INTRODUCTION TO MOS LSI DESIGN

J. Mavor, M. A. Jack and P. B. Denyer

Introduction to MOS LSI DESIGN

J. MAVOR, M.A. JACK, P.B. DENYER

University of Edinburgh

ADDISON-WESLEY PUBLISHING COMPANY

London · Reading, Massachusetts · Menlo Park, California · Amsterdam
Don Mills, Ontario · Manila · Singapore · Sydney · Tokyo

The authors dedicate this book to their long-suffering wives Sue, Marion and Fiona

Phototypeset by Tradespools Limited, Frome, Somerset.

Printed and bound in Great Britain by Collins, Glasgow.

Library of Congress Cataloging in Publication Data

Mavor, J.
Introduction to MOS LSI design.

(Microelectronics systems design series)
Bibliography: p.
Includes index.
1. Integrated circuits—Large scale integration.
2. Metal oxide semiconductors. I. Jack, Mervyn, 1949– . II. Denyer, Peter 1953– .
III. Title. IV. Title: Introduction to M.O.S. L.S.I. design. V. Series.
TK7874.M345 1983 621.381'73 82–8866
ISBN 0–201–14402–6

British Library Cataloguing in Publication Data

Mavor, John
Introduction to MOS LSI design.—(Microelectronics systems design series)
1. Integrated circuits 2. Metal oxide semiconductors
I. Title II. Jack, Mervyn
III. Denyer, Peter IV. Series
621.381'73 TK7871.85

ISBN 0–201–14402–6

ABCDEF 89876543

Preface

Integrated circuit technology has matured to the extent that silicon integrated circuit 'components' now form the key to efficient electronic systems implementation. It therefore follows that an increasing number of systems engineers need to be engaged on integrated circuit design for a vast range of emerging applications. These engineers need to possess a critical awareness of silicon technology to appreciate the potential and the limitations of integrated circuit design. Unfortunately, there is a serious shortage of engineers with such experience and although many university courses in both Computer Science and Electrical Engineering Departments are now being structured to satisfy the demand from industry for graduates with expertise in microelectronics, the shortage will exist for the foreseeable future.

This book has been written as a basic text for such courses as presented primarily to undergraduate students of electrical engineering and has been biased to give particular emphasis to the engineering design of integrated circuits, thereby complementing the overall systems design approach advocated by Mead and Conway in their text *Introduction to VLSI Systems**.

The contents of this book have been generated from a course in MOS LSI Circuit Design which has been developed in the Department of Electrical Engineering at Edinburgh University over a period of ten years. The course has been successfully run as part of the B.Sc. (Honours) Degree in Electrical Engineering and has also been presented to industry in the form of an intensive lecture series. More recently, the course has been offered as part of an M.Sc. Degree in the Department of Electrical Engineering. The material concentrates on the engineering of silicon-gate circuits in both NMOS and CMOS since these are recognised as technologies *with* a future—if not the technologies *of* the future—as device sizes reduce and VLSI circuits develop.

Starting with an overview of the LSI design route comparing design approaches, technologies and capabilities, MOS transistor fundamentals are introduced and their (first order) characteristics are examined. From this consideration of the basic MOS transistor, the treatment expands to consider the use of such devices in logic gates in terms of static and transient analysis for ratioed designs and ratio-less circuits. The need for simplification of the

*Mead, C. and Conway, L. 1980. *Introduction to VLSI Systems*. Reading, Mass., USA: Addison-Wesley Publishing Company.

wide range of design options is recognised and details of transportable design standards based on normalised feature sizes and transistor geometries are included.

Having established the basic design requirements of logic gates, the material develops into a consideration of the basic building blocks of LSI design—flip-flops, shift registers, arrays, counters and adders—emphasising design/performance tradeoffs for the various methods. The treatment of basic blocks continues with consideration of dynamic logic counter arrangements, ROM design and RAM design. Semi-custom methods for LSI design, such as the uncommitted logic array (or gate array) and functional cell library approaches, are included to create a perspective of where, and when, custom LSI fits a product definition or performance specification.

The requirement to actually test a finished chip is a feature which must be considered in the design from the outset, as opposed to reliance on test routines effected after the design has been completed. The philosophy of design for testability is covered to create an awareness of design testability requirements and several basic methods of design for testability are considered.

Computer-aided design (CAD) plays an all-important part in LSI design and, indeed, advanced CAD will prove crucial for development of higher complexity (VLSI) circuits. In recognition of these factors, Chapter 7 has been devoted to giving the reader an appreciation of the role of CAD techniques in circuit design.

MOS technology has a definite place in future integrated circuit developments, and the basic principles of MOS scaling to VLSI as device feature sizes shrink to micron and even sub-micron dimensions are included to emphasise the potential advantages and possible problem areas. Scaling to VLSI offers the possibility of integrating complete systems with very high performance digital operation. However, such systems require to interface with a real (analogue) world so the need for analogue circuitry (possibly not small geometry) associated with the VLSI architectures on a common chip will develop. Based on these considerations, a detailed treatment of analogue MOS design is included to cover switches, capacitors and operational amplifiers leading to switched-capacitor filter methods and coverage of the all-important analogue-to-digital and digital-to-analogue data converters.

In our experience it is vital for students new to the subject to perform a number of basic calculations relating to MOS transistors and simple circuits. A selection of tutorial examples is included in Appendix B for this purpose, with answers given as appropriate. The authors are prepared to generate full solutions if a suitable demand for them develops. In addition to these exercises, a simple worked example is included in Appendix A to introduce the reader to the procedures necessary in formulating an integrated design. A low complexity, 4-bit binary counter example has been chosen in order that every individual design aspect can be highlighted and illustrated. It is hoped that, taken together, both appendices will prove a useful feature of the book, especially for undergraduate students.

Although this text is focused on the engineering design of LSI density integrated circuits, the basic methods introduced can be extended as a foundation for VLSI circuits where detailed engineering design will continue to play a vital part.

The authors wish to acknowledge a considerable number of colleagues and students who, over the years, have contributed to this book. Firstly, we wish to thank Caroline Burns who carefully typed the text, and David Stewart-Robinson who provided an invaluable service in preparing the artwork to a high standard. Secondly, we wish to thank Brian Flynn, Ross Mactaggart and Neil Henderson who undertook a number of circuit design tasks in connection with the book, and also David Renshaw who meticulously checked the manuscript. Thirdly, our thanks go to the following industrialists who, over a number of years, helped us to develop a course in MOS LSI Design and encouraged us at many levels and in various ways to formulate the preparation of this text: Mr J Dickson, Plessey Research (Caswell) Ltd; Mr A M G Gundlach, University of Edinburgh; Mr S Kelly, GTE Microcircuits; Dr G Vanstone, Racal Microelectronic Systems Ltd; Dr A D Milne, Wolfson Microelectronics Institute; Mr K Murray, Hughes Microelectronics Ltd; Mr R Thomson, Wolfson Microelectronics Institute; and Dr C Walmsley, Denyer Walmsley Microelectronics Ltd. Finally, we acknowledge the University of Edinburgh for their early foresight in encouraging our teaching and research in microelectronics and providing conditions conducive to the efficient preparation of this book. The support of Professor J H Collins, Head of the Electrical Engineering Department at the University of Edinburgh, is also gratefully acknowledged.

University of Edinburgh
1982

John Mavor
Mervyn Jack
Peter Denyer

Contents

MICROELECTRONICS SYSTEMS DESIGN SERIES

Consulting editor **G. Musgrave**
Brunel University

1

The LSI design route

1.1 Introduction to integrated circuit technology

The cost, time investment and technical effort required to engineer an integrated circuit (IC) is currently a major undertaking: it may consume many man-years of design effort alone and a financial investment of around $250000. The functional complexity of the modern integrated circuit, which by the mid-1980's will exceed one million transistors, is now so great that this 'component' can contain complete systems. Indeed, recently reported 32-bit microprocessor circuits are complete computing systems on a single silicon integrated circuit.

Until lately, it was commonplace to design and fabricate an integrated circuit with perhaps up to a few thousand gates of 'random' logic. This relative level of integration or functional complexity is widely referred to as large scale integration (LSI), and such circuit densities were developed during the late 1970's and early 1980's. The integrated circuit industry has indeed matured since the late 1950's when the revolutionary *planar* technology concept was introduced, which first enabled an electronic circuit in *monolithic* form to be produced. At that time, only low functional densities at so-called small scale integration (SSI) were achievable and, typically, several gates were integrated. However, as the technology developed during the late 1960's, the era of medium scale integration (MSI) appeared where several hundred components per integrated circuit could be fabricated. Already in the 1980's prototype circuits exceeding one million components are in design and the era of very large scale integration (VLSI) has arrived. However, the design potential of LSI has not yet been fully exploited for the myriad of possible applications, owing to a world shortage of LSI designers. Further, at the VLSI complexity level the design problem itself will escalate, calling for new, desperately needed and still awaited computer-aided design (CAD) tools.

The purpose of this book is to introduce the reader to all aspects pertinent to the engineering design of LSI-density integrated circuits or 'chips', as they are often called. Because of the present acute shortage of skilled integrated circuit designers, the text has been written from the point of view of an intending designer. Although several books are already available

which treat general aspects of microelectronics technology and devices (Grove, 1967; Cobbold, 1970; Sze, 1969; and Glaser and Subak-Sharpe, 1977), this book is intended to be a comprehensive engineering text on LSI design aimed primarily at engineering students or systems engineers with no previous integrated circuit design experience.

Further, this book is focused exclusively on metal-oxide-semiconductor (MOS) technology where the semiconductor substrate material is single-crystal *silicon*, typically available as wafers of diameter 75 mm or 125 mm. The alternative, bipolar junction transistor (BJT) technology is still widely used and will continue to be applied particularly for high speed integrated circuit applications. Although MOS transistors and bipolar transistors can be fabricated on the same wafer for applications which require the features of both technologies, this is rarely done because of the increased processing complexity. Broadly, for the highest speed digital applications or low noise amplifier design BJT technology is preferred, whereas for high complexity and medium speed design MOS technology is favoured. The attraction of MOS technology stems from the fact that its transistor structure is inherently much simpler than the BJT structure. This is certainly true for MOS transistors with dimensions of 5 microns (or micrometres, 10^{-6} m) typical of LSI; however, for MOS devices at feature sizes of 1–3 microns (VLSI) the basic model of the MOS transistor becomes very complex and comparable with the BJT. Further, the modern MOS fabrication process, involving perhaps one or two interconnection levels (see Chapter 3) is as complex as a typical bipolar process. In spite of this undesirable trend, the semiconductor engineer has many unique device structures within the generic MOS family to call upon. For example, using charge-coupled device (CCD) technology, solid-state imagers and high density transversal filters can be designed; switched-capacitor (SC) filters can be designed in MOS technology for frequency-selective filtering requirements; VMOS technology can be used for switching high voltages; floating-gate avalanche (FAMOS) technology is available for implementing non-volatile memory, etc. It will be shown later that, from a circuit standpoint, the fundamental reason why MOS technology is preferred over BJT technology for many applications is because the MOS transistor incorporates, inherently, a near perfect input capacitor which permits input signals to be stored. Thus, the design of high performance dynamic circuitry is routine in MOS. This is in marked difference to bipolar designs using dynamic circuit principles where low impedance levels require that more complex circuit structures be employed. Because a variety of different structures can be designed in MOS technology, novel integrated circuits can be realised which, for example, may combine an imaging section coupled with focal-plane signal processing and associated memory, all tailored or custom designed for the particular requirement. Such flexibility is not as evident in bipolar technology, but bipolar technology is complementary to MOS technology in many applications areas. In Table 1.1 a broad comparison is made of the features of digital logic technology for two bipolar families (so-called transistor transistor logic (TTL) and integrated injection

logic (I^2L)), and two MOS technologies (NMOS which employs only *n*-channel devices, and CMOS which has complementary *n*-channel and *p*-channel devices). The comparison is made for the electrical properties of the basic logic structure referred to as the *logic gate*, and also for technological features such as the number of wafer masking steps required and the number of diffusions and implants employed in their fabrication sequence. It is evident from Table 1.1 that, in the main, bipolar circuits offer a speed advantage over MOS circuits, but at the expense of layout area per gate and power. In comparison, MOS technology has a potential for low power, high circuit density applications. Table 1.1 also reflects the fact that, currently, the processing complexity of either technology is broadly comparable. The manufacturing cost per 75 mm silicon wafer is typically about $100 irrespective of the process type, although for a more primitive process, such as metal-gate NMOS, the cost will be slightly cheaper and a more sophisticated process, such as silicon-gate CMOS, will be more expensive. However, as the cost per wafer is relatively constant at any time, the task of the integrated circuit designer is to maximise the density of logic functions per unit silicon area, and the fabrication engineer thereafter will try to minimise the number of defective circuits and thereby optimise the device yield.

Table 1.1 Digital technology comparison

	TTL	NMOS	CMOS	I^2L
Typical area/gate (μm^2)	27 000–36 000	1000	1500	5000
Propagation delay/gate (ns)	3–10	15–60	10–40	15
Static power/gate (mW)	2–19	0.2–0.5	>0.001	0.13
Speed-power product (pJ)	18–100	1–5	0.5–3	2
Masking steps	7–8	6	7	7
Diffusions and implants	5	3	4–6	4

The yield is determined by the number of random defects occurring per unit area on the silicon wafer (usually expressed as the number per cm^2), and is given as the percentage of perfectly functioning circuits to the number of potential candidates on a given wafer. For example, a 75 mm diameter wafer having 100 potentially working chips of 25 mm^2 area, fabricated in MOS technology, may exhibit a yield of 30%, or about 30 working circuits. As bipolar processes are inherently more complicated than MOS, and BJT device parameters are more critically dependent on processing, it is not surprising that the corresponding yield in bipolar technology may be nearer to 10% for the above example.

The presence of defects causes an integrated circuit to malfunction, mainly due to a number of unwanted open and short circuits. These are caused primarily by pinholes in gate oxides and excess reverse leakage currents in *pn* junction diodes from which the transistors are formed. In addition, breaks in metallisation tracks and defective contacts between

conducting layers will also drastically impair circuit performance. The main origins of the defects may be classed as:

(a) Poor starting materials, and contaminated chemicals and gases used in wafer fabrication. Of particular concern here are irregularities in the otherwise crystalline wafer which cause conducting 'spikes' in *pn* junctions. Contamination of furnaces by phosphorus can also adversely affect device parameters.

(b) Poor processing of the wafers. Defective contacts can be caused by incorrect etching, and misalignment of masks during processing can be responsible for a host of problems.

(c) Poor environmental control of processing area. The performance of organic resists used to define device structures is dependent upon the humidity and temperature of the processing area, as are most other fabrication steps.

To enable the process yield to be monitored concisely on a daily basis, it is mandatory to fabricate basic test structures with each production run, or batch, of 30 to 50 wafers. Usually, this procedure takes the form of several 'drop in' circuits per wafer, each of which contain standard transistor formations, contact tests, conductor paths to evaluate resistivity, etc. The process engineer can measure these test circuits routinely on a test computer before the production batch is released from the manufacturing area. Of particular interest is the statistical variance of the processing from the target specification, as well as the stability and reliability of the structures made on it. Given a stable process, new circuit designs can be evaluated with confidence and the performance of prototype circuits will be typical of later production.

The integrated circuit industry thrives on a high wafer throughput, and to obtain high yield consistently over more than 50 individual steps involved in the manufacturing process requires discipline and clean environmental conditions. A fully-equipped semiconductor fabrication plant costs about $10–15 million and has a daily wafer throughput of several thousand wafers. It is only by maintaining such a high throughput rate of wafers that the unit costs can be reduced, the enormous capital costs can be recovered and the considerable design overheads amortised over the production runs.

Silicon technology now exists as a pervasive, base technology for the implementation of BJT and MOS circuits for the following reasons:

(a) The silicon device processing industry has now been established for more than 20 years.

(b) Device yields are relatively high, thus integrated circuits are cheap to produce.

(c) The physics of silicon *pn* junctions and transistor structures are well researched.

(d) Raw materials such as silicon and aluminium are relatively inexpensive and available at high purity.

(e) Industry 'standard' processes are available throughout the world, permitting 'second sourcing' of production arrangements.

(f) A wide variety of devices and circuit techniques are available to the designer.

(g) Design rules are relatively 'standard'.

These key factors are certain to ensure that silicon technology will continue to be applied well into the next decade because of the maturity of the industry.

1.2 LSI design options

The advantages afforded by LSI technology can only be achieved when the majority of a system's electronics is integrated within a small number of integrated circuits. Printed circuit boards containing large numbers of integrated circuits and other discrete components such as diodes and capacitors are costly to assemble, and so the modern trend is to integrate all the electronics monolithically such that the power, space and assembly cost reductions can be optimised and the full impact of integrated circuit technology can be achieved.

The cheapest solution to a systems integration problem occurs where the system can be formed from several standard or 'catalogue' circuits. This is a very unusual but fortunate situation, as there are no integrated circuit design costs and the component costs are at a minimum because standard parts usually attract high volume sales and, therefore, can be sold by the manufacturer at a low profit margin. Apart from this case, the integration of a more general system can be realised via one of three main LSI design options:

(a) Design involvement at the IC interconnection level—here referred to as customised *interconnect*.

(b) Design involvement at the software level—here referred to as customised *program*.

(c) Design involvement at LSI circuit design level—here referred to as customised *circuit*.

1.2.1 Customised interconnect

Customised interconnect, is the term normally used to define the external connection of a number of integrated circuits by specially routed wires or printed circuit board tracks. However, the term is used here to infer customised interconnection tracks within the silicon integrated circuit itself, around standard and proven cells which have previously been made at high yield. The risk in this approach to systems integration is extremely low because computer aids may be employed to route the tracks automatically, and the cost and development time can also be relatively low. Indeed, companies which offer such a service normally encourage the customer to

undertake the design himself, and the time taken to obtain working circuits from a logic diagram may be only 10 weeks. Also, the systems engineer essentially takes full responsibility for the design work (and usually the testing as well), and so this approach, when applicable, is particularly attractive to both systems companies and semiconductor companies.

There are basically two distinctly different techniques for producing a custom-interconnected integrated circuit: (a) the standard cell approach; and (b) the gate array or uncommitted logic array approach. In method (a) a full range of standard cells, including a range of combinational and sequential logic functions plus peripheral pads for power supplies and input/output signals, are stored in a computer database. According to the design requirement, these cells can be called up on a computer terminal and displayed on a visual display unit (VDU). The cells can then be interconnected semi-automatically to a standard 'grid' of conductor tracks. With this approach, the silicon wafer is regarded essentially as a 'silicon printed circuit board', and the cells are equivalent to off-the-shelf TTL or CMOS standard circuits. The design cycle is inherently extremely rapid with this approach and a complete custom integrated circuit can be designed in a matter of weeks and then fitted within a standard pad frame (which enables input/output signals to be supplied via wire bonds at assembly). The design then proceeds to be fabricated in the normal manner as for any integrated circuit.

The alternative technique (b), based on gate arrays, is also a rapid design approach where the cell structures are more primitive and exist either as groups of 4–6 transistors or simple combinational logic gates in a regular two-dimensional array. The wiring procedures involved with gate arrays are clearly more complicated than with standard cell approaches because the functional cells themselves have to be constructed from a number of array cells. The gate array approach, therefore, relies heavily on computer-aided layout for interconnect routing and many such programs are commercially available or in development. Although the gate array approach may appear less desirable than the standard cell approach, gate array design is very flexible and functional blocks of any description can be composed from the primitive cells. The main significance of the gate array lies in the fact that only one step in processing needs to be undertaken for customising. Thus the gate array approach is low cost and the design cycle time is extremely short, which makes it attractive to systems engineers. This approach is treated in more detail in Chapter 6 where design examples are described.

1.2.2 Customised program

Currently, there is some debate as to which of the following two fundamentally different solutions to solving an LSI design problem is more efficient:

(a) Having a *general purpose* computational integrated circuit and designing its operational software *after* manufacture (the prime example being a microprocessor).

(b) Having a *totally dedicated*, custom-designed integrated circuit and designing its internal details *before* manufacture.

Clearly, the microprocessor approach is more flexible since it can be reprogrammed over and over again for a limitless variety of tasks. However, the sacrifice to be paid for a general purpose machine is that part of its computing power is concerned with controlling itself from its associated software and, in any event, only a fraction of the remainder of its power will in general be used to perform any task. Thus microprocessor-based systems have at present relatively poor performance, but balanced against this is the distinct advantage that their operating function may be changed easily and quickly via the controlling software. However, a further concern is involved with the security and testing of software programs, which, arguably, can be corrupted more easily and which are more volatile than a dedicated integrated circuit that has been proved to be fully functional after manufacture.

Microprocessor circuits are in a state of continual development and, recently, particularly useful 32-bit (digital word) designs have been introduced. Undoubtedly, in the future, more powerful circuits with more efficient and standardised, high level operating software will become available ensuring that this design approach will remain in widespread use, particularly for equipment controllers and in prototype systems design where flexibility is crucial.

1.2.3 Customised circuits

The most efficient way of realising any electronic function in integrated form is to design a dedicated, custom-designed circuit. Until recently, however, the effort and finance required for this task has not been justified unless the requirement for the circuit is predicted to exceed about 50000 parts per year or unless high performance is demanded. Only in these cases would the cost per function be acceptable when amortised over the production quantity. Relatively recently, there have been moves to address this problem; moves which include the setting up of 'silicon foundries' which could potentially attract an overall high throughput rate composed of a large number of low volume jobs. When such concepts are more widely accepted, significantly more custom-designed circuits may be produced for the many applications where the potential volume is modest, e.g. medical electronics and military systems.

A comparison of the relative merits of the three LSI design options for integrating logic, as a function of device volume and technology is presented in Table 1.2. It is reasonable to deduce from this table that for applications requiring the highest speed of operation, bipolar technology is a sure solution. However, as developments in CMOS fabricated on low capacitance, silicon on sapphire (CMOS-SOS) wafers and isolated-dielectric (ISO) CMOS technology progress, production processes which enable circuits to operate at

Table 1.2 The technology/design choice as a function of quantity requirement and performance (after Roberts, 1980)

	Requirement / *Technology & design approach*	*High speed*	*High packing density*	*Low power*
	Technology	*bipolar/* CMOS-SOS	NMOS I^2L	CMOS I^2L
Quantity of integrated circuits	*Technology* / *Design approach*			
1000	gate array	bipolar	ISO-CMOS I^2L	ISO-CMOS
10000	cellular	—	NMOS	CMOS
100000	fully customised	bipolar (lower density)	NMOS ISO-CMOS	ISO-CMOS I^2L
1000000	standard product	bipolar (lower density)	NMOS	ISO-CMOS NMOS I^2L

clock rates greater than 100 MHz will become available. Also for GHz clocking rates, gallium arsenide (GaAs) technology (not given in Table 1.2) may offer an attractive future alternative. Broadly, for all other applications requiring low power dissipation and a high circuit packing density, circuits produced in MOS technology are preferred.

With regard to the LSI design approach, Table 1.2 indicates that for small volume applications, perhaps a total demand of 1–5 thousand circuits, gate array techniques offer a rapid solution at minimum cost. Between this requirement of several thousand circuits up to a level of about 100000 parts a standard, cellular approach is suitable. Equivalent circuit densities of 2–3 thousand gates are typical and include a vast range of printed circuit board replacement applications. As discussed earlier in this chapter, because with this approach a new circuit design is engineered every time—albeit with a rapid design method—the fabrication timescale is at least several weeks, unlike the gate array approach where a turn around time of several days is not unknown. For any potential LSI development where 100000 parts are to be required then a full custom design is warranted, as the initial, substantial design costs can be reduced to a modest figure per device over this volume. This is ultimately the most efficient approach as the circuit 'hardware' is designed specifically for the function to be performed. However, the design cannot be modified easily after manufacture. For this reason, such a design exercise is nowadays often preceded by a microprocessor 'model' or, more rarely, by a 'breadboard model'. By definition, the microprocessor is software programmable and enables at least some of the circuit to be emulated before a custom LSI design is commenced. Finally, as Table 1.2 indicates, any custom-designed circuit which reaches a production level of one million parts, and is widely available, effectively becomes a standard circuit which is almost