




VLSI Technology, Inc.



INTRODUCTION TO VLSI SYSTEMS DESIGN



INTRODUCTION TO VLSI SYSTEMS DESIGN

Instructors:

Douglas G. Fairbairn — Director of VLSI Activities
VLSI Technology, Inc.

Robert Mathews — Acting Assistant Professor
Stanford University

John Newkirk — Acting Assistant Professor
Stanford University

VLSI Technology, Inc. reserves the right to make changes in the contents of this book without notice.
VLSI Technology, Inc. assumes no responsibility for any errors that may appear in this document.

Copyright ©1982 by VLSI Technology, Inc. All rights reserved. No part of this publication may be reproduced in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior approval of VLSI Technology, Inc.

INTRODUCTION TO VLSI SYSTEMS DESIGN

INSTRUCTORS:

Douglas Fairbairn is Manager of VLSI Engineering at VLSI Technology, Inc. Mr. Fairbairn left the Xerox Palo Alto Research Center in January 1980 to join VTI. While at Xerox, he participated in a wide variety of systems and IC design projects. The systems projects included an office information system based on high-resolution video terminals and a distributed network of minicomputers. Mr. Fairbairn was also chief architect and project manager of a program which developed a sophisticated portable computer. Since 1976, he has been involved in a variety of activities relating to custom integrated circuit design. These include architecting and laying out two custom integrated circuits and aiding in the design and coding of an interactive graphics program for IC design. Mr. Fairbairn has taught IC design at both Xerox and California Institute of Technology. He is co-founder and publisher of a new integrated circuit design magazine, VLSI Design.

Dr. Robert Mathews and Dr. John Newkirk are Acting Assistant Professors at Stanford University, where they co-teach the VLSI Systems course. They are engaged in VLSI research and the Information Systems Laboratory. Their interests include applying compiling and statistical techniques to VLSI problems and developing algorithms and computational building blocks suitable to VLSI. Dr. Mathews' previous research involved speech processing and psychophysical testing for an auditory prosthesis, while Dr. Newkirk's centered on computational aspects of fast algorithms for system identification. Both have broad programming backgrounds, including experience in industry over the past 15 years.

INTRODUCTION TO VLSI SYSTEMS DESIGN: A VIDEOTAPE COURSE

VLSI design education is a critical issue in an increasing number of high-technology electronics companies. As complete systems move onto single chips, it is vitally important that all those involved in the system design process have intimate familiarity with VLSI design techniques and trade-offs. It was in answer to this need for system engineers, managers, and CAD programmers to gain a first-hand knowledge of VLSI design that VLSI Technology, Inc. created the video-tape course "Introduction to VLSI Systems Design."

Primarily aimed at those who want to become involved in the design process, the course emphasizes the importance of the "learn by doing" approach. There are homework and project assignments which are important parts of the learning process. Students start by designing simple functional blocks such as decoders and parity generators and progress to the design of complete subsystems with complexities on the order of several thousand transistors. Each student or pair of students goes through the complete design cycle from systems architecture to "stick diagrams" (logic design) and final layout.

The importance of developing architectures appropriate to VLSI is very strongly emphasized. Appropriate architectures combined with very careful chip "floor planning" and regularized structures are the key elements of the design methodology presented in this video-tape course. The combination of these techniques help ensure the best trade-offs in system performance, chip size, and design time.

Examples and homework given in the course are based upon nMOS technology. This technology was chosen because it is relatively easy to design with and is widely available. However, most of the basic course materials apply to other MOS technologies as well. Many of the concepts are even applicable to bipolar design.

"Introduction to VLSI Systems Design" is specifically aimed at individuals with a hardware or software design background who are interested in applying their design or management skills in the VLSI domain. No previous IC design experience is assumed. In contrast to most VLSI seminars or tutorials, this course focuses on getting you involved in the design process. The course is not a survey of what others are doing, but rather teaches a perspective and a set of design skills which will allow you to make critical decisions about how best to include VLSI in your company's future plans. The perspective offered by the course will allow you to better plan and execute a VLSI design or design automation program within your company. Even if you do not intend to do VLSI design in-house, the knowledge gained in this course will allow you to make better decisions about outside integrated circuit design and fabrication vendors.

This video tape course was prepared for Hewlett-Packard, and is now being used by HP on a regular basis. The course is based on similar courses developed by Professor Carver Mead of Caltech and Lynn Conway of Xerox while teaching at M.I.T. The principle instructors for this course are Douglas Fairbairn, Director of VLSI Activities at VLSI Technology, Inc., and Professors John Newkirk and Robert Mathews of Stanford University. Guest lecturers include Professor Carver Mead of California Institute of Technology, Lynn Conway of Xerox, Professor Richard Newton of U.C. Berkeley, and Gunnar Wetlesen, Vice President of VLSI Technology, Inc. The course consists of 27 tapes containing 21 hours of lecture.

The course text is *Introduction to VLSI Systems* by Mead and Conway (Addison-Wesley, 1980). In addition to the text, VLSI Technology, Inc. offers this complete set of lecture notes that include all the visual aids used during the lectures, as well as a set of homework assignments. The course is greatly enhanced if the students have access to CAD tools and fast-turnaround chip fabrication. Any general-purpose computer can serve as a powerful VLSI design system. Appropriate software and plotting capabilities are the only special resources required.

Similar VLSI design courses based on the same text and design methodology are now being taught at over 100 different universities in the U.S. and abroad, including Stanford, Caltech, U.C. Berkeley, M.I.T., University of Illinois, and many others. In addition, the course has been taught for IBM, Xerox, Bell Labs, Hewlett-Packard, DEC, Varian Associates, Boeing Aerospace, and others.

VLSI Technology, Inc., was formed in 1979 to support the design and implementation of VLSI circuits. In addition to comprehensive educational courses and design tools, the company will offer electronic network access to a high-technology semiconductor fabrication facility built specifically to address the needs of the VLSI system designer. This facility will provide turnaround as short as three weeks from design input to finished chips.

For more information on how to overcome the VLSI education and design barrier in your company, complete and mail the postcard in the back of the book or write to VLSI Activities, D. Hungerford at VLSI Technology, P.O. Box 4788, Santa Clara, CA 95054-0788.

SYNOPSIS OF VIDEO TAPES

Lecture 1.1 — Overview

Douglas Fairbairn (20 mins.)

This lecture provides an overview of the topics to be covered in the course "Introduction to VLSI Systems Design." Specific topics include the introduction of the instructors, an outline of the course material, and a discussion of the mechanics of the course. The structured design methodology used throughout the course is introduced.

The purpose of the course is to teach software and hardware designers with no previous integrated circuit design background to architect, design, and layout fully-custom integrated circuits in the nMOS technology. The course is supported by the text *Introduction to VLSI Systems* by Carver Mead and Lynn Conway, and by a set of lecture notes and homework assignments. Homework and IC design projects are important parts of the instructional process.

Lecture 1.2 — VLSI Design: A Perspective

Professor Carver Mead (47 mins.)

In this introductory lecture Dr. Carver Mead first takes a look at the history of integrated circuits, and uses that history to establish a perspective on where the field is today. He reviews the scaling laws, which he helped develop in the early 1970's, and makes some projections of where the technology will eventually lead us. Dr. Mead points out, however, that in order to effectively use this technology we will have to solve the very real problems of designing complex chips.

Throughout the lecture, Dr. Mead develops the arguments which support the adoption of the basic design principles used in this course as the best solution to the VLSI design dilemma.

Lecture 1.3 Part 1 — nMOS Logic

Professor Robert Mathews (45 mins.)

This program, together with lecture 1.3 part 2, covers the basic notions of nMOS design: stick diagrams, switch logic and gate logic. Ratio rules are explained qualitatively, using a very simple transistor model. After this lecture, the student has all the information needed to begin designing logic in nMOS.

Lecture 1.3 Part 2 — nMOS Logic

Professor Robert Mathews (24 mins.)

This program, together with 1.3 part 1, covers the basic notions of nMOS design: stick diagrams, switch logic and gate logic. Ratio rules are explained qualitatively, using a very simple transistor model. After this lecture, the student has all the information needed to begin designing logic in nMOS.

Lecture 2.1 — Registers and Two-Phase Clocking

Professor John Newkirk (56 mins.)

This program introduces a simple, robust clocking and synchronization scheme: synchronous, two-phase, non-overlapping clocks with storage registers separated by combinational logic. Following a discussion of the properties and implications of this methodology, several simple register subsystems are developed. As a final exercise, these examples are used to construct a first-in, first-out stack.

Lecture 2.2 — PLA's and Finite State Machines

Professor Robert Mathews (40 mins.)

PLA's provide a way to mechanize sets of arbitrary Boolean equations in a geometrically and conceptually regular way. Moreover, with the addition of clocking, they become Finite State Machines. This lecture covers how to convert sets of equations into PLA's and how to describe and implement FSM's. A traffic light controller is used as an example.

Lecture 3.1 Part 1 — Subsystems and Floor Plans

Douglas Fairbairn (47 mins.)

The opening segment of this lecture describes the solutions to problems 2, 3 and 4 of the second homework assignment. The main body of the lecture introduces the student to some basic building blocks which are useful in nMOS design. Topics covered include precharging, function blocks, and the ALU cell used in the OM data path chip.

Lecture 3.1 Part 2 — Subsystems and Floor Plans

Douglas Fairbairn (38 mins.)

This lecture is a continuation of the lecture on subsystem building blocks and chip floor planning. The OM barrel shifter is described, followed by a discussion of floor planning and how to choose the appropriate layers for routing power, clocks, and signals. The lecture concludes with a careful description of the third homework assignment, which allows the student to apply the building blocks and the floor planning ideas discussed in this lecture.

Lecture 3.2 — nMOS Processing and Design Rules

Professor John Newkirk (52 mins.)

The objective of this lecture is to define a set of design rules: geometric constraints on line width, line separation, etc., that characterize the fabrication process from the designer's perspective. A brief overview of the fabrication process is presented, followed by a more detailed development of the patterning process and of the patterning sequence used in fabricating nMOS devices. Several physical constraints arising from this process are briefly examined and used to develop and justify a set of design rules.

Lecture 4.1 — Chip Planning and Layout

Professor Robert Mathews (44 mins.)

In this lecture, case studies are presented exploring the planning and layout of several subsystems, including a shift-register array and a memory based on 3-transistor cells. Among the topics discussed are how to assign layers for different purposes, how to count ratios in gates, and how to apply the nMOS design rules. The latter is illustrated by walking through a layout.

Lecture 4.2 — Timing, Delays and Power Dissipation

Professor Robert Mathews (50 mins.)

Adequate power bussing is essential if a design is to work. Delay calculations are important if a design is to meet speed specifications. This lecture covers resistance and current calculations, metal-migration limits, capacitance calculations, and delay estimates. Calculations are expressed in normalized units wherever possible.

Lecture 4.3 — nMOS Scaling and Its Design and Process Implications

Douglas Fairbairn, Gunnar Wetlesen (57 mins.)

This lecture is divided into two segments. The first is done by Douglas Fairbairn and describes the details of scaling nMOS circuits and the benefits that accrue from this scaling. Scaling of nMOS circuits calls for reducing all the dimensions of the transistors as well as other physical parameters such as the doping densities. The problems associated with the scaling of interconnect as well as transistors are also covered. The first segment of this lecture generally treats scaling from an ideal point of view.

The second segment is presented by Gunnar Wetlesen and describes the differences between ideal scaling principles and those actually applied on the production line. Problems in scaling buried contacts and other structures are also discussed.

Lecture 5.1 — Memory Subsystems

Douglas Fairbairn (54 mins.)

This lecture introduces the student to the alternative and trade-offs for designing memory cells and memory subsystems. Alternative memory-cell designs are compared in terms of area, power, and design complexity. Examples of the design and use of "smart memories," such as stack cells, are also given.

Lecture 5.2 — Project Introduction

Douglas Fairbairn (54 mins.)

It is assumed that students taking the VLSI design course will have the opportunity to design their own integrated subsystem during the class. This lecture describes what size and complexity of project is appropriate and some of the pre-defined cells which should be available to use in class projects. These cells include input and output pads, PLA cells, and shift register cells. Other topics include hints on chip size estimation, design pitfalls, testing strategies, and a reasonable design schedule. Wafer dicing and packaging are also touched on.

Lecture 6.1 Part 1 — Signal Processing with VLSI

Richard F. Lyon (37 mins.)

Dick Lyon provides a very short introduction to the subject of digital signal processing. He then discusses the reasons behind his choice of a special-purpose, bit-serial architecture for the audio-rate, signal-processing system he has developed. Mr Lyon has adopted a modularized approach which allows many different building blocks to be connected together easily to perform special functions. Some of the basic building blocks include serial multipliers, delay (memory) elements, and adder/subtractors. The lecture details the operation of the elements as well as the overall timing strategy which was adopted to ensure that all building blocks would communicate reliably. Examples of VLSI signal-processing systems which were built with these building blocks are described.

Lecture 6.1 Part 2 — Signal Processing with VLSI

Richard F. Lyon (40 mins.)

This tape is the second part of a two-part lecture by Dick Lyon on digital signal processing with custom VLSI chips. Mr. Lyon has adopted a modularized approach for this signal-processing system which allows many different building blocks to be connected together easily to perform special functions. Some of the basic building blocks include serial multipliers, delay (memory) elements, and adder/subtractors. The lecture details the operation of the elements as well as the overall timing strategy which was adopted to ensure that all building blocks would communicate reliably. Examples of VLSI signal-processing systems which were built with these building blocks are described.

Lecture 6.2 — What is Silicon Compilation?

Dave Johannsen (30 mins.)

This lecture describes the concept of a silicon compiler as well as the first example of such a program. This compiler, known as Bristle Blocks, was developed by the lecturer at Caltech as a first cut at the solution to the VLSI design problem. Key concepts in the Bristle Blocks program include regularized architecture, highly "parameterized cells" which, when compiled, can customize themselves to the application, and multiple representations of the same function (logic drawings, layout, simulation code). These multiple representations enable design, simulation, and documentation of the chip to be done in a hierarchical fashion, and they help guarantee the correctness of the chip by construction rather than post-design analysis.

Lecture 7.2 — Correctness by Construction

Professor Carver Mead (25 mins.)

One of the major differences between LSI design and VLSI design is its sheer complexity. In this tape, Carver Mead advocates the adoption of very carefully-defined disciplines of design which, when followed, help guarantee the correctness of a design. A simple example of such a discipline is the two-phase clocking scheme presented in this course. The price for ignoring these design disciplines in complex circuits is high: the circuits may never be made to work because new problems are constantly found faster than old ones are corrected.

Lecture 7.3 — Understanding Hierarchical Design

Dr. James A Rowson (51 mins.)

Hierarchical design is an often-used term that is commonly misunderstood. In this lecture, Dr. Rowson clarifies the meaning of the term and describes the key characteristics that a hierarchical design system ought to have. He offers a way of mathematically proving the equivalence of different hierarchical descriptions, and then describes some basic composition rules which can help guarantee the correctness of a hierarchical design.

Lecture 7.4 — Introduction to Simulation for VLSI

Professor Richard Newton (58 mins.)

This is an excellent overview of the LSI/VLSI simulation problem. Dr. Newton defines simulation, describes when it is appropriate to simulate, at what level, and also summarizes the cost of simulation at the circuit, logic, and functional levels. He describes the advantages and problems of various simulators such as SPICE, SPLICE, MOTISC, and SALOGS IV, and gives examples of their application. This lecture can stand alone as a general introduction to the topic or it can be used in conjunction with the other lectures on design automation (6.2, 7.2, 7.3).

Lecture 8.1 — Review of Scaling: Synchronous Systems

Professor Charles L. Seitz (58 mins.)

This is the first in a two-lecture series by Dr. Seitz on system timing issues related to VLSI design. The two lectures are related, but not so closely that they can't be viewed independently. In this first lecture, Dr. Seitz describes the concept of a synchronous digital system as well as the background and justification for the two-phase clocking system described in the course text (*Introduction to VLSI Systems* by Mead and Conway). The significant problem of clock skew in VLSI chips and the effect further scaling of chip dimensions will have on the problem are also discussed. The lecture concludes with a description of, and warning about, the synchronization problem which exists between any two digital subsystems with different clocks. Lecture 8.2 presents a self-timed approach to solving the problems discussed in this lecture.

Lecture 8.2 — Self-Timed Systems

Professor Charles L. Seitz (57 mins.)

This second lecture on system timing for VLSI describes the concept of self-timed systems. Such systems are made up of modules which operate at their own inherent speed and use special synchronizing hand-shake signals to indicate when any particular module is done computing and has a valid result available. This timing methodology is an alternative to globally-synchronous systems and is a potential solution to the clock distribution and synchronization problems described in lecture 8.1.

Lecture 8.3 — A VLSI Geometry Subsystem for Computer Graphics

Professor James Clark (56 mins.)

Dr. Clark describes the architecture and implementation of a VLSI system which performs the transformation, clipping, and scaling functions required in high-performance, 3-dimensional graphics display systems. He tailored the architecture of this system specifically for VLSI. The architecture chosen will allow the system to make use of future advances in chip scaling with no redesign required. This architecture also allows one chip to be used in different configurations to perform all the above functions in a pipelined system. This geometry subsystem uses many of the timing concepts discussed in lectures 8.1 and 8.2 and should be viewed in conjunction with those lectures.

Lecture 9.1 Part 1 — The SCHEME-79 Architecture and Design Methodology

Jack Holloway (55 mins.)

This lecture describes the architecture and design methodology used by Jack Holloway and other researchers at MIT on the SCHEME-79 LISP microprocessor chip. This chip directly executes a dialect of LISP known as SCHEME. In addition to this interpreter, the chip contains an automatic storage allocator, an integral garbage collector, and an interrupt system for interrupt handlers written in SCHEME. The design tools used in the project were largely LISP-based and include a symbolic, procedural layout language embedded in LISP, a microlanguage for specifying microcode in the form of a high-level language, and a silicon compiler for implementing a particular class of architectures. This two-part lecture must be viewed as a whole.

Lecture 9.1 Part 2 — The SCHEME-79 Architecture and Design Methodology

Jack Holloway (43 mins.)

This is the second part of a two-part lecture describing the architecture and design methodology used by Jack Holloway and other researchers at MIT on the SCHEME-79 LISP microprocessor chip. This chip directly executes a dialect of LISP known as SCHEME. In addition to this interpreter, the chip contains an automatic storage allocator, an integral garbage collector, and an interrupt system for interrupt handlers written in SCHEME. The design tools used in the project were largely LISP-based and include a symbolic, procedural layout language embedded in LISP, a microlanguage for specifying microcode in the form of a high-level language, and a silicon compiler for implementing a particular class of architectures. This two-part lecture must be viewed as a whole.

Lecture 10.1 Part 1 — New Methods for Implementing VLSI Designs

Lynn Conway (39 mins.)

Lynn Conway, coauthor of the course text, *Introduction to VLSI Systems*, describes the issues and problems associated with implementing VLSI systems. She first reviews such items as the role of the Caltech Intermediate Form (CIF), the starting frame, optical pattern generation vs. electron-beam masks, and packaging. This discussion is followed by a detailed description of the Multi-Project Chip (MPC) operating system developed at Xerox and the crucial role of such a system in making VLSI fabrication available to a large community of designers.

This two-part lecture should be viewed as a whole. This lecture may be shown in conjunction with the Project Introduction lecture (5.2) or at the end of the course as originally scheduled.

Lecture 10.1 Part 2 — New Methods for Implementing VLSI Designs

Lynn Conway (55 mins.)

This is the second part of a two-part lecture by Lynn Conway in which she describes the issues and problems associated with implementing VLSI systems. She first reviews such items as the role of the Caltech Intermediate Form (CIF), the starting frame, optical pattern generation vs. electron-beam masks, and packaging. This discussion is followed by a detailed description of the Multi-Project Chip (MPC) operating system developed at Xerox and the crucial role of such a system in making VLSI fabrication available to a large community of designers.

This two-part lecture should be viewed as a whole. This lecture may be shown in conjunction with the Project Introduction lecture (5.2) or at the end of the course as originally scheduled.

TABLE OF CONTENTS

Section 1	Lectures	1.1 Overview and Course Goal Douglas Fairbairn, VLSI Technology, Inc.	1
		1.2 VLSI Design: A Perspective Professor Carver Mead, Caltech	
		1.3 nMOS Logic Professor Robert Mathews, Stanford	
Section 2	Lectures	2.1 Registers and 2-Phase Clocking Professor John Newkirk, Stanford	2
		2.2 PLAs and Finite State Machines Professor Robert Mathews, Stanford	
Section 3	Lectures	3.1 Subsystems and Floor Plans Douglas Fairbairn, VLSI Technology, Inc.	3
		3.2 nMOS Processing and Design Rules Professor John Newkirk, Stanford	
Section 4	Lectures	4.1 Chip Planning and Layout Professor Robert Mathews, Stanford	4
		4.2 Timing, Delays and Power Dissipation Professor Robert Mathews, Stanford	
		4.3 nMOS Scaling and Its Design and Process Implications Douglas Fairbairn and Gunnar Wetlesen, VLSI Technology, Inc.	5
Section 5	Lectures	5.1 Memory Systems Douglas Fairbairn, VLSI Technology, Inc.	6
		5.2 Project Introduction Douglas Fairbairn, VLSI Technology, Inc.	
Section 6	Lectures	6.1 Signal Processing with VLSI Richard F. Lyon — Xerox Palo Alto Research Center	7
		6.2 What is Silicon Compilation Dave Johannsen, Caltech	
Section 7	Lectures	7.2 Correctness by Construction Professor Carver Mead, Caltech	8
		7.3 Understanding Hierarchical Design Dr. James Rowson, VLSI Technology, Inc.	
		7.4 Introduction to Simulation for VLSI Professor Richard Newton, U.C. Berkeley	
Section 8	Lectures	8.1 Review of Scaling: Synchronous Systems Professor Charles L. Seitz, Caltech	9
		8.2 Self-Timed Systems Professor Charles L. Seitz, Caltech	
		8.3 A VLSI Geometry Subsystem for Computer Graphics Professor James Clark, Stanford	10
Section 9	Lecture	9.1 The SCHEME-79 Architecture of Design Methodology Professor Jack Holloway, M.I.T.	11
Section 10	Lecture	10.1 New Methods for Implementing VLSI Designs Lynn Conway, Xerox Palo Alto Research Center	12
Section 11		Homework	
Section 12		Color Plates	