INTRODUCTION TO VLSI SYSTEMS

CARVER MEAD • LYNN CONWAY

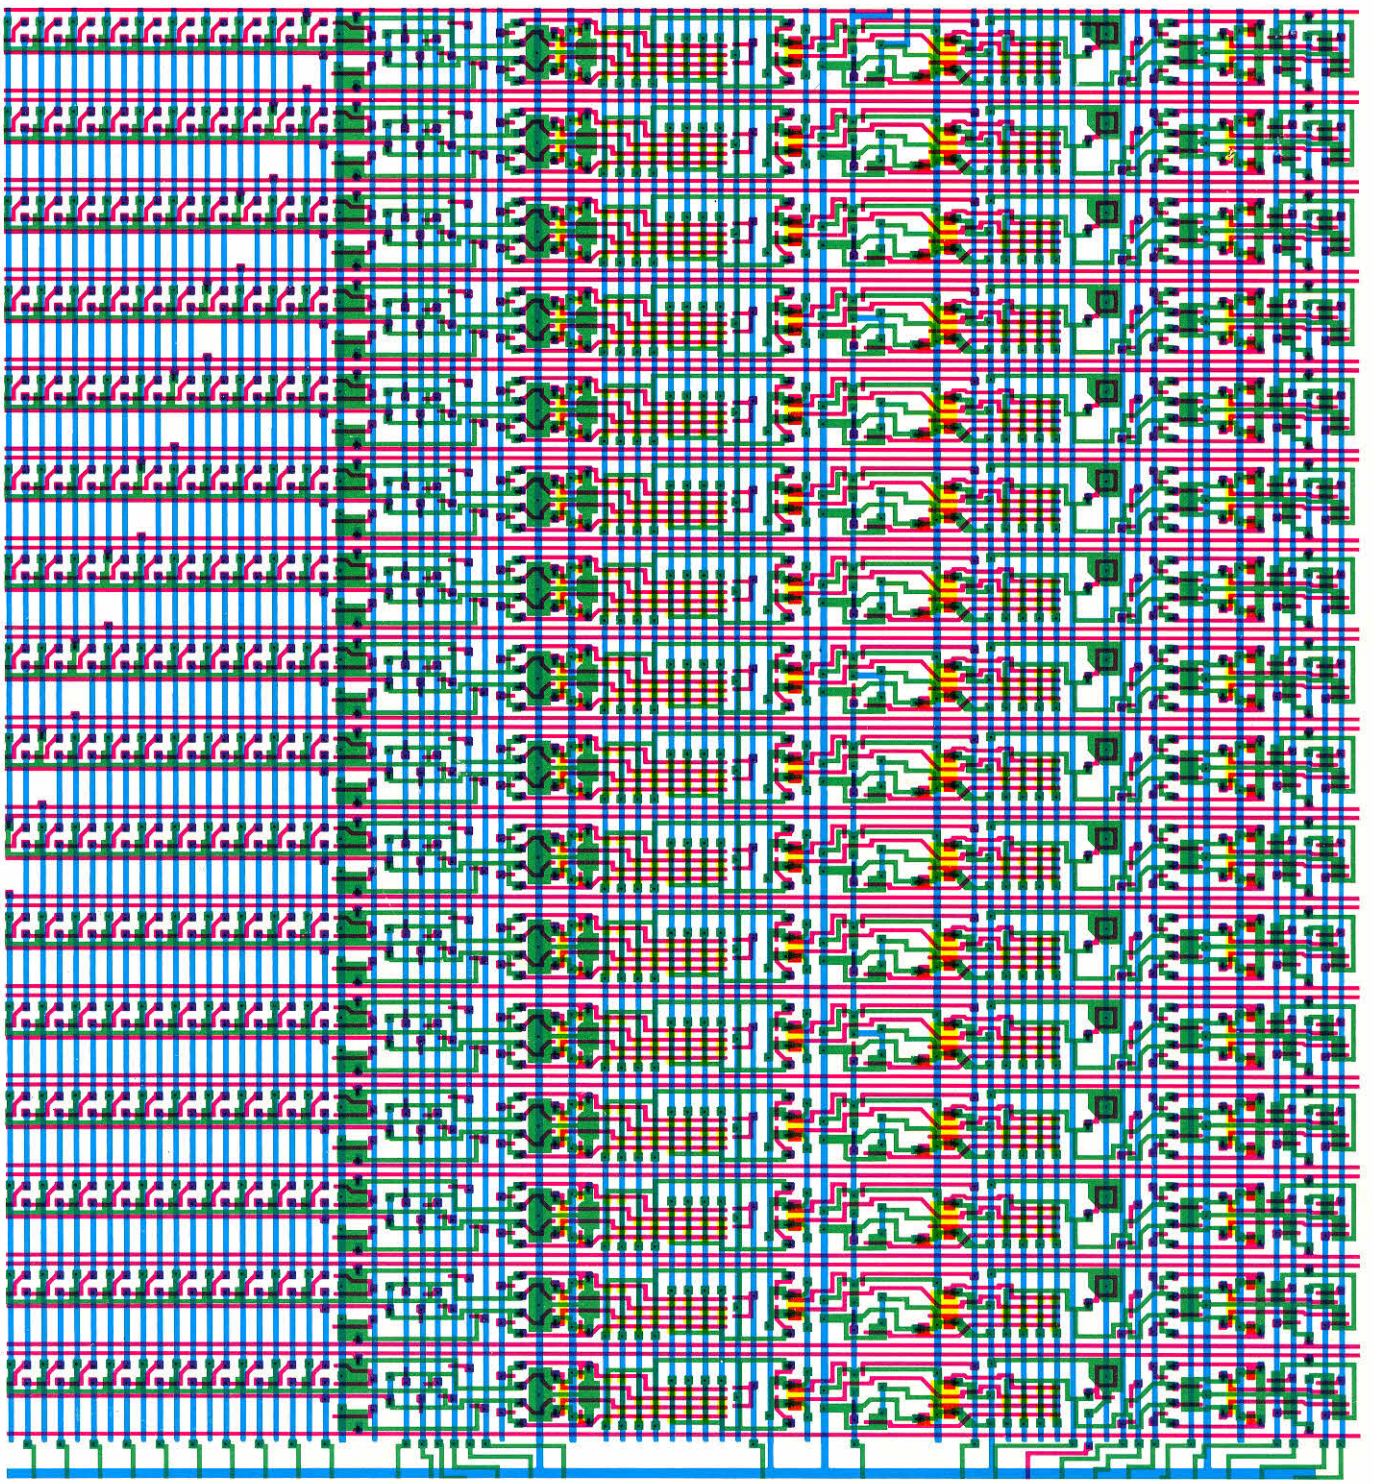

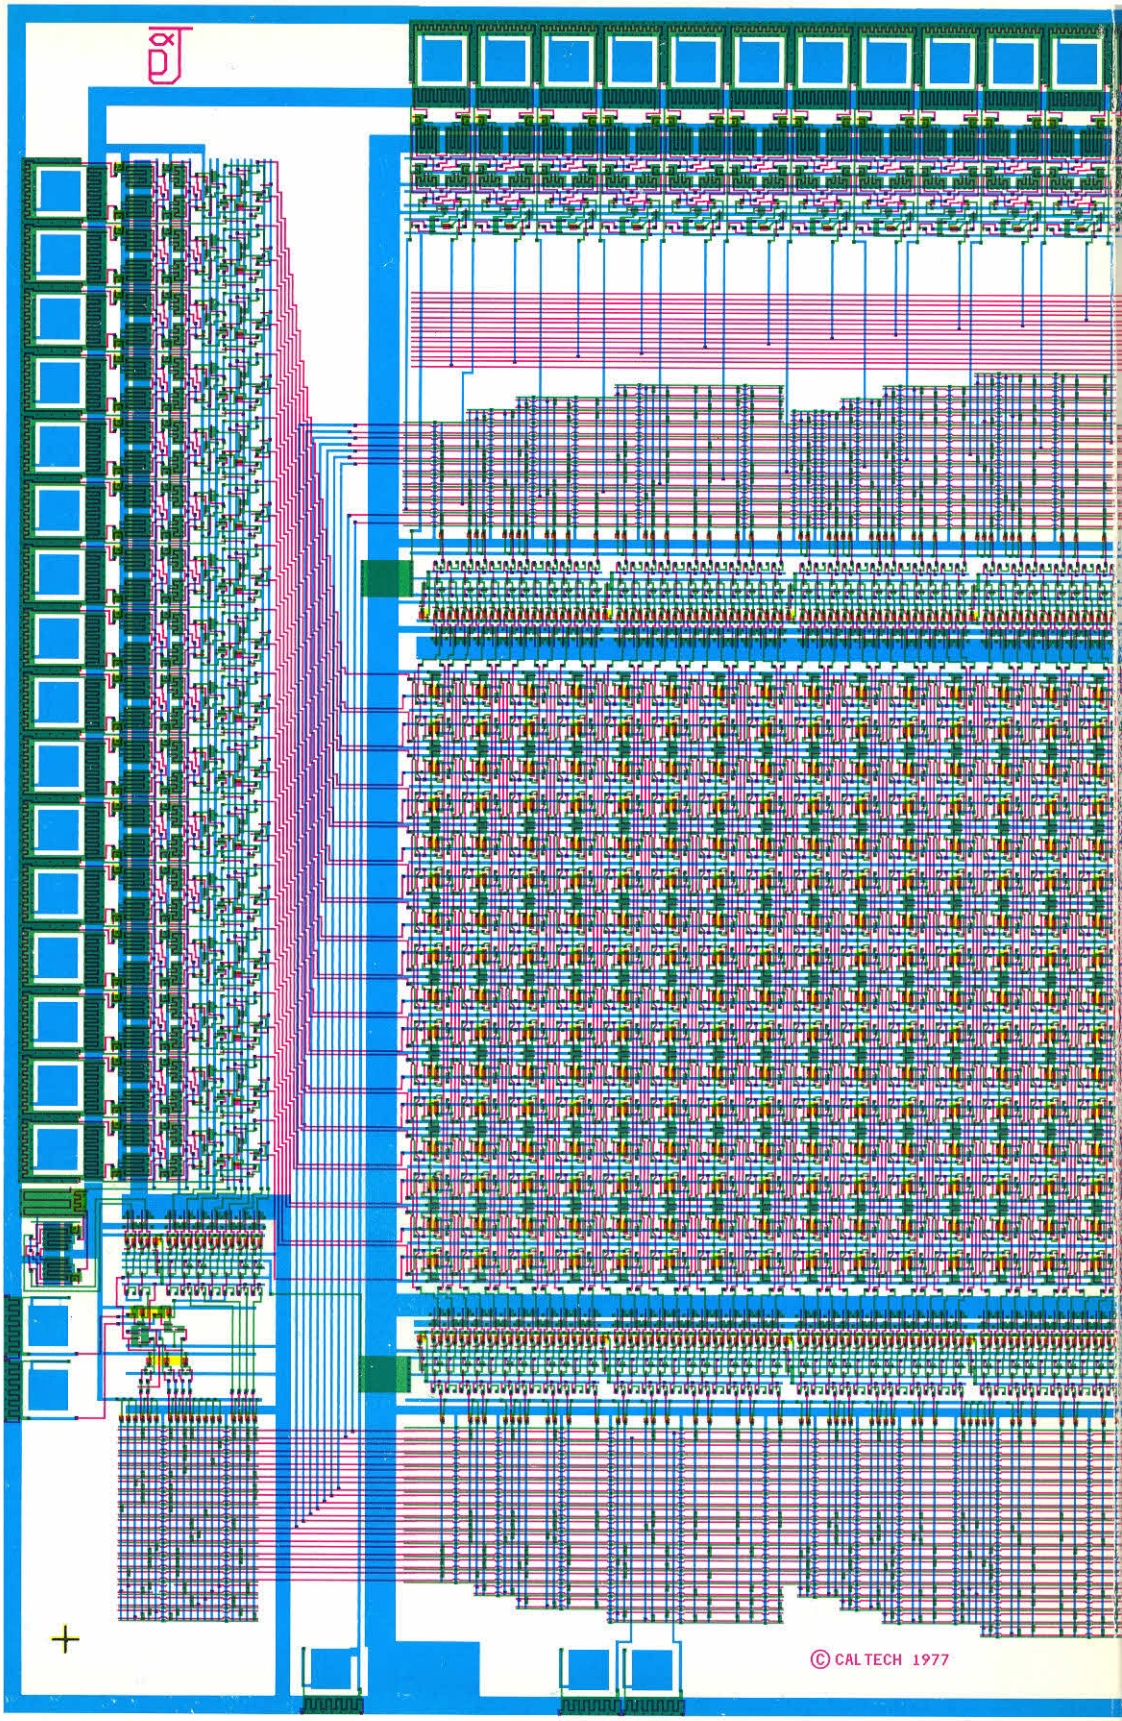
© CALTECH 1977

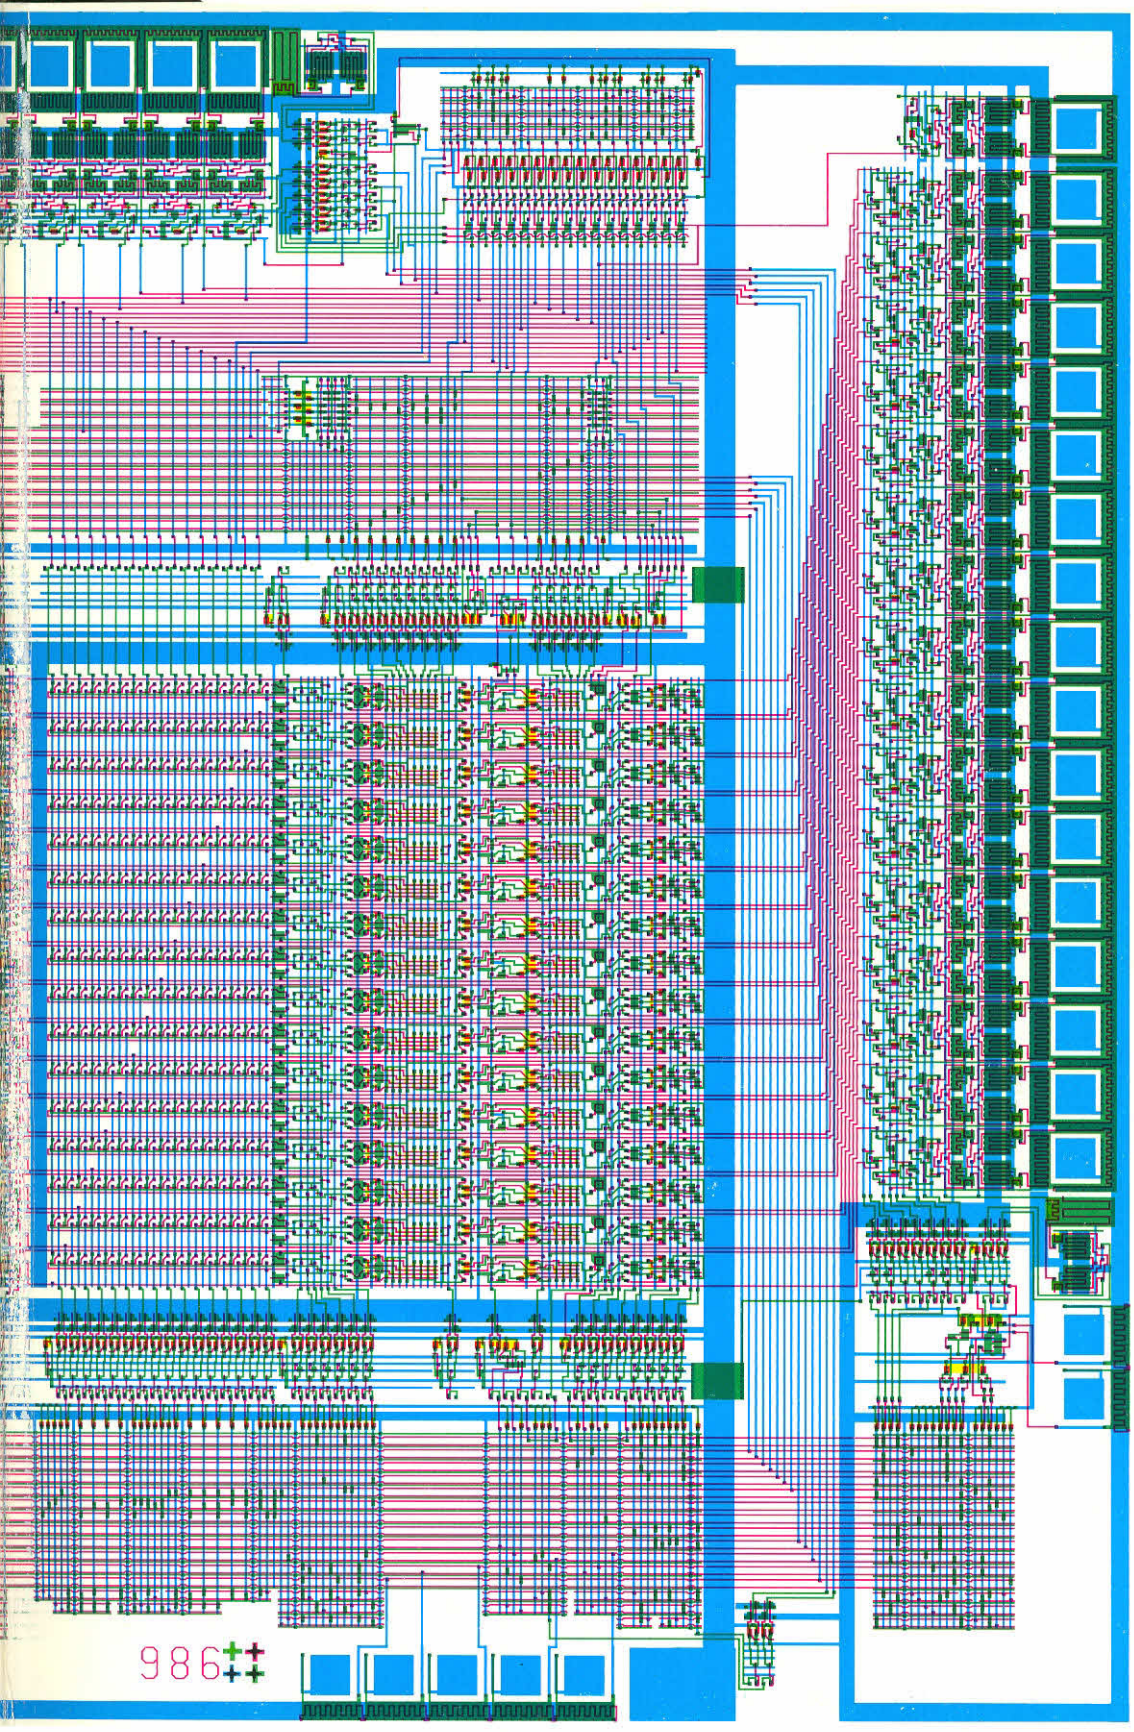
986

INTRODUCTION TO VLSI SYSTEMS

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

LYNN CONWAY
Research Fellow, and
Manager, VLSI System Design Area
Palo Alto Research Center, Xerox Corporation

ADDISON-WESLEY PUBLISHING COMPANY

Reading, Massachusetts • Menlo Park, California
London • Amsterdam • Don Mills, Ontario • Sydney

This book is in the
Addison-Wesley Series in Computer Science

Consulting Editor
Michael A. Harrison

Library of Congress Cataloging in Publication Data

Mead, Carver A
Introduction to VLSI systems.

1. Integrated circuits—Large scale integration.
2. Microcomputers. 3. Digital electronics.
4. Computer architecture. I. Conway, Lynn A., joint author. II. Title.
TK7874.M37 621.3819'535 78-74688
ISBN 0-201-04358-0

Second printing, October 1980

ISBN 0-201-04358-0
ABCDEFGHIJK-HA-89876543210

TO W. R. SUTHERLAND

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 doubles with each passing year or two. Physical principles indicate that transistors can be scaled down to less than 1/100th of their present area and still function as the sort of switching elements with which we can build digital systems. By the late 1980s it will be possible to fabricate chips containing millions of transistors. The devices and interconnections in such very large scale integrated (VLSI) systems will have linear dimensions smaller than the wavelength of visible light. New high-resolution lithographic techniques have already been demonstrated that will enable fabrication of such circuitry.

VLSI electronics presents a challenge, not only to those involved in the development of 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 trade-offs between hardware and software, and the design of computational algorithms will all be greatly affected by the coming changes in integrated electronics. We believe this will be a major area of activity in computer science on through the 1980s.

Until recently the design of integrated circuitry has been the province of circuit and logic designers working within semiconductor firms. Computer architects have traditionally composed systems from standard integrated circuits designed and manufactured by these firms but have seldom participated in the specification and design of these circuits. Electrical Engineering and Computer Science (EE/CS) curricula reflect this tradition, with courses in device physics and integrated circuit design aimed at a different group of students than those interested in digital system architecture and computer science.

This text is written to fill a current gap in the literature and to introduce all EE/CS students to integrated system architecture and design. Combined with individual study in related research areas and participation in large system design projects, this text provides the basis for a graduate course-sequence in integrated systems. However, it is primarily intended for use in intensive undergraduate courses on the subject. The material can also be used to augment courses on computer architecture. We assume the reader's background contains the equivalent of introductory courses in computer science, electronic circuits, and digital design.

There have been major obstacles in the way of those seeking an overall understanding of integrated systems. Integrated electronics has developed in a heatedly competitive and often secretive business environment. There has been a proliferation of different device technologies, circuit design families, logic design techniques, maskmaking and wafer fabrication techniques, etc. Many of these technologies have sprung up from the grass roots of "Silicon Valley" in the San Francisco Bay Area of California, and thus many of the "experts" are located in that one region. Most workers in the industry have concentrated on narrow specialties. Separate integrated electronics cultures have independently evolved within many companies, and thus the terminology and practices of the specialties vary from company to company.

As a result of this background, texts on integrated electronics have tended to give detailed accounts of some very narrow horizontal segment of the overall subject, such as device physics or circuit design, and are often tied in subtle ways to some specific context, thus limiting their general applicability.

We have chosen instead to provide just enough essential information about devices, circuits, fabrication technology, logic design techniques, and system architecture to enable the reader to fully span the entire range of abstractions from the underlying physics to complete VLSI digital computer systems. A rather small set of key concepts is sufficient. Only by learning the essence of each topic, and by carrying along the least amount of mental baggage at each step, will the student emerge with a good overall understanding of the subject. This understanding can then be mapped into the reader's own space of application, technology, and technical culture.

The high rate of change of integrated electronics presents another obstacle: information often becomes obsolete very rapidly. The major force for obsolescence is the ongoing improvement in fabrication technology, leading to smaller and smaller devices as time passes and thus to a constant change in device characteristics. We attack this obstacle by stressing the effects of the scaling-down of device dimensions. Many of the coming changes in system architectural parameters are thus anticipated. The reader will learn what is common to systems composed of 6 μm, 2 μm, and 0.5 μm devices, and what is not.

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, circuits, and fabrication; Chapters 3 and 4 give the basics of system design and implementation; Chapters 5 and 6 present an example of LSI system design; topics of current research interest are discussed in Chapters 7, 8, and 9. We recommend that readers start in the chapter where they are most knowledgeable, and read until information is required from an adjacent area described in some other chapter. By using this algorithm and consulting the suggested references where necessary, readers can gradually work through the primary material of all chapters. Although much of the material in this text is previously unpublished, it nevertheless contains only basic concepts. However, these concepts cover quite a wide range of disciplines and are easily visualized only after the overall context becomes clear.

In any given technology, form follows function in a particular way. The most efficient first step towards understanding the architectural possibilities of a technology is the study of carefully selected existing designs. However, system architecture and design, like any art, can only be learned by doing. Carrying a small design from conception through to successful completion provides the confidence necessary to undertake larger designs. The space of possibilities unfolds only as the medium is worked. This book provides a set of selected design examples and also describes procedures for implementing one's own designs. Because of the density, speed, and topological properties of *n*MOS, and the easy access to *n*MOS wafer fabrication, that technology is used for our examples. The architectural skill of mapping function into form, when once acquired, can then be extended to other technologies.

The general availability of courses in VLSI system design at major universities marks the beginning of a new era in electronics. The rate of systems innovation using this remarkable technology need no longer be limited by the perceptions of a handful of semiconductor companies and large computer manufacturers. New metaphors for computation, new design methodologies, and an abundance of new application areas are already arising within the universities, within many system firms, and within a multitude of new small enterprises. There may never have been a greater opportunity for free enterprise than that presented by these circumstances.

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

Pasadena, California C.M.
Belmont, California L.C.
July 1979

BACKGROUND

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 has evolved within this milieu. These early courses greatly benefited from interactions with friends in industry, particularly Robert Noyce, Gordon Moore, Frederico Faggin, Dov Frohman-Bentchkowsky, Ted Jenkins, and Joel Sorem.

A separate Computer Science activity was created at Caltech in 1976, with integrated systems as a focus. An early, informal association was formed with systems architects in industry, in particular with the then newly formed LSI Systems Area, led by Lynn Conway, at Xerox Palo Alto Research Center (PARC). The increased interaction of Caltech students and faculty with industrial researchers stimulated the research on both sides.

Work on this text began in August 1977. Collaborators from a number of universities and industrial firms joined in the enterprise. Prior to commercial publication, several limited printings were distributed to a selected group of universities as notes for courses on integrated systems. 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 Séquin at U.C. Berkeley. The first five chapters were used during the spring of 1978 in courses given by Ivan Sutherland and Amr Mohsen at Caltech, by Robert Sproull at Carnegie-Mellon University, by Dov Frohman-Bentchkowsky at Hebrew University, Jerusalem, and by Fred Rosenberger at Washington University, St. Louis. The third and final prepublication printing of all nine chapters, in the fall of 1978, was used in the courses at Caltech and U.C. Berkeley, and in new courses by Kent Smith at the University of Utah, and by Lynn Conway, while visiting at M.I.T.

The 1978 M.I.T. course provided a final test, prior to publication of this text, to confirm the transportability of the project-oriented form of the course (in which as much emphasis is placed on creative architectural activity as on formal analysis), and to also confirm the technical and economic feasibility of the remote-entry of student LSI designs to a central facility for fast-turnaround implementation. In the future, M.I.T. will offer this as course 6.371, with Jonathan Allen teaching it in the fall of 1979.

The following information concerning the M.I.T. experience may be useful to those planning similar activities. The course began in mid-September and was attended by 30 students (mostly graduate EE/CS students). Most of the formal lecture material necessary for undertaking projects (covering selected portions of Chapters 1 through 6 of this text) was completed by early November. The students then defined and began work on their LSI design projects. A design cut-off date of December 5 was set, and most designs were completed by that time. The projects included a LISP Microprocessor, a Graphics Memory Subsystem for mirroring and rotating bit-map data, a Writeable PLA project, the Data Path for a Bit-Slice Microprocessor, an LRU Virtual-Memory Paging subsystem, a Bus-Interfaceable Real-Time Clock-Calendar, a Multifunction Smart Memory, several digital signal-processing projects, several subsystems for data-base operations, and many other innovative designs.

Students described their layouts in a simple subset of CIF2.0, using a standard text-editor running on a DEC-20. The only hardware added to the DEC-20 system to support the course were several CRT terminals, two HP four-color pen plotters, and a connection to the local ARPANET host machine. The only software developed were programs for parsing the CIF subset, for instantiating data for plotting, and for driving the plotters. A small library of useful cells, namely input pads with lightning arrestors, output pads with cascaded drivers, and a set of PLA cells, were made available in CIF form. Some students developed their own symbolic layout languages and translaters to CIF, in order to make layout encoding less tedious. By using a structured design methodology the students were able to complete substantial LSI projects in a short period of time, using only primitive design tools. Each project contained on the order of several hundred to several thousand transistors. The logistics of interacting with the large group of student designers to organize the multiproject chip (updating the rules of the game, selecting projects for inclusion, negotiating space allocations, answering individual questions, etc.) were expedited by using the message system on the DEC-20. The corresponding interactions with the remote implementation facility (PARC) were handled via electronic mail.

On December 6, 1978, the individual CIF2.0 design files were transmitted from M.I.T. via the ARPANET to Xerox PARC. This was an additional test by these ARPA contractors of the use of such packet-switching networks for transmitting LSI design files and organizing multiproject chip sets. At PARC all the student designs were merged into a multiproject chip design file, from which

masks were generated using Micro Mask, Inc.'s electron-beam maskmaking facility. Wafers were fabricated at Hewlett-Packard's Deer Creek Laboratory, which cooperated in the feasibility test. The wafers were returned to M.I.T. and electrically characterized by tests of the project-set electrical test patterns. They were then diced and the resulting chips packaged. Packaged chips, with wire-bonding customized for each project, were made available to all students by January 18, 1979. Many of the projects have since undergone thorough functional testing by the students. A number of these function completely correctly. Most exhibited only minor bugs, typically at the logic level of abstraction, of a sort reminiscent of one's first efforts at constructing large programs in a new language.

As a common VLSI system design culture spreads, as higher-level design aids are developed and shared within this culture, and as standard-interface commercially-accessible implementation facilities are established, we will undoubtedly see far more ambitious courses, projects, and research activities undertaken by students and faculty within the universities. And thus the period of exploration begins.

ACKNOWLEDGEMENTS

We wish to express our gratitude to the many individuals who have contributed their ideas, their time, and their energy toward 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 contributions to Chapter 8; David Johannsen, Caltech, for his major contributions to Chapters 5 and 6; H. T. Kung and Charles Leiserson, Carnegie-Mellon University, for permission to reprint their original, copyrighted work, *Algorithms for VLSI Processor Arrays*, as Section 3 of Chapter 8; Robert Sproull, CMU, and Richard Lyon, Xerox PARC, for the CIF section; Carlo Séquin, U.C. Berkeley, for his contributions to Chapter 1, his detailed review of the text, and his many suggestions for improvements; John Best, Chuck Seitz, Jim Kajiya, and Tom McGill, Caltech, and Johan de Kleer, Xerox PARC, for their stimulating discussions and suggestions for Chapter 9; Robert Sproull and Wayne Wilner, for their assistance in the preparation of Chapter 8; Hank Smith, M.I.T., for his contributions to the section on high-resolution lithography; Douglas Fairbairn, Xerox PARC, and James Rowson, Caltech, for the ICARUS section; Dale Green and Liz Bond, Xerox EOS, for technical assistance in the preparation of the text, figures, and color plates; Barbara Baird of Xerox PARC, for overseeing the printing and distribution of the prepublication versions of the text.

For valuable discussions and comments: Robert Sproull; Fred Rosenberger and Charles Molnar, Washington University; Richard Lyon, Doug Fairbairn, Wayne Wilner, and Leo Guibas, Xerox PARC; Chuck Seitz and Ivan Sutherland,

Caltech; Craig Mudge, DEC; Bill Heller, IBM; Gerry Parker, Dick Pashley, and John Wipfli, Intel; Harry Peterson, BNR Ltd.; Al Perlis, Yale University; Wesley Clark.

For participation in the OM projects: Dave Johannsen, Mike Tolle, Chris Carroll, Rod Masumoto, Ivan Sutherland, Chuck Seitz, Danny Cohen, and Leslie Froisland.

For helping to establish the multiproject chip capability: Richard Lyon, Doug Fairbairn, and Alan Bell, Xerox PARC; Bob Hon, CMU; Carlo Séquin, U.C. Berkeley; Ted Jenkins, Intel; Jim Rowson, Ron Ayres, and Steve Trimberger, Caltech.

For contributions to the M.I.T. 1978 VLSI design course: Jonathan Allen, Paul Penfield, Jr., Paul E. Gray, Fernando Corbató, Bill Henke, Glen Miranker, Joy Thompson, Dimitri Antoniadis, Stephen Senturia, Gerald Sussman, Jack Holloway, and Tom Knight, M.I.T.; Robert Noyce and Gordon Moore, Intel Corporation; Merrill Brooksby, and Patricia Castro, Hewlett-Packard; Richard Lyon and Alan Bell, Xerox PARC; and, especially, all the students who participated in course 6.978.

For long-standing support of integrated systems research at Caltech: The Office of Naval Research; Robert Noyce and Gordon Moore, Intel Corporation.

We are very grateful to Xerox Corporation for providing us access to the *Alto-Ethernet-Dover* research-prototype office systems at Xerox PARC and elsewhere within Xerox. Throughout the development of the ideas for this textbook, these personal, distributed computer systems enabled a large, scattered group of collaborators to function as a closely knit research community. These systems also enabled rapid creation and distribution of the prepublication printings of the text, thus helping us obtain valuable feedback from those teaching the early courses in the universities.

We are especially grateful to W. R. Sutherland, Manager, Systems Science Laboratory of Xerox PARC, for providing us with inspiration, guidance, and support; to Ivan Sutherland, Fletcher Jones Professor of Computer Science, Caltech, for his role in starting the Computer Science activities at Caltech, and for establishing links between the computer science community and the emerging activities in integrated systems; to Robert E. Kahn, Chief Scientist and Director for Information Processing Techniques, Defense Advanced Research Projects Agency, for encouraging and supporting integrated systems research in the universities; and to George E. Pake, Vice-President, Xerox Corporate Research, for creating the research environment that made this book possible.

CONTENTS