did so with V_{dd} within 0.2 V of 5 V. None of the other TOC chips showed any improvement at higher supply voltages.

To check our hypothesis regarding the failure mechanisms of the TOC chips, the pads have been modified by hand and substituted directly in the layout file. This new layout should correct the voltage drop problem. The new design is being submitted to the next 3- μ m MOSIS run, currently scheduled for May 1984.

B. OPTICAL PROBE

Lincoln Laboratory is preparing a manual which describes the assembly and operation of an optical prober suitable for examining the states of the internal nodes of a packaged integrated circuit while it is being excited at its input pins by test vectors under normal operating conditions. It is meant to replace the microprober commonly used as a debugging aid by IC designers attempting to diagnose new designs which fail to operate as anticipated. It should provide an inexpensive, non-invasive means for probing any node in a complex circuit, thereby expediting design verification and debugging.

It is intended to provide enough information in this manual so that DARPA contractors will be able to purchase the component parts and assemble a complete prober with a few hours of electronic technician help.

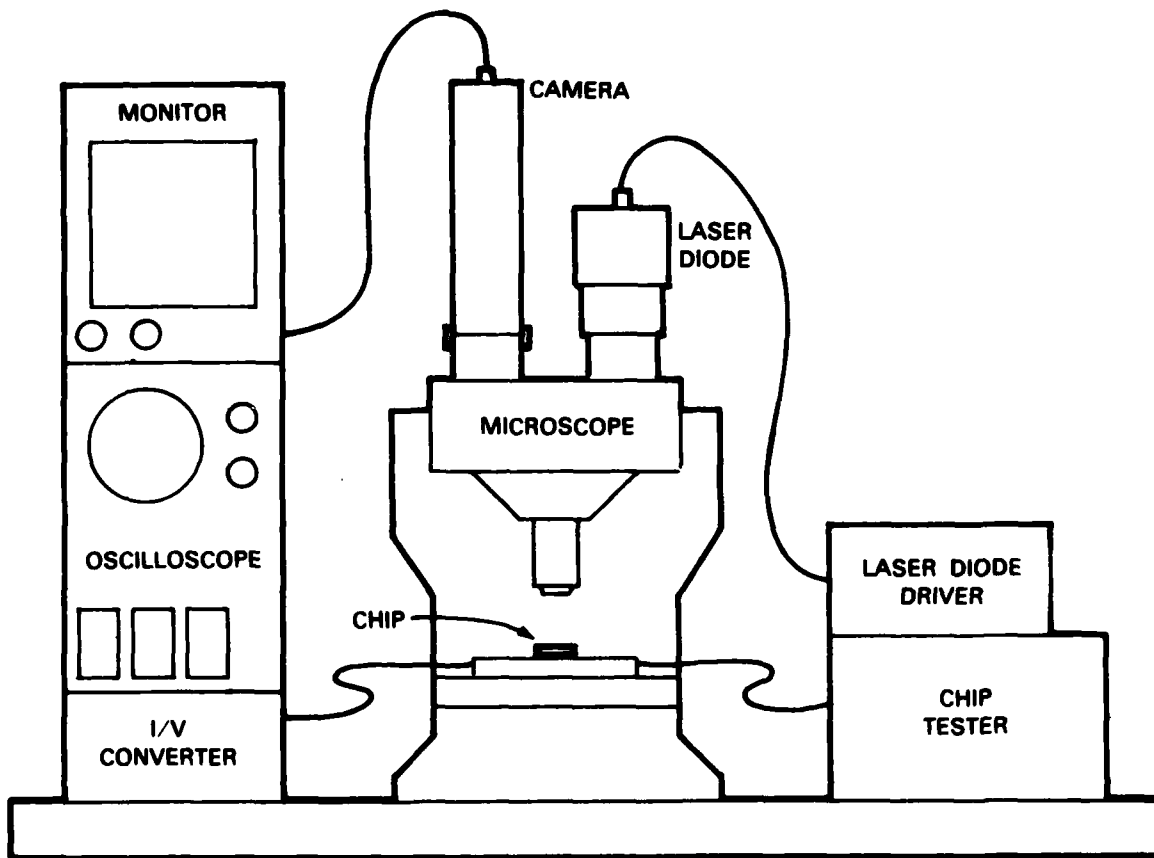


Figure 4. Assembled optical prober.

A sketch of the assembled optical prober is shown in Figure 4. It comprises the following parts which will be described in detail below:

- (1) Light Source Assembly
- (2) Microscope
- (3) TV Camera and Monitor
- (4) Detector Circuit
- (5) Optical Probe Driver.

1. Light Source Assembly

The light source consists of a laser diode, a protection circuit, and an objective lens used as a collimator, all mounted in two concentric aluminum cylinders which are machined to allow easy

insertion over the eyepiece of a conventional binocular microscope. The assembly was specifically sized to fit into a Leitz microscope but should be compatible with most manufacturers' models of similar type. The aluminum cylinder and diode mount will be provided by Lincoln Laboratory since they are the only parts which require custom machine work.

2. Microscope

The microscope used for most of the experiments with this light source was a model LABOR-LUX 12 ME made by Leitz. In order to provide a high resolution image for light probe placement at the appropriate node of a transistor, a 50X objective was used for an overall magnification of 500X at the eyepiece. Other models with somewhat different magnifications may be perfectly suitable, however. It is recommended that a binocular instrument be used in order to guarantee that users do not attempt to view the IC directly through the microscope eyepiece. The laser diode radiates in the near infrared and is therefore not visible; eye damage could result from direct exposure to this highly focused spot. Microscopes with a third port for a TV camera would leave one eyepiece open for viewing, whereas a binocular instrument would have both eyepieces occupied, one by the light source and the other by the TV camera.

3. TV Camera and Monitor

The vidicon, a Gaertner Scientific Corporation Model M3000 Series, was chosen because of its light weight which allows mounting directly on the eyepiece of the binocular microscope. Other models may be suitable as long as they meet this requirement. Note that the camera must be sensitive in the IR region in order to permit TV display of the spot position.

4. Detector Circuit

The detector circuit is a transconductance amplifier which converts a measurement of optical probe current to a voltage which can be displayed on an oscilloscope.

5. Optical Probe Driver

The optical probe driver must provide a current pulse of about 25 mA at a voltage of ~ 10 V. A standard pulse generator such as the DATAPULSE 101 Pulse Generator provides a convenient means for providing such a pulse with the additional convenience of variable pulse width and delay. The electrical and optical characteristics of the laser diode used are shown in Figure 5.

C. TEST VECTOR GENERATION

A new computer program called BANDITS (Boolean Analysis of Digital Timeless Systems) has been introduced and made available for general use. This program accepts the gate-level description of a logic circuit and produces a reduced set of input patterns that will drive the circuit's output(s) to the logic 0 and 1 states. BANDITS is not a Boolean minimization program. However, based on a circuit's given description, it generates input patterns with minimal constraints (i.e., containing as

PARAMETER	SYMBOL	VALUE	UNIT	CONDITION
THRESHOLD CURRENT	I_{th}	13	mA	CW
OPERATING CURRENT	I_{opr}	22	mA	$p = 3$ mW, CW
OPERATING VOLTAGE	V_{opr}	1.66	V	$p = 3$ mW, CW
MONITORING OUTPUT	I_m	0.57	mA	$p = 3$ mW, CW
WAVELENGTH	λ_L	834	nm	$p = 3$ mW, CW

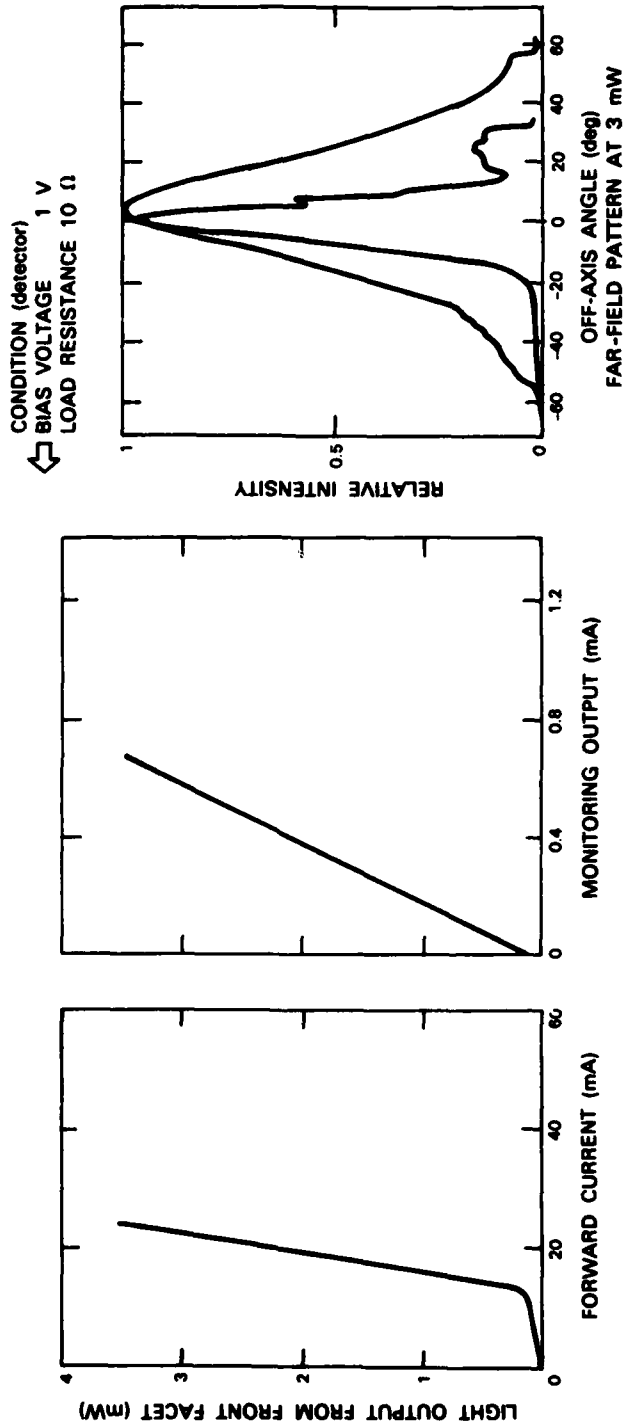


Figure 5. Electrical and optical characteristics of a laser source.

140611-N

many unassigned input variables as possible) which will exercise the circuit's output(s) in both the logic 0 and 1 states. Whereas it accepts only gate type (i.e., non-memory) circuit elements, BANDITS will automatically detect any feedback paths in a circuit and obtain the correct solutions for the output nodes, in terms of the circuit's primary inputs and feedback variables. The program generates excitation patterns for solving a network in the steady state (hence, the term Timeless Systems). It will ignore the effects of propagation delays on a network's response to input stimuli as well as rejecting patterns that create oscillatory conditions within the network.

BANDITS can solve circuits with up to 2000 nodes (gates) and having up to 30 primary inputs and/or state variables. In general, execution time does not depend upon the number of input/state variables since the program operates on all input/state variables in parallel, using word manipulation techniques. Typically, the program generates around 60 patterns per second on a VAX 11/780 computer. For example, a total of 4120, 14-bit patterns are computed in 72 s for the TI 74181 ALU chip which has 14 inputs and 8 outputs. The total number of patterns necessary to drive each of the 8 outputs individually to the logic 0 and logic 1 states is 2568 (the remaining $4120 - 2568 = 1552$ patterns control the internal nodes).

BANDITS uses the well-known cross-product and union operations to compute the necessary patterns. Initially, an L-dimensional vector is assigned to each input node I_i . This vector contains "don't care" (-) values in all bit positions except position i , which contains either a 0 or 1, depending upon whether it is meant to be the pattern for setting that node to logic 0 or 1, respectively. Then, the control patterns for an AND gate with two primary inputs feeding it are computed using

AND-at-0 \equiv Union of the 0-vectors at its inputs,
AND-at-1 \equiv Cross-product of the 1-vectors at its inputs.

Patterns for the other gate types are computed in a similar manner, using cross-product/union and 0-vectors/1-vectors, as appropriate for that gate's type. Techniques used in implementing the necessary operation are many and too complex to include in this report.

The program may be used for the following purposes:

- (1) Obtain, in a single pass, the results of simulating a circuit under all possible input/state combinations.
- (2) Given two different versions of the same circuit, verify that they implement the same function and/or find cases where they differ. To achieve this, both versions should be described so that they share the same set of primary input terminals but retain their individual outputs. Then, tie the corresponding outputs from each circuit to EQUIVALENCE gates and use BANDITS to generate patterns that will drive the EQUIVALENCE gate outputs to logic 0. Any such pattern found indicates a case where the two circuits produce different output values.
- (3) Check a given circuit to see if certain conditions always hold true. For example, if setting certain inputs to some values should prevent an output

node from being set to some logic value, tie the appropriate inputs to the logic 0 and 1 states and use BANDITS to generate the controlling patterns for the output node.

- (4) Perform design verification between a gate-level description of a circuit and its high-level description by simulating the high-level description with the pattern computed using BANDITS. Note that this requires a three-valued (0, 1, unknown) simulator.

A user's manual for internal use provides instructions for using BANDITS and is separately available.

D. TESTING STRATEGIES

Testing of complex VLSI components/systems requires a radically different approach than those presently used for testing MSI/LSI-based digital systems. Indeed, the "difference" must be more than just a clever technique that enables existing (or improved) automatic test pattern generation (ATPG) algorithms to perform better. That is, our focus should be on not only improving algorithm efficiency, but also on improving the testability of VLSI designs through changes in their implementation.

A very good start has been made in this direction by the introduction of Level Sensitive Scan Design (LSSD) rules. Currently, several major digital systems manufacturers are using (variations of) these rules in their designs. Using the LSSD rules enables the designers to eliminate many potential timing problems and makes it possible to implement a "scan path" whereby each and every individual bi-stable element in the circuit becomes separately controllable and observable. Generically, this is achieved by configuring the device-under-test (DUT) such that, for testing purposes, all of its latches become part of a single shift register, called the scan register. The serial data input and output terminals of the scan register are made accessible from two of the external I/O pins of the device. Then, any combination of bit values can be loaded into the scan latches by serially shifting the desired combination into the scan register. The values stored in the scan latches act as input patterns to the combinational part of the DUT. During testing, first the desired bit pattern is shifted into the scan register. Next, the combinational circuit outputs are latched (in parallel) into the scan register. Finally, as the next bit pattern is being shifted in, the results of the previous test become available at the output of the scan register, one bit at a time. This technique, which is commonly referred to as the scan-set technique, reduces the problem of testing a complex digital system to that of testing only the combinational part of its circuitry. However, apart from increasing the controllability/observability of the internal nodes of a system, scan-set does not offer a new approach to the ATPG problem.

Given that modern VLSI systems are capable of operating at very high clock rates, a natural extension of the scan-set approach is to drop the ATPG altogether and exercise the combinational part of a digital system exhaustively. Despite the potentially very large size of such combinational circuits, exhaustive testing appears to be feasible. This can be seen by observing that a multi-input multi-output combinational circuit consists of multiple single-output circuits, each of which

may receive inputs from only a subset of the bits of the scan register. For example, a scan register may have several thousand bits but any single-output logic cone may use only 30 of these as its inputs. Then, if a new input test pattern can be generated at every clock period, applying all 2^{30} -bit permutations to such a logic cone would take less than 2 min., if the clock rate is 10 MHz. However, to achieve this, it is necessary that test results (i.e., combinational circuit output values) should not be latched back into the scan register as this prevents generation of a new test pattern with each clock. Instead, utilization can be made of a separate "signature" register where test results can be accumulated.

Consider the case where the scan register has m -bits feeding an m -input/ m -output combinational circuit and assume that the combinational circuit consists of m -many single output logic cones, the largest of which has t -inputs from the scan register. Then, exhaustive test patterns can be generated using an n -bit ($n \geq t$) prime linear feedback shift register (LFSR) whose output is shifted along the scan register. In this case, the following properties can be stated:

- (1) Any logic cone whose t -inputs fall within some consecutive n -bits of the scan register will be exhaustively tested.
- (2) If the t -inputs to some logic cone do not all lie within some n consecutive bits of the scan register, then the probability that the given cone will not be exhaustively tested depends on " $n - t$." This probability rapidly gets smaller and becomes < 0.03 when $n - t \geq 5$. Furthermore, even more favorable results can be expected if several LFSRs are used to generate the test patterns in such a fashion that after LFSR _{i} has produced its $2^n - 1$ patterns, we change to using LFSR _{$i+1$} .

Test results can be collected from the m -outputs of the combinational circuit by loading these outputs into an m -bit parallel input signature register. In this case, the signature register is implemented using an m -bit LFSR such that the next state of the signature register is determined by the EXCLUSIVE-OR of the next state of the LFSR and the m -bit combinational circuit outputs. This way, fault detection becomes possible by comparing the final state (signature) of the signature register with its expected value, which can be computed via "good network" simulation or experimentation.

It can be shown that the probability of some errors being masked by the above described signature mechanism is 2^{-m} , which for practical values of m , can be considered as insignificant.

Analyzing the expected behavior of the signature register reveals that if this register is provided with m -bit input patterns with a uniform distribution over a period of 4×2^m shift cycles, it will visit 98 percent of its total 2^m states. This implies that the signature register itself may be usable as an input source for generating (almost exhaustive) input test patterns. Whereas m may be too large to allow 4×2^m to be selected as the length of a test sequence, the distribution of the states visited by the signature register can be shown to be uniform. Thus, any subset of t -many bit positions will go through 98 percent of 2^t possible permutations if the test length is chosen to be 4×2^t . This important result enables the removal of the signature register so that combinational circuit outputs can be latched back into the scan register. Then, the current value of the scan register is used as a

test pattern which exercises the combinational circuit. The results of this test are stored back in the scan register after they have been EXCLUSIVE-OR'ed (i.e., hashed) with the current contents of that register. The resulting signature is then used as the next test pattern, and so on. The test sequence length is chosen to be $L \geq 4 \times 2^t$, where t is the maximum number of inputs that any single output logic cone may receive from the scan register. However, in order to improve the probability that the combined scan/signature register approach might work as expected, it is necessary to randomize the combinational circuit outputs by passing these through a hashing circuit. The hashing circuit serves to decrease the correlations between the m -bit combinational circuit inputs and the m -bit outputs. Experiments have confirmed that indeed 98 percent coverage is achieved at the inputs of any t -input logic cone when the combinational circuit outputs are randomized before being fed into the scan register. A more detailed description of the results presented here can be found in Reference 6.

REFERENCES

1. K.W. Crouch, "L5 User's Guide," Project Report RVLSI-5, Lincoln Laboratory, M.I.T. (7 March 1984).
2. Restructurable VLSI Program Semiannual Technical Summary, Lincoln Laboratory, M.I.T. (30 September 1983), DTIC AD-A139959.
3. J.F. Freidin, "RWED Reference Manual," Project Report RVLSI-6, Lincoln Laboratory, M.I.T. (2 March 1984).
4. H. Murveit, "An Integrated Circuit Based Speech-Recognition System," Ph.D. Thesis, University of California at Berkeley, 14 November 1983.
5. C.S. Myers and L.R. Rabiner, "A Level Building Dynamic Time Warping Algorithm for Connected Word Recognition," IEEE Trans. Acoust., Speech, and Signal Processing ASSP-29, 284-297 (April 1983).
6. B. Dervisoglu, "VLSI Self-Testing Using Exhaustive Bit Patterns," Technical Report 685, Lincoln Laboratory, M.I.T. (to be published).

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER ESD-TR-84-021	2. GOVT ACCESSION NO. AD-A148145	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Restructurable VLSI Program		5. TYPE OF REPORT & PERIOD COVERED Semiannual Technical Summary 1 October 1983 — 31 March 1984
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Peter E. Blankenship		8. CONTRACT OR GRANT NUMBER(s) F19628-85-C-0002
9. PERFORMING ORGANIZATION NAME AND ADDRESS Lincoln Laboratory, M.I.T. P.O. Box 73 Lexington, MA 02173-0073		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS ARPA Order 3797 Program Element No. 61101E Project No. 3D30
11. CONTROLLING OFFICE NAME AND ADDRESS Defense Advanced Research Projects Agency 1400 Wilson Boulevard Arlington, VA 22209		12. REPORT DATE 31 March 1984
		13. NUMBER OF PAGES 36
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Electronic Systems Division Hanscom AFB, MA 01731		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES None		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)		
VLSI Restructurable VLSI (RVLSI) programmable interconnect defect avoidance	customization hardware description language placement routing	systemic array integrator wafer-scale systems speech recognition
20. ABSTRACT (Continue on reverse side if necessary and identify by block number)		
<p align="center">This report describes work performed on the Restructurable VLSI Program sponsored by the Information Processing Techniques Office of the Defense Advanced Research Projects Agency during the semiannual period 1 October 1983 — 31 March 1984.</p>		