


INTRODUCTION TO VLSI SYSTEMS


CARVER A. MEAD

*Professor of Computer Science,
Electrical Engineering,
and Applied Physics,
California Institute of Technology*

LYNN A. CONWAY

*Manager, LSI Systems Area,
Palo Alto Research Center,
Xerox Corporation*






986

© CA TECH 1977

80

+



PREFACE:


As a result of improvements in fabrication technology, Large Scale Integrated (LSI) electronic circuitry has become so dense that a single silicon LSI chip may contain tens of thousands of transistors. Many LSI chips, such as microprocessors, now consist of multiple, complex subsystems, and thus are really *integrated systems* rather than integrated circuits.

What we have seen so far is only the beginning. Achievable circuit density now approximately doubles with each passing year. How long can this continue, and how small can the transistor be made? From the physics we find that the linear dimensions of transistors can be reduced to less than 1/10 of those in current integrated systems before they cease to function as the sort of switching elements from which we can easily build digital systems. It will eventually be possible to fabricate chips with hundreds of times as many components as today's. The transistors in such very large scale integrated (VLSI) systems will ultimately have linear dimensions smaller than the wave-length of visible light. New, non-optical, high-resolution lithographic techniques are now being developed by many firms to enable fabrication of such circuitry.

The emerging VLSI presents a challenge not only to those involved in development of appropriate fabrication technology, but also to computer scientists and computer architects. The ways in which digital systems are structured, the procedures used to design them, the tradeoffs between hardware and software, and the methodology and metrics of analysis of algorithms will all be greatly affected by the maturation of electronic technology towards its maximum density. We believe this will be an important area of activity for computer science on through the 1980's.

Until recently, the design of integrated electronic circuitry has been largely the province of circuit and logic designers working within semiconductor firms. Computer system architects have traditionally built systems from standard integrated circuits designed and manufactured by these firms, but haven't often participated in the specification and design of these integrated circuits. Electrical Engineering and Computer Science (EE/CS) curricula have reflected this tradition, with courses in device physics and integrated circuit design aimed at and generally taken by different students than those interested in digital system architecture and computer science.

This text is written to fill a current gap in the literature, and provide students of computer science and electrical engineering with an introduction to integrated system architecture and design. Combined with individual study in related research areas and participation in actual system design projects, the text could serve as a basis for a graduate course sequence in integrated systems. Portions could be used for an undergraduate text on the subject, or to augment a graduate course on computer architecture. It could also be used to extend, in the system direction, a classical electrical engineering course in integrated circuits. We assume the reader's background contains the equivalent of an introductory course sequence in computer science, and introductory courses in electronic circuits and digital design.



Up till now there have been major obstacles in the path of those attempting to gain an overall understanding of integrated systems. Integrated electronics, developing in a heatedly competitive business environment, has proliferated into a large array of different device technologies, circuit design families, logic design techniques, maskmaking and wafer fabrication techniques, etc. The technologies have sprung up from the grass roots of "Silicon Valley" in California. Most participants in the industry have of necessity concentrated on rather narrow specialties. Texts on the subject have tended to give detailed accounts of some very narrow horizontal segment of the overall subject, such as device physics or circuit design.

We have chosen instead to provide the minimum of basic information about devices, circuits, fabrication technology, logic design techniques, and system architecture, which is sufficient to enable the reader to span fully the entire range of abstractions, from the underlying physics to complete VLSI digital computer systems. A rather small set of carefully selected key concepts is all that is necessary for this purpose. We believe that only by carrying along the least amount of unnecessary mental baggage at each step in such a study, will the student emerge with a good overall understanding of the subject. Once this range of abstractions is spanned, the sequence of concepts can then be mapped into the reader's own space of application and technology.

Another major obstacle has been the high rate of change of integrated electronics. The uninitiated could easily get the feeling that much energy could be invested in learning material which becomes obsolete as rapidly as it is assimilated. The major driving mechanism in all this change is the continual improvement in fabrication technology. This evolutionary process results in the feasibility of manufacturing smaller and smaller devices as time passes. By including the effects of scaling down device dimensions as an essential ingredient of all topics in this text, many of the important changes of the architectural parameters of the technology are predicted, expected, and indeed hoped for.

The key concepts of this text are illustrated by way of specific examples. In any given technology, form follows function in a particular way. The skill of mapping function into form, when once acquired, can be readily applied to any technology. Because of its density, speed, topological properties, and general availability of wafer fabrication, nMOS has been chosen as the technology in which examples are implemented.

An atmosphere of excitement and anticipation pervades this field. Workers from many backgrounds, computer scientists, electrical engineers, and physicists, are collaborating on a common problem area which has not yet become classical. The territory is vast, and largely unexplored. The rewards are great for those who simply press forward.

Carver A. Mead
Pasadena, California

Lynn A. Conway
Belmont, California




BACKGROUND

Prior to the commercial publication of this textbook, this limited printing is being distributed to a selected group of universities as course notes for graduate courses on integrated systems, for the purpose of obtaining critical reviews. Copies are also being distributed to the industrial participants in the Caltech Silicon Structures Project, and to selected individuals in universities and in industry for their review. The authors welcome any and all comments and suggestions from readers. We are especially interested in hearing of the experiences of those teaching from the text. Notifications and corrections of errors, ideas for improvements in the tutorial techniques used, and suggestions of instructive problems for each chapter would be greatly appreciated.

While the material in this text is presented in a particular order, it need not be read in that order. Each chapter presents material from a distinct level in the hierarchy of disciplines involved in integrated systems. The material falls into four major groupings: chapters 1 and 2 provide the basics of devices and fabrication, chapters 3 and 4 give the basics of design and implementation, chapters 5 and 6 present an example of LSI system design, and chapters 7, 8, and 9 discuss topics of current interest in integrated systems research. We recommend that the reader start in the chapter where he or she is most knowledgeable, and read until information is required from an adjacent area described in some other chapter. By using this algorithm and consulting the reading references where necessary, the reader can gradually work through the primary material of all chapters. Although much of the material in this text is previously unpublished, it contains only fundamental concepts. However, these concepts cover a wide range of disciplines, and are easily visualized only after the overall context becomes clear.

This text has its origins in a series of courses in integrated circuit design given by Carver Mead at Caltech, beginning in 1970. Starting in 1971, students in these courses designed and debugged their own integrated circuits. The students undertook increasingly complex system designs, using only rather simple implementation aids. The structured design methodology presented in this text evolved within this milieu. A separate Computer Science department was formed at Caltech in 1976, with integrated systems as a focus. An early association was formed with systems architects in industry. Interaction of Caltech students and faculty with industrial researchers resulted in the idea of a joint Caltech-industry cooperative program. The initial industrial participants in this program, known as the Caltech Silicon Structures Project, are Intel, DEC, IBM, Hewlett-Packard, Xerox, and Burroughs.

Work on this text began in August 1977. The first three chapters were used as course notes during the fall of 1977, in courses given by Carver Mead at Caltech and by Carlo Sequin at U. C. Berkeley. The first five chapters were used during the spring of 1978 in courses given by Ivan Sutherland at Caltech, by Robert Sproull at Carnegie-Mellon University, and by Fred Rosenberger at Washington University, St. Louis. This printing is being used again during the fall of 1978 in the courses at Caltech and U.C. Berkeley, and in new courses by Lynn Conway, while visiting at M.I.T., by Kent Smith at University of Utah, and by Bob Bower at U.C.L.A.



ACKNOWLEDGEMENTS

We wish to express our gratitude to the many individuals who have contributed ideas, time, and energy towards the creation of this textbook. In particular we wish to thank the following:

For contributions to the text:

Chuck Seitz, California Institute of Technology, for contributing chapter 7, *System Timing*, to the text; Martin Rem, Eindhoven University of Technology, and Sally Browning, Caltech, for their major contributions to chapter 8; H. T. Kung and Charles Leiserson, Carnegie-Mellon University, for permission to reprint their original, copyrighted work, *Algorithms for VLSI Processor Arrays*, as section 2 of chapter 8; David Johannsen, Caltech, for the OM2 Specification section, and for his assistance in the preparation chapter 6; Robert Sproull, Carnegie-Mellon University, and Richard Lyon, Xerox PARC, for the CIF section; Carlo Sequin, U. C. Berkeley, for his contributions to chapter 1, his detailed review of the text, and his many suggestions for improvements; Robert Sproull and Wayne Wilner, for their assistance in the preparation of chapter 8; Hank Smith, M.I.T. Lincoln Laboratory, for his contributions of ideas for the section on x-ray lithography; Douglas Fairbairn, Xerox PARC, and James Rowson, Caltech, for the ICARUS section; Dale Green and his group at Xerox EOS for technical assistance in preparation of the chapter 5 color plates; California Institute of Technology, for permission to reprint the OM2 layout figures and OM2 specifications; Barbara Baird of Xerox PARC, for her efforts in overseeing this printing of the text.

For reviewing drafts:

Robert Sproull; Ivan Sutherland and Chuck Seitz, Caltech; Fred Rosenberger and Charles Molnar, Washington Univ.; Richard Lyon, Chuck Thacker, Wayne Wilner, and Alan Kay, Xerox.

For participation in the OM projects:

Dave Johannsen, Mike Tolle, Chris Carroll, Rod Masumoto, Ivan Sutherland, and Chuck Seitz.

For helping define the starting frame and establish the multi-project chip capability:

Carl Pompei, NBK; Joel Sorem, Micro Mask; Vir Dhaka, Doug Fairbairn, Bill Winfield, Xerox; Jim Rowson, Ron Ayres, Steve Trimberger, Caltech; Bob Hon, C.M.U.; Bob Baldwin, M.I.T.; Dan Izumi and George Marr, Maruman Integrated Circuits; Bill Lattin and Ted Jenkins, Intel.

For support of integrated systems research at Caltech:

The Office of Naval Research; Robert Noyce and Gordon Moore, Intel Corporation.

We are especially grateful to W. R. Sutherland, Manager, Systems Science Laboratory of Xerox PARC, for his encouragement and support, and to George E. Pake, Vice-President, Xerox Corporate Research, who created the environment which made this book possible.




TABLE OF CONTENTS

Chapter 1: MOS Devices and Circuits

The MOS Transistor - - - The Basic Inverter - - - Inverter Delay - - - Parasitic Effects - - - Driving Large Capacitive Loads - - - Space vs Time - - - Basic NAND and NOR Logic Circuits - - - Super Buffers - - - A Closer Look at the Electrical Parameters - - - Depletion Mode vs Enhancement Mode Pullups - - - Delays in Another Form of Logic Circuitry - - - Pullup/Pulldown Ratios for Inverting Logic Coupled by Pass Transistors - - - Transit Times and Clock Periods - - - Properties of Cross Coupled Circuits - - - A Fluid Model for Visualizing MOS Transistor Behavior - - - Effects of Scaling Down the Dimensions of MOS Circuits and Systems

Chapter 2: Integrated System Fabrication

Patterning - - - Scaling of Patterning Technology - - - The Silicon Gate n-Channel Process - - - Yield Statistics - - - Scaling of the Processing Technology - - - Design Rules - - - Formal Description of Design Rules - - - Electrical Parameters - - - Current Limitations in Conductors - - - A Closer Look at Some Details - - - Choice of Technology

Chapter 3: Data and Control Flow in Systematic Structures

Notation - - - Two Phase Clocks - - - The Shift Register - - - Relating Different Levels of Abstraction - - - Implementing Dynamic Registers - - - Designing a Subsystem - - - Register to Register Transfer - - - Combinational Logic - - - The Programmable Logic Array - - - Finite State Machines - - - Towards a Structured Design Methodology

Chapter 4: Implementing Integrated System Designs

Patterning and Fabrication - - - Hand Layout and Digitization using a Symbolic Layout Language - - - An Interactive Layout System - - - The Caltech Intermediate Form - - - The Multi-Project Chip - - - Examples - - - Patterning and Fabrication in the Future - - - Fully Integrated, Interactive Design Systems - - - System Simulation, Test Generation, and Testing

Chapter 5: Overview of an LSI Computer System, and Design of the OM2 Data Path Chip

The OM Project at Caltech - - - System Overview - - - The Overall Structure of the Data Path - - - The Arithmetic Logic Unit - - - ALU Registers - - - Buses - - - Barrel Shifter - - - Register Array - - - Communication with the Outside World - - - Data Path Operation Encoding - - - Functional Specification of the OM2 Data Path Chip

Chapter 6: Architecture and Design of System Controllers, and the Design of the OM2 Controller Chip

Alternative Control Structures - - - The Stored Program Machine - - - Microprogrammed Control - - - Design of the OM2 Controller Chip - - - Examples of Controller Operation - - - Some Reflections on the Classical Stored Program Machine

Chapter 7: System Timing , by Charles L. Seitz

The Third Dimension - - - Synchronous Systems - - - Clock Distribution - - - Clock Generation - - - Synchronization Failure - - - Self-Timed Systems - - - Signaling Conventions - - - Synchronous Elements - - - Asynchronous Elements - - - Arbitration

Chapter 8: Highly Concurrent Systems

Introduction - - - Communication and Concurrency in Conventional Computers - - - Algorithms for VLSI Processor Arrays - - - Hierarchically Organized Machines - - - Highly Concurrent Systems with Global Communications - - - Challenges for the Future

Chapter 9: Physics of Computational Systems

Thermodynamic View of Computation - - - Energetics of Bistable Devices - - - Thermal Limit - - - Quantum Limits - - - Granularity of Charge - - - Voltage Limits - - - An Example - - - Energy Management - - - Discreteness in Quantum Mechanical Systems - - - Conclusion

